The Mother of All Leakages: How to Simulate Noisy Leakages via Bounded Leakage (Almost) for Free

Gianluca Brian[1], Antonio Faonio[2], Maciej Obremski[3], João Ribeiro[4], Mark Simkin[5], Maciej Skórski[6], and Daniele Venturi[1]

[1] Sapienza University of Rome, Rome, Italy.
brian@di.uniroma1.it, venturi@di.uniroma1.it
[2] EURECOM, Sophia-Antipolis, France.
antonio.faonio@eurecom.fr
[3] National University of Singapore, Singapore, Singapore.
obremski.math@gmail.com
[4] Imperial College London, London, UK.
j.lourenco-ribeiro17@imperial.ac.uk
[5] Aarhus University, Aarhus, Denmark.
simkin@cs.au.dk
[6] University of Luxembourg, Luxembourg, Luxembourg.
maciej.skorski@uni.lu

Abstract. We show that the most common flavors of noisy leakage can be simulated in the information-theoretic setting using a single query of bounded leakage, up to a small statistical simulation error and a slight loss in the leakage parameter. The latter holds true in particular for one of the most used noisy-leakage models, where the noisiness is measured using the conditional average min-entropy (Naor and Segev, CRYPTO'09 and SICOMP'12).

Our reductions between noisy and bounded leakage are achieved in two steps. First, we put forward a new leakage model (dubbed the *dense leakage* model) and prove that dense leakage can be simulated in the information-theoretic setting using a single query of bounded leakage, up to small statistical distance. Second, we show that the most common noisy-leakage models fall within the class of dense leakage, with good parameters. Third, we prove lower bounds on the amount of bounded leakage required for simulation with sub-constant error, showing that our reductions are nearly optimal. In particular, our results imply that useful general simulation of noisy leakage based on statistical distance and mutual information is impossible. We also provide a complete picture of the relationships between different noisy-leakage models.

Our result finds applications to leakage-resilient cryptography, where we are often able to lift security in the presence of bounded leakage to security in the presence of noisy leakage, both in the information-theoretic and in the computational setting. Additionally, we show how to use lower bounds in communication complexity to prove that bounded-collusion protocols (Kumar, Meka, and Sahai, FOCS'19) for certain functions do not only require long transcripts, but also necessarily need to reveal enough information about the inputs.

1 Introduction

1.1 Background

The security analysis of cryptographic primitives typically relies on the assumption that the underlying secrets (including, e.g., secret keys and internal randomness) are uniformly random to the eyes of the attacker. In reality, however, this assumption may simply be false due to the presence of so-called side-channel attacks [36, 37, 4], where an adversary can obtain partial information (also known as leakage) on the secret state of an implementation of a cryptographic scheme, by exploiting physical phenomena.

Leakage-resilient cryptography [34, 43, 28] aims at bridging this gap by allowing the adversary to launch leakage attacks in theoretical models too. The last decade has seen an impressive amount of work in this area, thanks to which we now dispose of a large number of leakage-resilient cryptographic primitives in different leakage models. We refer the reader to the recent survey by Kalai and Reyzin [35] for an overview of these results.

From an abstract viewpoint, we can think of the leakage on a random variable X (corresponding, say, to the secret key of an encryption scheme) as a correlated random variable $Z = f(X)$ for some leakage function f that can be chosen by the adversary. Depending on the restriction[7] we put on f, we obtain different leakage models. The first such restriction, introduced for the first time by Dziembowski and Pietrzak [28], is to simply assume that the length $\ell \in \mathbb{N}$ of the leakage Z is small enough. This yields the so-called *Bounded Leakage Model.* Thanks to its simplicity and versatility, this model has been used to construct many cryptographic primitives that remain secure in the presence of bounded leakage.

A considerable limitation of the Bounded Leakage Model is the fact that, in real-world side-channel attacks, the leakage obtained by the attacker is rarely bounded in length. For instance, the power trace on a physical implementation of AES typically consists of several Megabytes of information, which is much larger than the length of the secret key.

This motivates a more general notion of *noisy leakage*, where there is no upper bound on the length of Z but instead we assume the leakage is somewhat noisy, in the sense that it does not reveal too much information about X. It turns out that the level of noisiness of the leakage can be measured in several ways, each yielding a different leakage model. The first such model, proposed for the first time by Naor and Segev [44] in the setting of leakage-resilient public-key encryption, assumes that the uncertainty of X given Z drops at most by some parameter $\ell \in \mathbb{R}_{>0}$. The latter can be formalized by means of conditional[8] average min-entropy [22], i.e. by requiring that $\widetilde{\mathbb{H}}_\infty(X|Z) \geq \mathbb{H}_\infty(X) - \ell$. In this work, we will refer to this model as the *Min-Entropy-Noisy (ME-Noisy) Leakage Model.* Dodis, Haralambiev, López-Alt, and Wichs [20] considered a similar model, which

[7] Clearly, there must be some restriction as otherwise $f(X) = X$ and there is no hope for security.

[8] Intuitively, the conditional average min-entropy of a random variable X given Z measures how hard it is to predict X given Z on average (by an unbounded predictor).

we refer to as the *Uniform-Noisy (U-Noisy) Leakage Model*, where the condition about the min-entropy drop is defined w.r.t. the uniform distribution U (rather than on X which may not[9] be uniform).

Another variant of noisy leakage was pioneered by Prouff and Rivain [47] (building on previous work by Chari, Jutla, Rao, and Rohatgi [16]), who suggested to measure the noisiness of the leakage by bounding the Euclidean norm between the joint distribution P_{XZ} and the product distribution $P_X \otimes P_Z$ with some parameter $\eta \in (0, 1)$. Follow-up works by Duc, Dziembowski, and Faust [24] and by Prest, Goudarzi, Martinelli, and Passelègue [46] replaced the Euclidean norm, respectively, with the statistical distance and the mutual information, yielding what we refer to as the *SD-Noisy Leakage* and the *MI-Noisy Leakage Models*. More precisely,[10] Duc, Dziembowski, and Faust considered a strict subset of SD-noisy leakage—hereafter dubbed *DDF-noisy leakage*—for the special case where $X = (X_1, \ldots, X_n)$, for some fixed parameter $n \in \mathbb{N}$, and the function f has a type $f = (f_1, \ldots, f_n)$ such that $\Delta(P_{X_i} \otimes P_{Z_i}, P_{X_i Z_i}) \leq \eta$ for each X_i and $Z_i = f_i(X_i)$. All of these works studied noisy leakage in the setting of leakage-resilient circuit compilers (see §1.4).

The different flavors of noisy leakage discussed above capture either a more general class of leakage functions than bounded leakage (as in the case of ME-noisy and U-noisy leakage), or an orthogonal class of leakage functions (as in the case of SD-noisy and MI-noisy leakage). On the other hand, it is usually easiest (and most common) to prove security of a cryptographic primitive against bounded leakage, whereas extending the analysis to other types of noisy leakage requires non-trivial specialized proofs for each primitive. Motivated by this situation, we consider the following question: *Can we reduce noisy-leakage resilience to bounded-leakage resilience in a general way?*

1.2 Our Results

In this work, we answer the above question to the positive in the information-theoretic setting. In a nutshell, we achieve this by proving that a novel and very general leakage model, which we refer to as the *Dense Leakage Model* and that encompasses all the aforementioned noisy-leakage models, can be simulated almost for free (albeit possibly inefficiently) using a single query of bounded leakage. Our result allows us to show in a streamlined way that many cryptographic primitives which have only been proved to be resilient against bounded leakage are also secure against noisy leakage, with only a small loss in parameters. Importantly, the latter does not only hold for cryptographic schemes with information-theoretic security, but also for ones with computational security only. We elaborate on our contributions in more details in the paragraphs below, and refer the reader to §1.3 for a more technical overview.

[9] For instance, in the setting of public-key encryption [44], the random variable X corresponds to the distribution of the secret key SK given the public key PK, which may not be uniform.

[10] The work by Prest, Goudarzi, Martinelli, and Passelègue considered a similar restriction for MI-noisy leakage.

Simulating dense leakage with bounded leakage. As the starting point for our work, in §3, we introduce a meaningful simulation paradigm between leakage models. Informally, given some random variable X and two families of leakage functions $\mathcal{F}$ and $\mathcal{G}$ on X, we say $\mathcal{F}$ is ε-simulatable from $\mathcal{G}$ if for every $f \in \mathcal{F}$ we can simulate $(X, f(X))$ to within statistical distance ε using a *single* query of the form $g(X)$ for some $g \in \mathcal{G}$.

Taking into account the above simulation paradigm, the question we tackle is whether we can have simulation theorems stating that different noisy-leakage families $\mathcal{F}$ are ε-simulatable from the the family $\mathcal{G}$ of ℓ-bounded leakage (for some small ε). We prove such a simulation theorem for a new leakage model that we call *dense leakage*.

In order to define the Dense Leakage Model, we begin with the concept of *δ-density*: Given two distributions P and P' over a discrete set $\mathcal{X}$, we say P is δ-dense in P' if $P(x) \leq \frac{P'(x)}{\delta}$ for all $x \in \mathcal{X}$. In particular, δ-density implies that $P(x) = 0$ whenever $P'(x) = 0$, and thus this concept is connected to the notion of absolute continuity of one measure with respect to another. Given this notion, it is simple to describe the Dense Leakage Model. If $Z = f(X)$ denotes some leakage from X, then Z is (p, γ, δ)-dense leakage from X if, with probability $1-p$ over the choice of $X = x$, we have $P_{Z|X=x}(z) \leq \frac{P_Z(z)}{\delta}$ with probability $1 - \gamma$ over the choice of $Z = z$. Intuitively, Z being a dense leakage of X essentially corresponds to the distributions $P_{Z|X=x}$ being "approximately" dense in the marginal distribution P_Z for most choices of $x \in \mathcal{X}$.

Our first result is a simulation theorem for dense leakage with respect to bounded leakage, which we state in simplified form below.

Theorem 1 (Informal) *For any random variable X, and every parameter $\varepsilon \in (0,1)$, the family of (p, γ, δ)-dense leakage functions on X is $(\varepsilon + \varepsilon^{1/4\delta} + \gamma + p)$-simulatable from the family of ℓ-bounded leakage functions on X, so long as*

$$\ell \geq \log(1/\delta) + \log\log(1/\varepsilon) + 2\log\left(\frac{1}{1-\gamma}\right) + 2.$$

On the power of dense leakage. Second, we show that dense leakage captures all of the noisy-leakage models considered above. In particular, we obtain the following informal result.

Theorem 2 (Informal) *The families of ME-noisy, U-noisy, and DDF-noisy leakages fall within the family of dense leakage with good*[11] *parameters.*

By combining Theorem 1 and Theorem 2, we obtain non-trivial simulation theorems for the families of ME-noisy, U-noisy, and DDF-noisy leakage from bounded leakage, with small simulation error and small bounded leakage parameter. It is worth mentioning that, for the specific case of ME-noisy leakage, Theorem 2 only holds for distributions X that are almost flat. As we shall prove, this restriction is

[11] In particular, small enough in order to be combined with Theorem 1 yielding interesting applications.

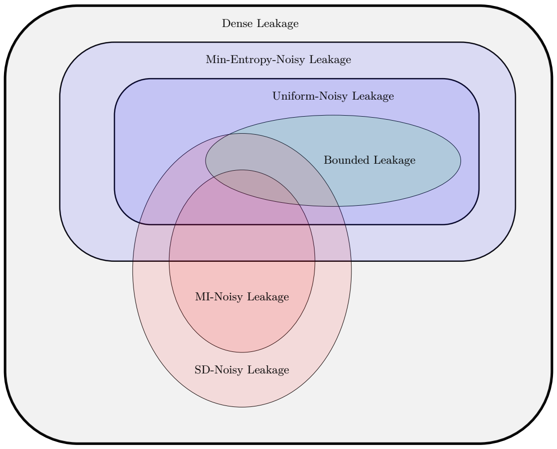

Fig. 1. Containment of the different leakage models considered in this paper. Our main result is that a single query of bounded leakage is enough to simulate dense leakage to within small statistical distance.

nearly optimal in the sense that there exist "non-flat" distributions X for which we cannot simulate ME-noisy leakage on X from bounded leakage on X with good parameters, *even when the drop in min-entropy is minimal.*

Fundamental limitations of SD-noisy and MI-noisy leakages. Turning to the families of SD-noisy and MI-noisy leakage, one can show that they fall within the family of dense leakage too. However, the parameters we obtain in this case are not good enough to be combined with Theorem 1 in order to yield interesting applications. In fact, we prove that the families of η-SD-noisy and η-MI-noisy leakage are trivially simulatable with statistical error roughly η even from the degenerate family of 0-bounded leakage. Unfortunately, this is inherent for the general form of SD-noisy and MI-noisy leakage we consider: we prove that no simulator can achieve simulation error significantly smaller than η even when leaking almost all of the input. In contrast, Duc, Dziembowski, and Faust [23, 24] gave a non-trivial[12] simulation theorem for the family of DDF-noisy leakage (which is a strict subset of SD-noisy leakage) from a special type of bounded leakage called *threshold probing leakage.* Consistently, Theorem 2 establishes that DDF-noisy leakage is dense leakage with good parameters which in combination with Theorem 1 gives an alternative (non-trivial) simulation theorem for DDF-noisy leakage from bounded leakage. While this result is not new, we believe it showcases the generality of our techniques.

A complete picture, and near-optimality of our simulation theorems. We also provide a complete picture of inclusions and separations between the different

[12] In particular, with negligible simulation error and small bounded leakage parameter even for constant η.

leakage models, as depicted in Figure 1. Some of these relationships were already known (e.g., the fact that the family of U-noisy leakage is a strict subset of the family of ME-noisy leakage), and some are new (e.g., the separations between the family of SD-noisy leakage and the families of ME-noisy and MI-noisy leakage).

Moreover, we prove a series of results showing that the amount of bounded leakage we use in our simulation theorems is nearly optimal with respect to the desired simulation error.

Applications in brief. Next, we explore applications of our results to leakage-resilient cryptography. Intuitively, the reason why the simulation paradigm is useful is that it may allow us to reduce leakage resilience of a cryptographic scheme against $\mathcal{F}$ to leakage resilience against $\mathcal{G}$. In particular, when $\mathcal{G}$ is taken to be the family of bounded-leakage functions, we obtain that many primitives which were already known to be secure against bounded leakage are also secure against dense (and thus noisy) leakage. Examples include forward-secure storage [26], leakage-resilient one-way functions and public-key encryption [5], cylinder-intersection extractors [38], symmetric non-interactive key exchange [40], leakage-resilient secret sharing [9, 48, 1, 38, 41] and two-party computation [32].

1.3 Technical Overview

Due to space constraints, most proofs have been deferred to the full version of this paper [13].

Simulation via rejection sampling. We begin by giving an overview of the approach we use to simulate dense leakage from bounded leakage. As discussed before, our goal is to show that, for a random variable X and some associated dense leakage function f (where f may be randomized), there is a (possibly inefficient) simulator that makes *at most one* black-box query $g(X)$ for some ℓ-bounded leakage function $g : \mathcal{X} \to \{0,1\}^\ell$ and outputs $\widetilde{Z}$ such that

$$(X, f(X)) \approx_\varepsilon (X, \widetilde{Z}), \tag{1}$$

where $\approx_\varepsilon$ denotes statistical distance at most ε. For simplicity, we focus here on the setting where f is "exactly" δ-dense leakage from X, meaning that, if $Z = f(X)$, we have $P_{Z|X=x}(z) \leq \frac{P_Z(z)}{\delta}$ for all x and z. This setting is already appropriate to showcase our main ideas.

The key observation that enables the design of our simulator, as we formalize in §2, is that if a distribution P is δ-dense in P', then it is possible to sample $\widetilde{P}$ satisfying $\widetilde{P} \approx_\varepsilon P$ with access only to $s = \frac{\log(1/\varepsilon)}{\delta}$ independent and identically distributed (i.i.d.) samples from P', say $z_1, z_2, \ldots, z_s$, and knowledge of the distribution P, via *rejection sampling*: For $i = 1, 2, \ldots, s$, either output z_i with probability $\delta P(z_i)/P'(z_i) \leq 1$, or move to $i+1$ otherwise (if $i = s+1$, abort).

This suggests the following simulator for f exploiting δ-density: The simulator generates s i.i.d. samples $\mathbf{z} = (z_1, z_2, \ldots, z_s)$ from P_Z. Then, it queries the bounded-leakage oracle with the randomized function $g_{\mathbf{z}}$ which, with full knowledge of x, performs rejection sampling of $P_{Z|X=x}$ from P_Z using $\mathbf{z}$. If rejection sampling outputs z_i, then $g_{\mathbf{z}}(x) = i$, and if rejection sampling aborts we may set $g_{\mathbf{z}}(x) = \perp$. In particular, $g_{\mathbf{z}}$ has $1 + s$ possible outputs, and so it is ℓ-bounded-leakage from X with $\ell = \log(1 + s) \leq \log(1/\delta) + \log\log(1/\varepsilon) + 1$. The behavior of the simulator is now clear: Since it knows $\mathbf{z}$, it can simply output $\widetilde{Z} = z_i$ (or $\widetilde{Z} = \perp$ if rejection sampling aborted). The discussion above guarantees that the output of the simulator is ε-close in statistical distance to $f(x)$, which yields Eq. (1).

As previously discussed, in the actual proof (which appears in §4.1) we must deal with an approximate variant of δ-density. However, we show that the above approach still works in the setting of approximate density at the price of some additional small terms in the simulation error and in the bounded leakage length.

Noisy leakage is dense leakage. As an example of how we manage to frame many types of noisy leakage as dense leakage with good parameters, we discuss how this can be accomplished for ME-noisy leakage assuming X satisfies a property we call α-*semi-flatness*. The full proof appears in §4.2. The property states that X satisfies $P_X(x) \leq 2^\alpha \cdot P_X(x')$ for all $x, x' \in \text{supp}(X)$, and, as we shall see, it is usually satisfied in applications with small α (or even $\alpha = 0$, which corresponds to a flat distribution). We stress that for the case of U-noisy, DDF-noisy, SD-noisy, and MI-noisy leakages, *no* assumption is required on X to place these types of leakage inside the set of dense leakages. More details can be found in §4.3 and §4.4.

Consider some α-semi-flat X and leakage function f such that $Z = f(X)$ satisfies

$$\mathbb{H}_\infty(X|Z = z) \geq \mathbb{H}_\infty(X) - \ell \tag{2}$$

for some $\ell > 0$ and all z. Note that this is a special case of ME-noisy leakage, but it suffices to present the main ideas of our approach. Our goal is to show that f is $(0, 0, \delta)$-dense leakage of X for an appropriate parameter δ, meaning that we wish to prove that $P_{Z|X=x}(z) \leq \frac{P_Z(z)}{\delta}$ for all x and z. Observe that, by Eq. (2), we have $P_{X|Z=z}(x) \leq 2^\ell \max_{x'} P_X(x') \leq 2^{\ell+\alpha} P_X(x)$ for all x and z, where the rightmost inequality makes use of the fact that X is α-semi-flat. Rewriting the inequality above using Bayes' theorem yields $P_{Z|X=x}(z) \leq 2^{\ell+\alpha} P_Z(z)$, meaning that f is $(p = 0, \gamma = 0, \delta = 2^{-\ell-\alpha})$-dense leakage of X. By Theorem 1, we then have that $f(X)$ can be simulated with statistical error 2ε using $\ell' = \ell + \alpha + \log\log(1/\varepsilon) + 2$ bits of bounded leakage from X. This statement allows for significant flexibility in the choice of parameters. For example, setting $\varepsilon = 2^{-\lambda}$ for some security parameter λ yields negligible simulation error from $\ell + \alpha + \log(\lambda) + 2$ bits of bounded leakage. Since α is usually very small in applications (often we have $\alpha = 0$), in practice we can achieve negligible simulation error using $\ell + \log(\lambda) + O(1)$ bits of bounded leakage, i.e., by paying only an extra $\log(\lambda) + O(1)$ bits of leakage. Extending the argument above to general ME-noisy leakage from X requires the addition of small error terms p and γ, but setting parameters similarly to the above still allows us to simulate general ℓ-ME-noisy leakage from X using only, say, $\ell + O(\log^2(\lambda))$ bits of bounded leakage from X.

Trivial simulation of SD-noisy and MI-noisy leakages. Consider the trivial simulator that given the function f simply samples $\widetilde{X}$ according to the distribution of X and then outputs $\widetilde{Z} = f(\widetilde{X})$. Assuming f belongs to the family of η-SD-noisy leakage, the above gives a simulation theorem for SD-noisy leakage with simulation error η (and without requiring any leakage from X). By Pinsker inequality, the above also implies a simulation theorem for η-MI-noisy leakage with simulation error $\sqrt{2\eta}$ (again without leaking anything from X).

Unfortunately, it turns out that one cannot do much better than the trivial simulator (even when using large bounded leakage) for our general definition of SD-noisy leakage. More specifically, there exists some X such that any simulator for a function f that is η-SD-noisy leakage for X must incur a simulation error of at least $\eta/2$ even when leaking all but one bit from X. In the case of MI-noisy leakage, we prove a similar result: There exists an X such that any simulator must have simulation error at least $\frac{\eta}{2n}$ when simulating η-MI-noisy leakage from X, even when leaking all but one bit of X. Notably, this means that negligible simulation error is impossible to achieve when η is non-negligible, and thus one cannot do significantly better than the trivial simulator for MI-noisy leakage either.

It is instructive to compare the above trivial simulation theorem for SD-noisy leakage with the result by Duc, Dziembowski, and Faust [24], who gave a non-trivial simulation theorem for DDF-noisy leakage from a special case of bounded leakage known as threshold probing leakage. Notice that by the triangle inequality, the trivial simulation theorem for η-SD-noisy leakage implies a trivial simulation theorem for η-DDF-noisy leakage with large simulation error $n \cdot \eta$, which in particular becomes uninteresting as soon as η is non-negligible.

Nevertheless, in [13], we show that the family of η-DDF-noisy leakage falls within the family of U-noisy (and thus dense) leakage with good parameters, which in turn gives a non-trivial simulation theorem for η-DDF-noisy leakage from ℓ-bounded leakage with negligible simulation error and for small bounded leakage parameter ℓ, even when $\eta \in (0,1)$ is constant.

Separations between leakage families, and tradeoffs between simulation error and bounded leakage parameter. We complement our positive results in several ways. First, we present missing separations between the different types of leakages we consider in [13], leading to a complete picture of their relationships (as depicted in Figure 1). Second, we study the minimum amount of bounded leakage required to simulate different types of noisy leakage with a given simulation error, and show that our simulation theorems are close to optimal. For example, in the case of ME-noisy leakage, for a large range of ℓ and α we show that $\ell + \alpha - O(1)$ bits of bounded leakage are required to simulate ℓ-ME-noisy leakage from some α semi-flat X. In contrast, as discussed above, our simulation theorem states that approximately $\ell + \alpha$ bits of bounded leakage are sufficient to achieve negligible simulation error.

To showcase our approach towards obtaining tradeoffs between simulation error and the bounded leakage parameter, we discuss here one particularly insightful implication of a more general theorem we obtain, which states that enforcing α-semi-flatness of X is necessary to obtain a non-trivial simulation theorem for ME-noisy leakage with sub-constant simulation error. More precisely, there exists X with support in $\{0,1\}^n$ with an associated 0-noisy leakage function f (meaning that $\widetilde{\mathbb{H}}_\infty(X|f(X)) = \mathbb{H}_\infty(X)$) with the property that simulating $Z = f(X)$ with simulation error less than $1/4$ requires one ℓ'-bounded-leakage query for $\ell' \geq n-2$. In other words, to achieve small simulation error without semi-flatness, we must leak almost all of the input X. The statement above is proved as follows. Consider $X \in \{0,1\}^n$ satisfying $P_X(0^n) = 1/2$ and $P_X(x) = \frac{1}{2(2^n-1)}$ for $x \neq 0^n$. Moreover, set $Z = f(X)$ for a leakage function f such that $f(0^n)$ is uniformly distributed over $\{0,1\}^n \setminus \{0^n\}$ and $f(x) = x$ with probability 1 for $x \neq 0^n$. Routine calculations show that $\mathbb{H}_\infty(X) = 1$ and $\mathbb{H}_\infty(X|Z = z) = 1$ for all z, meaning that $\widetilde{\mathbb{H}}_\infty(X|Z = z) = 1 = \mathbb{H}_\infty(X)$, as desired. Finally, every simulator for (X, Z) above with access to one query of ℓ'-bounded-leakage for $\ell' \leq n-2$ must have simulation error $1/4$ because, conditioned on $X \neq 0^n$ (which holds with probability $1/2$), we have $f(X) = X$ and X uniform over $\{0,1\}^n \setminus \{0^n\}$. Therefore, under this conditioning, we can only correctly guess $f(X)$ with probability at most $1/2$ from any one $(n-2)$-bounded-leakage query of X.

Sample Application: leakage-resilient secret sharing. We now explain how to use our result in order to lift bounded-leakage resilience to noisy-leakage resilience (almost) for free in cryptographic applications. In fact, in the information-theoretic setting, the latter is an almost immediate consequence of our result.

For the purpose of this overview, let us focus on the concrete setting of secret sharing schemes with local leakage resilience [9]. Briefly, a t-out-of-n secret sharing scheme allows to share a message y into n shares $(x_1, \ldots, x_n)$ in such a way that y can be efficiently recovered using any subset of t shares. Local leakage resilience intuitively says that no unbounded attacker obtaining in full all of the shares $x_{\mathcal{U}}$ within an unauthorized subset $\mathcal{U} \subset [n]$ of size $u < t$, and further leaking at most ℓ bits of information z_i from each of the shares x_i independently, should be able to tell apart a secret sharing of message y_0 from a secret sharing of message y_1. Benhamouda, Degwekar, Ishai and Rabin [9] recently proved that both Shamir secret sharing and additive secret sharing satisfy local leakage resilience for certain ranges of parameters.

Thanks to Theorem 1, in §5.1, we show that any secret sharing scheme meeting the above property continues to be secure even if the attacker obtains dense (rather than bounded) leakage on each of the shares x_i independently. The proof of this fact is simple. We move to a mental experiment in which leakages $(z_1, \ldots, z_n)$ corresponding to dense-leakage functions $(f_1, \ldots, f_n)$ are replaced by $(\tilde{z}_1, \ldots, \tilde{z}_n)$ obtained as follows: For each $i \in [n]$, first run the simulator guaranteed by Theorem 1 in order to obtain an ℓ'-bounded leakage function f'_i and compute $z'_i = f'_i(x_i)$; then, run the simulator upon input z'_i in order to obtain a simulated leakage $\tilde{z}_i$.

By a hybrid argument, the above experiment is statistically close to the original experiment. Furthermore, we can reduce a successful attacker in the mental experiment to an attacker breaking local bounded-leakage resilience. The proofs follows. Finally, thanks to Theorem 2, we can use the abstraction of dense leakage in order to obtain security also in the presence of ME-noisy and U-noisy leakage as well. Note that in the case of ME-noisy leakage, for the second step to work, we need that the distribution X_i of each share outside $\mathcal{U}$ given the shares $x_{\mathcal{U}}$ is almost flat, which is the case for Shamir and additive secret sharing.

Applications in the computational setting. The above proof technique can be essentially applied to any cryptographic primitive with bounded leakage resilience in the information-theoretic setting. Further examples include, e.g., forward-secure storage [26], leakage-resilient storage [19], leakage-resilient non-malleable codes [2], non-malleable secret sharing [38, 14] and algebraic manipulation detection codes [6, 42, 3]. (We work out the details for some of these primitives in [13].) However, we cannot apply the same trick in the computational setting or when in the proof of security we need to define an efficient simulator (e.g., for leakage-resilient non-interactive zero knowledge [7] and leakage-resilient multi party computation [9, 32]), as the simulation of dense leakage with bounded leakage guaranteed by Theorem 1 may not be efficient.

Nevertheless, we show that our results are still useful for lifting bounded-leakage to noisy-leakage resilience in the computational setting too. In particular, in [13], we exemplify how to do that for the concrete construction of leakage-resilient one-way functions in the floppy model proposed by Agrawal, Dodis, Vaikuntananthan, and Wichs [5], and in the setting of multi-party computation.

We give an overview of the former application, and refer to [13] for the latter. Let $\mathbb{G}$ be a cyclic group with generator g and prime order q, and define $g_i = g^{\tau_i}$ for each $i \in [n]$. Upon input a vector $\boldsymbol{x} = (x_1, \ldots, x_n)$, the one-way function outputs $y = \prod_{i=1}^{n} g_i^{x_i}$; moreover, there is a refreshing procedure that given y and $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$ can generate a fresh pre-image $\boldsymbol{x}'$ of y by simply letting $\boldsymbol{x}' = \boldsymbol{x} + \boldsymbol{\sigma}$ for randomly chosen $\boldsymbol{\sigma}$ orthogonal to $\boldsymbol{\tau}$. Here, one should think of $\boldsymbol{\tau}$ as a sort of master secret key to be stored in some secure hardware (i.e., the floppy). Agrawal, Dodis, Vaikuntananthan, and Wichs proved that, under the discrete logarithm assumption in $\mathbb{G}$, no efficient attacker can successfully invert y even when given ℓ-bounded leakage on $\boldsymbol{x}$, so long as $\ell \approx (n-3)\log(q)$ and assuming that after each leakage query the value $\boldsymbol{x}$ is refreshed using the floppy. The proof of this fact follows in two steps. First, we move to a mental experiment where each of the leakage queries is answered using a random $(n-2)$-dimensional subspace $\mathcal{S} \subseteq \ker(\boldsymbol{\tau})$. By the subspace hiding lemma [12], this experiment is *statistically close* to the original experiment. Thus, we can use Theorem 1 and Theorem 2 to show that the above still holds in the case of ME-noisy and U-noisy leakage.[13] Second, one finally reduces a successful attacker in the mental experiment to an efficient breaker for the discrete logarithm problem; in this last step, however, the reduction can trivially answer leakage queries by using $\mathcal{S}$, and thus it does not matter whether the leakage is bounded or noisy. We believe the

[13] The former requires the distribution of $\boldsymbol{x}$ given y and $(\mathbb{G}, g, g_1, \ldots, g_n, q)$ to be almost flat which is easily seen to be the case.

above blueprint can be applied to analyze other cryptographic primitives whose leakage resilience is derived through the subspace hiding lemma; we mention a few natural candidates in [13].

Bounded-collusion protocols. Finally, motivated by additional applications to leakage-resilient cryptography and by exploring new lower bounds in communication complexity [49], in §5.2, we investigate the setting of bounded-collusion protocols (BCPs) as proposed by Kumar, Meka, and Sahai [38]. Here, a set of n parties each holding an input x_i wishes to evaluate a Boolean function ϕ of their inputs by means of an interactive protocol π. At the j-th round, a subset of k parties (where $k < n$ is called the collusion bound) is selected, and appends to the protocol transcript τ an arbitrary (possibly unbounded) function f_j of their joint inputs. The goal is to minimize the size ℓ of the transcript, which leads to what we call an ℓ-bounded communication k-bounded collusion protocol (BC-BCP). BC-BCPs interpolate nicely between the well-studied number-in-hand (NIH) [45] (which corresponds to $k = 1$) and number-on-forehead (NOF) [15] (which corresponds to $k = n - 1$) models.

We put forward two natural generalizations of BC-BCPs, dubbed *dense* (resp. *noisy*) communication k-bounded collusion protocols (DC-BCPs, resp. NC-BCP), in which there is no restriction on the length of the final transcript τ but the round functions are either dense or U-noisy leakage functions. It is easy to see that any BC-BCP is also a NC-BCP as well as a DC-BCP. By Theorem 1 and Theorem 2, we are able to show that the converse is also true: namely, we can simulate[14] the transcript τ of any DC-BCP or NC-BCP π using the transcript τ' of a related BC-BCP π' up to a small statistical distance. Protocol π' roughly runs π and uses the simulation paradigm in order to translate the functions used within π into functions to be used within π'. The proof requires a hybrid argument, and thus the final simulation error grows linearly with the number of rounds of the underlying BC-BCP.

The above fact has two consequences. The first consequence is that we can translate communication complexity lower bounds for BC-BCPs into lower bounds on the noisiness of NC-BCPs. A communication complexity lower bound for a Boolean function ϕ says that any BC-BCP computing ϕ with good probability must have long transcripts (i.e., large ℓ). Concrete examples of such functions ϕ include those based on the generalized inner product and on quadratic residues in the NOF model with logarithmic (in the input length) number of parties [18, 8], and more recently a new function (based on the Bourgain extractor [11]) for more general values of k and even for super-logarithmic number of parties [39]. Note that the above lower bounds do not necessarily say how much information a transcript must reveal about the inputs. Thanks to our results, we can show that any NC-BCP (i.e., where there is no upper bound on the transcript length)

[14] The reason for not considering NC-BCPs where the round functions are ME-noisy (instead of U-noisy) leakage functions is that simulating ME-noisy leakage with bounded leakage inherently requires semi-flatness, but we cannot ensure this condition is maintained throughout the entire execution of a leakage protocol.

computing the above functions with good probability must also in some sense reveal enough information about the inputs. However, for technical reasons, the latter holds true only so long as the number of rounds is not too large. We refer the reader to §5.2 for further details.

The second consequence is that we can lift the security of cryptographic primitives whose leakage resilience is modeled as a BC-BCP (which intuitively corresponds to security against adaptive bounded joint leakage) to the more general setting where leakage resilience is modeled as a NC-BCP or DC-BCP (which intuitively corresponds to security against adaptive noisy joint leakage). Examples include secret sharing with security against adaptive joint leakage [38, 17, 39] (see §5.2), extractors for cylinder-intersection sources [38, 40, 17, 39] (see [13]), and leakage-resilient non-interactive key exchange [40] (see [13]). Interestingly, the security of these applications in the bounded-leakage setting has been derived exploiting communication complexity lower bounds for BC-BCPs. We can instead directly lift security to the dense and U-noisy leakage setting in a fully black-box way, and thus without re-doing the analysis.

1.4 Related Work

Naor and Segev [44] conjectured that ME-noisy leakage may be compressed to small leakage in the information-theoretic setting. In this light, our results prove this conjecture to be false for arbitrary distributions and establish the exact conditions under which the above statement holds true not only in the case of ME-noisy leakage, but also for U-noisy leakage.

Most relevant to our work is the line of research on leakage-resilient circuit compilers (see, e.g., [34, 29, 30]), where the equivalence of different leakage models has also been explored. For instance, the beautiful work by Duc, Dziembowski, and Faust [23, 24] shows that DDF-noisy leakage on masking schemes used to protect the internal values within a cryptographic circuit can be simulated by probing a limited number of wires (which can be thought of as bounded leakage in the circuit setting). The notion of DDF-noisy leakage was studied further, both experimentally and theoretically, by Duc, Faust, and Standaert [25]. Follow-up work by Dziembowski, Faust, and Skórski [27] and by Prest, Goudarzi, Martinelli, and Passelègue [46] further improved the parameters of such a reduction and extended it to other noisy-leakage models as well. The difference between the above results and our work is that we prove simulation theorems between very abstract and general leakage models, which ultimately allows us to obtain a broad range of applications which goes beyond the setting of leakage-resilient circuits. In a complementary direction, Fuller and Hamlin [31] studied the relationship between different types of computational leakage.

Harsha, Ishai, Kilian, Nissim, and Venkatesh [33] investigate tradeoffs between communication complexity and time complexity in non-cryptographic settings, including deterministic two-party protocols, query complexity and property testing. Our simulation theorems can be thought of as similar tradeoffs in the cryptographic setting.

1.5 Notation

We denote by $[n]$ the set $\{1, \ldots, n\}$. For a string $x \in \{0,1\}^*$, we denote its length by $|x|$; if $\mathcal{X}$ is a set, $|\mathcal{X}|$ represents the number of elements in $\mathcal{X}$. When x is chosen randomly in $\mathcal{X}$, we write $x \leftarrow\!\!\$\ \mathcal{X}$. When A is a randomized algorithm, we write $y \leftarrow\!\!\$\ \mathsf{A}(x)$ to denote a run of A on input x (and implicit random coins r) and output y; the value y is a random variable and $\mathsf{A}(x;r)$ denotes a run of A on input x and randomness r. An algorithm A is *probabilistic polynomial-time* (PPT for short) if A is randomized and for any input $x, r \in \{0,1\}^*$, the computation of $\mathsf{A}(x;r)$ terminates in a polynomial number of steps (in the size of the input). For a random variable X, we write $\mathbb{P}[X = x]$ for the probability that X takes on a particular value $x \in \mathcal{X}$, with $\mathcal{X}$ being the set where X is defined. The probability mass function of X is denoted P_X, *i.e.*, $P_X(x) = \mathbb{P}[X = x]$ for all $x \in \mathcal{X}$; we sometimes omit X and just write P when X is clear from the context. For a set (or event) $\mathcal{S} \subseteq \mathcal{X}$, we write $P(\mathcal{S})$ for the probability of event $\mathcal{S}$, i.e., $P(\mathcal{S}) = \sum_{x \in \mathcal{S}} P(x)$. We denote the statistical distance between two distributions P and P' by $\Delta(P; P')$. The min-entropy of a random variable X is denoted by $\mathbb{H}_\infty(X)$, and the average conditional min-entropy of X given Y is denoted by $\widetilde{\mathbb{H}}_\infty(X|Y)$.

We denote with $\lambda \in \mathbb{N}$ the security parameter. A function p is *polynomial* (in the security parameter), denoted $p \in \mathrm{poly}(\lambda)$, if $p(\lambda) \in O(\lambda^c)$ for some constant $c > 0$. A function $\nu : \mathbb{N} \to [0,1]$ is *negligible* (in the security parameter) if it vanishes faster than the inverse of any polynomial in λ, i.e. $\nu(\lambda) \in O(1/p(\lambda))$ for all positive polynomials $p(\lambda)$. We sometimes write $\mathrm{negl}(\lambda)$ to denote an unspecified negligible function. Unless stated otherwise, throughout the paper, we implicitly assume that the security parameter is given as input (in unary) to all algorithms.

Basic definitions and lemmas in cryptography used throughout the paper are discussed in [13].

2 Rejection Sampling for Approximate Density

The problem that we consider in this section is the following: How can we sample from a distribution P with statistical error at most ε, given only *black-box* access to i.i.d. samples from another distribution P'?

It turns out that the problem above can be solved via rejection sampling, assuming that P is approximately dense in P' as defined below.

Definition 1 (δ-density) *Given distributions P and P' over a set $\mathcal{Z}$ and $\delta \in (0,1]$, we say P is* δ-dense in P' *if for every $z \in \mathcal{Z}$ it holds that $P(z) \leq \frac{P'(z)}{\delta}$.*

Definition 2 ((γ,δ)-density) *Given distributions P and P' over a set $\mathcal{Z}$ and $\gamma \in [0,1]$, $\delta \in (0,1]$, we say P is* γ-approximate δ-dense in P'*, or simply* (γ,δ)-dense in P'*, if there exists a set $\mathcal{S} \subseteq \mathcal{Z}$ such that $P(\mathcal{S}), P'(\mathcal{S}) \geq 1 - \gamma$, and for all $z \in \mathcal{S}$ it holds that $P(z) \leq \frac{P'(z)}{\delta}$.*

2.1 The Case of Exact Density

First, we consider the special case where P is δ-dense in P'.

Lemma 1 *Suppose P is δ-dense in P'. Then, for any $\varepsilon \in (0,1]$, it is possible to sample $\widetilde{P}$ such that $\widetilde{P} \approx_\varepsilon P$ given access to $s = \frac{\log(1/\varepsilon)}{\delta}$ i.i.d. samples from P'.*

Proof. Consider the following rejection sampling algorithm:

1. Sample $z_1, \ldots, z_s$ i.i.d. according to the distribution P', and set $y := \perp$;
2. For $i = 1, \ldots, s$ do the following: Set $B_i := 1$ with probability $p_i = \frac{\delta P(z_i)}{P'(z_i)}$ and $B_i := 0$ otherwise. If $B_i := 1$, set $y := z_i$ and stop the cycle;
3. Output y.

Observe that $\frac{\delta P(z_i)}{P'(z_i)} \leq 1$ for all z_i (hence the algorithm above is valid), and that the probability that the algorithm outputs some z in the i-th round is

$$\mathbb{P}[B_i = 1] = \sum_z P'(z) \cdot \frac{\delta P(z)}{P'(z)} = \delta. \tag{3}$$

Let $\widetilde{P}$ denote the distribution of the output of the algorithm above and let Y be the corresponding output. Observe that $(Y|Y \neq \perp)$ is distributed exactly like P. This holds because, in view of Eq. (3), the probability that the algorithm outputs z in the i-th round given that it stops in the i-th round is

$$\mathbb{P}[Y = z | B_i = 1, \forall j < i : B_j = 0] = \frac{(1-\delta)^{i-1} \cdot P'(z) \cdot \frac{\delta P(z)}{P'(z)}}{(1-\delta)^{i-1} \cdot \delta} = P(z).$$

Moreover, we have

$$\mathbb{P}[Y = \perp] = (1-\delta)^s \leq \exp(-\delta \cdot s) = \varepsilon.$$

From these observations, we conclude that $\mathit{\Delta}(\widetilde{P}; P) \leq \Pr[Y = \perp] \leq \varepsilon$.

2.2 The Case of Approximate Density

The analogous result for approximate density follows by a similar proof.

Lemma 2 *Suppose P is (γ, δ)-dense in P'. Then, for any $\varepsilon \in (0,1]$, it is possible to sample $\widetilde{P}$ such that $\widetilde{P} \approx_{\varepsilon + \varepsilon^{\frac{1}{4\delta}} + \gamma} P$ given access to $\frac{2\log(1/\varepsilon)}{\delta(1-\gamma)^2}$ i.i.d. samples from P'.*

3 Leakage Models

In this section, we review several leakage models from the literature, and introduce the simulation paradigm which will later allow us to draw connections between different leakage models. Our take is very general, in that we think as the leakage as a randomized function f on a random variable X, over a set $\mathcal{X}$, which yields a correlated random variable $Z = f(X)$. Different leakage models are then obtained by putting restrictions on the joint distribution (X, Z). We refer the reader to §5 for concrete examples of what the distribution X is in applications.

3.1 Bounded Leakage

A first natural restriction is to simply assume an upper bound $\ell \in \mathbb{N}$ on the total length of the leakage. This yields the so-called Bounded Leakage Model, which was formalized for the first time by Dziembowski and Pietrzak [28].

Definition 3 (Bounded leakage) *Given a random variable X over $\mathcal{X}$, we say a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is an ℓ-*bounded leakage function for *X if $\mathcal{Z} \subseteq \{0,1\}^\ell$. For fixed X, we denote the set of all its ℓ-bounded leakage functions by $\mathsf{Bounded}_\ell(X)$.*

3.2 Noisy Leakage

A considerable drawback of the Bounded Leakage Model is that physical leakage is rarely of bounded length. The Noisy Leakage Model overcomes this limitation by assuming that the length of the leakage is unbounded but somewhat *noisy*.

There are different ways from the literature how to measure the noisiness of the leakage. A first way, considered for the first time by Naor and Segev [44], is to assume that the leakage drops the min-entropy of X by at most $\ell \in \mathbb{R}_{>0}$ bits. We will refer to this model as the ME-Noisy Leakage Model.

Definition 4 (ME-noisy leakage) *Given a random variable X over $\mathcal{X}$, we say a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is an ℓ-*ME-noisy leakage function for *X if, denoting $Z = f(X)$, we have $\widetilde{\mathbb{H}}_\infty(X|Z) \geq \mathbb{H}_\infty(X) - \ell$. For fixed X, we denote the set of all its ℓ-ME-noisy leakage functions by $\mathsf{Noisy}_{\infty,\ell}(X)$.*

Dodis *et al.* [20] considered a slight variant of the above definition where the min-entropy drop is measured w.r.t. the uniform distribution U over $\mathcal{X}$ (rather than X itself). We will refer to this model as the U-Noisy Leakage Model.

Definition 5 (U-noisy leakage) *Given a random variable X over $\mathcal{X}$, we say a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is an ℓ-*U-noisy leakage function for *X if it holds that $\widetilde{\mathbb{H}}_\infty(U|f(U)) \geq \mathbb{H}_\infty(U) - \ell$, where U denotes the uniform distribution over $\mathcal{X}$. For fixed X, we denote the set of all its ℓ-U-noisy leakage functions by $\mathsf{UNoisy}_{\infty,\ell}(X)$.*

A second way to measure noisiness is to assume that the leakage only implies a bounded bias in the distribution X, which is formally defined as distributions P_{XZ} and $P_X \otimes P_Z$ being close according to some distance when seen as real-valued vectors. Prouff and Rivain [47] were the first to consider this restriction using the Euclidean norm (i.e., the ℓ_2-norm), whereas Duc, Dziembowski and Faust [24] used the statistical distance (i.e., the ℓ_1-norm). We recall the latter definition below, which we will refer to as the SD-Noisy Leakage Model.

Definition 6 (SD-noisy leakage) *Given a random variable X over $\mathcal{X}$, we say a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is an η-*SD-noisy leakage function for *X if, denoting $Z = f(X)$, it holds that $\Delta(P_{XZ}; P_X \otimes P_Z) \leq \eta$, where $P_X \otimes P_Z$ denotes the product distribution of X and Z. For fixed X, we denote the set of all its η-SD-noisy leakage functions by $\mathsf{Noisy}_{\Delta,\eta}(X)$.*

Duc, Dziembowski, and Faust [24] considered only a restricted subset of SD-noisy leakage, which we call *DDF-noisy leakage.* We discuss it in [13], placing it with respect to other leakage models and deriving an associated simulation theorem.

Alternatively, as suggested by Prest *et al.* [46], we can measure the noisiness of the leakage by looking at the mutual information between X and Z. We can define the mutual information between X and Z as $I(X;Z) = D_{\mathsf{KL}}(P_{XZ} \| P_X \otimes P_Z)$, where $D_{\mathsf{KL}}(P\|P') = \sum_{x\in\mathcal{X}} P(x)\log\left(\frac{P(x)}{P'(x)}\right)$ is the Kullback-Leibler divergence between P and P'.

Definition 7 (MI-noisy leakage) *Given a random variable X over $\mathcal{X}$, we say a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is an η-*MI-noisy leakage function for X *if, denoting $Z = f(X)$, it holds that $I(X;Z) \leq \eta$. For fixed X, we denote the set of all its η-MI-noisy leakage functions by $\mathsf{MINoisy}_\eta(X)$.*

The well-known Pinsker inequality allows us to relate MI-noisy leakage to SD-noisy leakage.

Lemma 3 (Pinsker inequality) *For arbitrary distributions P and P' over a set $\mathcal{X}$ it holds that $\Delta(P;P') \leq \sqrt{2 \cdot D_{\mathsf{KL}}(P\|P')}$.*

As an immediate corollary of Lemma 3, we obtain the following result (which was observed also in [46]).

Corollary 1 *For any $\eta > 0$ and X we have $\mathsf{MINoisy}_\eta(X) \subseteq \mathsf{Noisy}_{\Delta,\sqrt{2\eta}}(X)$.*

3.3 Dense Leakage

Next, we introduce a new leakage model which we dub the Dense Leakage Model. This model intuitively says that the distribution of $Z|X = x$ is approximately dense in the distribution of Z for a large fraction of x's. Looking ahead, dense leakage will serve as a powerful abstraction to relate different leakage models.

Definition 8 (Dense leakage) *Given a random variable X over $\mathcal{X}$, we say a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is a (p,γ,δ)-*dense leakage function for X *if, denoting $Z = f(X)$, there exists a set $\mathcal{T} \subseteq \mathcal{X}$ with $P_X(\mathcal{T}) \geq 1 - p$ such that $P_{Z|X=x}$ is (γ,δ)-dense in P_Z for all $x \in \mathcal{T}$. For fixed X, we denote the set of all its (p,γ,δ)-dense leakage functions by $\mathsf{Dense}_{p,\gamma,\delta}(X)$.*

3.4 The Simulation Paradigm

Finally, we define the simulation paradigm which allows to draw connections between different leakage models. Intuitively, for any random variable X, we will say that a leakage family $\mathcal{F}(X)$ is simulatable from another leakage family $\mathcal{G}(X)$ if for all functions $f \in \mathcal{F}(X)$ there exists a simulator Sim_f which can generate $\widetilde{Z}$ such that (X,Z) and $(X,\widetilde{Z})$ are statistically close, using a single sample $g(X)$ for some function $g \in \mathcal{G}(X)$.

Definition 9 (Leakage simulation) *Given a random variable X and two leakage families $\mathcal{F}(X)$ and $\mathcal{G}(X)$, we say $\mathcal{F}(X)$ is ε-*simulatable from $\mathcal{G}(X)$ *if for all $f \in \mathcal{F}(X)$ there is a (possibly inefficient) randomized algorithm Sim_f such that $(X, Z) \approx_\varepsilon (X, \mathsf{Sim}_f^{\mathsf{Leak}(X,\cdot)})$, where $Z = f(X)$ and the oracle $\mathsf{Leak}(X, \cdot)$ accepts a single query $g \in \mathcal{G}(X)$ and outputs $g(X)$.*

Remark 1 (On the simulator). Note that since the simulator Sim_f knows the distribution P_X of X and the leakage function f, it also knows the joint distribution $P_{X,Z}$ where $Z = f(X)$. We will use this fact to design our leakage simulators. We will also sometimes think of the simulator Sim_f as two machines with a shared random tape, where the first machine outputs the description of a leakage function $g \in \mathcal{G}(X)$, while the second machine outputs the simulated leakage $\tilde{Z}$ given the value $g(X)$.

4 Relating Different Leakage Models

In this section, we show both implications and separations between the leakage models defined in §3. In a nutshell, our implications show that all the noisy-leakage models from §3 can be simulated by bounded leakage with good parameters. We achieve this in two main steps: First, we prove that dense leakage can be simulated by bounded leakage with good parameters. Second, we show that dense leakage contains the other leakage models we have previously defined. Combining the two steps above, we conclude that many different leakage models can be simulated by bounded leakage with good parameters. To complement these results, our separations show that the containment of the different leakage models in dense leakage are essentially the best we can hope for in general.

The simulation theorem for the case of ME-noisy leakage only holds for certain distributions X, which are nevertheless the most relevant in applications. In particular, we will require to assume that the random variable X is semi-flat, as formally defined below.

Definition 10 (Semi-flat distribution) *We say that X is α-*semi-flat *if for all $x, x' \in \operatorname{supp}(X)$ we have $P_X(x) \leq 2^\alpha \cdot P_X(x')$.*

4.1 Simulating Dense Leakage with Bounded Leakage

The following theorem states that one dense leakage query can be simulated with one bounded leakage query to within small statistical error. The efficiency of the simulator and the bounded leakage function is essentially governed by the density parameter δ.

Theorem 3 *For arbitrary X, and for any $\varepsilon \in (0, 1]$, the set of dense leakages $\mathsf{Dense}_{p,\gamma,\delta}(X)$ is $(\varepsilon + \varepsilon^{1/4\delta} + \gamma + p)$-simulatable from $\mathsf{Bounded}_\ell(X)$ with*

$$\ell = 1 + \log\left(\frac{2\log(1/\varepsilon)}{(1-\gamma)^2\delta}\right) = \log(1/\delta) + \log\log(1/\varepsilon) + 2\log\left(\frac{1}{1-\gamma}\right) + 2.$$

Proof. Fix any $f \in \mathsf{Dense}_{p,\gamma,\delta}(X)$. By hypothesis, there is a set $\mathcal{T} \subseteq \mathcal{X}$ such that $P_X(\mathcal{T}) \geq 1 - p$ and $P_{Z|X=x}$ is (γ, δ)-dense in P_Z for all $x \in \mathcal{T}$. Thus, we may assume that $X \in \mathcal{T}$ by adding p to the simulation error.

We consider the simulator Sim_f which, given the distribution P_{XZ}, samples $s^* = \frac{2\log(1/\varepsilon)}{(1-\gamma)\delta}$ i.i.d. samples $\mathbf{z} = (z_1, z_2, \ldots, z_{s^*})$ from P_Z. Then, Sim_f makes a query to $Z' = g_{\mathbf{z}}(X) \in \mathsf{Bounded}_\ell(X)$, where $\ell = 1 + \log s^*$ and $g_{\mathbf{z}} : \mathcal{X} \to \{0,1\}^\ell$ on input $x \in \mathcal{T}$ runs the rejection sampling algorithm from the proof of Lemma 2 to sample from $P_{Z|X=x}$ to within statistical error $\varepsilon + \varepsilon^{1/4\delta} + \gamma$ using the s^* i.i.d. samples $(z_1, \ldots, z_{s^*})$ from P_Z, and outputs the index $i \leq s^*$ such that z_i is output by the rejection sampling algorithm, or $s^* + 1$ if this algorithm outputs $\perp$. Finally, if $I = g_{\mathbf{z}}(X) \leq s^*$, then Sim_f outputs z_I, and otherwise it outputs $\perp$. Let $\widetilde{Z}$ the random variable corresponding to the output of the simulator. Summing up all simulation errors, Lemma 2 guarantees that $(X, Z) \approx_{\varepsilon + \varepsilon^{1/4\delta} + \gamma + p} (X, \widetilde{Z})$, which completes the proof. □

Remark 2 (On useful parameters). The statement of Theorem 3 is most useful when ε, γ, and p are negligible in the security parameter, so as to obtain negligible simulation error. The parameter δ essentially dictates the number of bits of bounded leakage required to simulate a given class of dense leakages. Indeed, it is usually the case that $\log\log(1/\varepsilon) + 2\log\left(\frac{1}{1-\gamma}\right)$ *is much smaller than* $\log(1/\delta)$.

Remark 3 (On efficiency of the simulation). The complexity of the simulator from Theorem 3 is dominated by the complexity of computing the distributions P_Z (possible with knowledge of P_X and f) and $P_{Z|X=x}$ (possible with knowledge of X and f), and of sampling both the z_i according to P_Z and the decision in each step of rejection sampling. If these steps can be implemented in polynomial time with respect to some parameter of interest, then the simulator is efficient.

4.2 Min-Entropy-Noisy Leakage is Dense Leakage

The following theorem states that all ME-noisy leakage is also dense leakage for semi-flat distributions. Looking ahead, we will later establish that the semi-flatness condition is necessary.

Theorem 4 *Suppose X is α-semi-flat. Then, for every $\beta > 0$ and $\ell > 0$, and for $p = 2^{-\beta/2}$, $\gamma = 2^{-\beta/2}$ and $\delta = 2^{-(\ell+\beta+\alpha)}$, we have* $\mathsf{Noisy}_{\infty,\ell}(X) \subseteq \mathsf{Dense}_{p,\gamma,\delta}(X)$.

Combining Theorem 3 and Theorem 4, we immediately obtain the following corollary.

Corollary 2 *If X is α-semi-flat, then, for any $\beta > 0$ and $\varepsilon > 0$, the set of leakages* $\mathsf{Noisy}_{\infty,\ell}(X)$ *is* $(\varepsilon + \varepsilon^{2^{\ell+\beta+\alpha-2}} + 2^{-\beta/2+1})$*-simulatable from* $\mathsf{Bounded}_{\ell'}(X)$ *with* $\ell' = \ell + \beta + \alpha + \log\log(1/\varepsilon) + 2\log\left(\frac{1}{1-2^{-\beta/2}}\right) + 2$.

The remark below says that there is a natural tradeoff between the simulation error in the above corollary and the leakage bound.

Remark 4 (Trading simulation error with ME-noisy leakage). By choosing $\varepsilon = 2^{-\lambda}$ and $\beta = 2 + \log^2(\lambda)$ in Corollary 2, we can obtain negligible simulation error $\varepsilon' = \lambda^{-\omega(1)}$ with leakage[15] $\ell' = \ell + O(\log^2(\lambda) + \alpha)$. By choosing $\beta = \lambda$, we can instead obtain a much smaller simulation error of $\varepsilon' = 2^{-\Omega(\lambda)}$ with larger leakage $\ell' = \ell + O(\lambda + \alpha)$.

Near-optimality of simulation theorem for ME-noisy leakage. We now show that our simulation result for ME-noisy leakage (Corollary 2) is essentially optimal. More precisely, we obtain the following result.

Theorem 5 *For every n and $\alpha, \ell > 0$ such that $\ell + \alpha < n - 2$ there exists an $(\alpha+1)$-semi-flat random variable X and $f \in \mathsf{Noisy}_{\infty,\ell+2}(X)$ such that simulating $f(X)$ with error less than $1/4$ requires one ℓ'-bounded leakage query for $\ell' \geq \ell + \alpha - 1$.*

Essentially, Theorem 5 states that $\ell+\alpha-O(1)$ bits of bounded leakage are required to simulate ℓ-ME-noisy leakage from an α-semi-flat random variable X with useful simulation error. Our simulation theorem from Corollary 2 complements this negative result, showing that $\ell' \approx \ell + \alpha$ bits of bounded leakage are enough even with *negligible* simulation error.

Necessity of the semi-flatness assumption in Corollary 2. Theorem 5 implies that assuming α-semi-flatness of X is necessary to obtain a non-trivial simulation theorem for ME-noisy leakage, even when we are attempting to simulate only 0-ME-noisy leakage functions. Indeed, setting $\ell = 0$ and $\alpha = n - 3$ in Theorem 5, we conclude that there exists a random variable X along with an associated 0-ME-noisy-leakage function $f \in \mathsf{Noisy}_{\infty,0}(X)$ that requires $n - O(1)$ bits of bounded leakage from X in order to be simulated with error less than $1/4$.

Note also that the proof of Theorem 5 shows the impossibility of non-trivial simulation theorems for ME-noisy leakage even for a restricted subset of semi-flat distributions X for which there exists x^* such that $P_X(x^*)$ may be large but $(X|X \neq x^*)$ is flat.

4.3 Uniform-Noisy Leakage is Also Dense Leakage

There is a known connection between U-noisy and ME-noisy leakage, i.e., every U-noisy leakage function is also a ME-noisy leakage function by itself.

Lemma 4 ([20]) *Given any randomized function $f : \mathcal{X} \to \mathcal{Z}$, if it holds that $\widetilde{\mathbb{H}}_\infty(U|f(U)) \geq \mathbb{H}_\infty(U) - \ell$, then for any X over $\mathcal{X}$ it is the case that $\widetilde{\mathbb{H}}_\infty(X|f(X)) \geq \mathbb{H}_\infty(X) - \ell$. In particular, this implies that $\mathsf{UNoisy}_{\infty,\ell}(X) \subseteq \mathsf{Noisy}_{\infty,\ell}(X)$.*

[15] In fact, we can push the leakage bound down to $\ell' = \ell + O(\log\log(\lambda)\log(\lambda) + \alpha)$ or even $\ell' = \ell + O(\log^*(\lambda)\log(\lambda) + \alpha)$, while still obtaining negligible simulation error.

We remark that there also exist some X and a leakage function f such that $f \in \mathsf{Noisy}_{\infty,\ell}(X)$ but $f \notin \mathsf{UNoisy}_{\infty,\ell}(X)$ (such an example is provided in [20]). This shows that the containment of U-noisy leakage in ME-noisy leakage may be strict for some X.

Although Lemma 4 immediately yields an analogue of Corollary 2 for U-noisy leakage, we can obtain a better result by arguing directly that every U-noisy leakage function is also a dense leakage function *for arbitrary* X, i.e., without requiring that X be semi-flat. Our result is stated formally in the next theorem.

Theorem 6 *For every $\beta > 0$ and X, we have $\mathsf{UNoisy}_{\infty,\ell}(X) \subseteq \mathsf{Dense}_{p,\gamma,\delta}(X)$, where $p = 2^{-\beta/2}$, $\gamma = 2^{-\beta/2}$ and $\delta = 2^{-(\ell+\beta)}$.*

Combining Theorem 3 and Theorem 6 immediately yields the following result.

Corollary 3 *For every X and every $\beta > 0$ and $\varepsilon > 0$, the set of leakages $\mathsf{UNoisy}_{\infty,\ell}(X)$ is $(\varepsilon + \varepsilon^{2^{\ell+\beta-2}} + 2^{-\beta/2+1})$-simulatable from $\mathsf{Bounded}_{\ell'}(X)$ with $\ell' = \ell + \beta + \log\log(1/\varepsilon) + 2\log\left(\frac{1}{1-2^{-\beta/2}}\right) + 2$.*

The remark below says that there is a natural tradeoff between the simulation error in the above corollary and the leakage bound.

Remark 5 (Trading simulation error with U-noisy leakage). By choosing $\varepsilon = 2^{-\lambda}$ and $\beta = 2 + \log^2(\lambda)$ in Corollary 3, we can obtain negligible simulation error $\varepsilon' = \lambda^{-\omega(1)}$ with leakage $\ell' = \ell + O(\log^2(\lambda))$. By choosing $\beta = \lambda$, we can instead obtain a much smaller simulation error of $\varepsilon' = 2^{-\Omega(\lambda)}$ with larger leakage $\ell' = \ell + O(\lambda)$.

Near-optimality of simulation theorem for U-noisy leakage. We now show that in order to simulate arbitrary ℓ-U-noisy leakage from X uniformly distributed over $\{0,1\}^n$ with simulation error less than $1/2$, we need access to one query of ℓ'-bounded leakage from X for $\ell' \geq \ell - 1$. As we shall see, this result implies that our simulation theorem from Corollary 3 is nearly optimal.

Theorem 7 *For X uniform over $\{0,1\}^n$ and every $\ell \geq 1$ there exists $f \in \mathsf{UNoisy}_{\infty,\ell}(X) \subseteq \mathsf{Noisy}_{\infty,\ell}(X)$ such that $f(X)$ cannot be simulated with error less than $1/2$ by one $(\ell-1)$-bounded leakage query to X. Moreover, it also holds that $f \in \mathsf{Dense}_{p=0,\gamma=0,\delta=2^{-\ell}}(X)$.*

Comparing Theorem 7 with Corollary 3, we see that our simulation theorem for U-noisy leakage is nearly optimal with respect to the bounded leakage parameter, since we only require approximately ℓ bits of bounded leakage to simulate U-noisy leakage for uniform X. Furthermore, we can achieve this result with negligible simulation error.

4.4 SD-Noisy and MI-Noisy Leakage are Also Dense Leakage

We now proceed to relate SD-noisy leakage and dense leakage.

Theorem 8 *For every $\gamma > 0$ and X, we have $\mathsf{Noisy}_{\Delta,\eta}(X) \subseteq \mathsf{Dense}_{p,\gamma,\delta}(X)$ with $p = 2\eta/\gamma$ and $\delta = 1/2$.*

By combining Corollary 1 and Theorem 8, we immediately obtain an analogous result for MI-noisy leakage.

Theorem 9 *For every $\gamma > 0$ and X, we have $\mathsf{MINoisy}_{\eta}(X) \subseteq \mathsf{Dense}_{p,\gamma,\delta}(X)$ with $p = \sqrt{8\eta}/\gamma$ and $\delta = 1/2$.*

Near-optimality of trivial simulator for SD-noisy and MI-noisy leakages. While one can combine Theorem 8 and Theorem 9 with Theorem 3 in order to obtain simulation theorems for SD-noisy and MI-noisy leakage from bounded leakage, it turns out that these simulation theorems do not perform better than the trivial simulator that makes no bounded leakage queries to X: Sample $\widetilde{X}$ according to P_X, and output $\widetilde{Z} = f(\widetilde{X})$. In [13], we show that this is inherent, since the trivial simulator is nearly optimal for SD-noisy and MI-noisy leakages.

5 Applications

In this section we show that our results have interesting implications for so-called leakage-resilient cryptography. In particular, we will show that many cryptographic primitives that have been shown to be resilient to bounded leakage are also resilient to different forms of noisy leakage, with only a small loss in parameters.

5.1 Secret Sharing with Local Leakage Resilience

In this section, we consider the following kind of local leakage attack on a threshold secret sharing scheme: after seeing an unauthorized subset of shares, the adversary performs one query of leakage from all the shares independently.

Definition 11 (Local leakage-resilient secret sharing) *Let $t, n, u \in \mathbb{N}$ be parameters such that $u < t \leq n$, and let $\Sigma = (\mathsf{Share}, \mathsf{Rec})$ be a t-out-of-n secret sharing scheme. We say that Σ is a (p, γ, δ)-dense u-local ε-leakage-resilient secret sharing scheme (or $(u, p, \gamma, \delta, \varepsilon)$-DLLR-SS for short) if for all messages $y_0, y_1 \in \{0,1\}^m$, all unauthorized subsets $\mathcal{U} \subseteq [n]$ such that $|\mathcal{U}| \leq u$, and every tuple of leakage functions $(f_1, \ldots, f_n)$ such that f_i is (p, γ, δ)-dense for all $i \in [n]$, we have $\Delta\big((X^0_{\mathcal{U}}, (f_i(X^0_{\mathcal{U}}, X^0_i))_{i\in[n]}), (X^1_{\mathcal{U}}, (f_i(X^1_{\mathcal{U}}, X^1_i))_{i\in[n]})\big) \leq \varepsilon$, where $(X^b_1, \ldots, X^b_n) = \mathsf{Share}(y_b)$ for all $b \in \{0,1\}$.*

Moreover, in case the functions f_i in the above definition are:

- ℓ-bounded leakage functions, we say that Σ is ℓ-bounded u-local ε-leakage-resilient (or (u, ℓ, ε)-BLLR-SS);
- ℓ-ME-noisy leakage functions, we say that Σ is ℓ-min-entropy-noisy u-local ε-leakage-resilient (or (u, ℓ, ε)-ME-NLLR-SS);

- ℓ-U-noisy leakage functions, we say that Σ is ℓ-uniform-noisy u-local ε-leakage-resilient (or (u, ℓ, ε)-U-NLLR-SS).

The theorem below says that any bounded leakage-resilient secret sharing scheme is also secure in the presence of dense leakage.

Theorem 10 *Any (u, ℓ, ε)-BLLR-SS is also a $(u, p, \gamma, \delta, \varepsilon')$-DLLR-SS so long as*

$$\ell = \log(1/\delta) + \log\log(1/\varepsilon) + 2\log\left(\frac{1}{1-\gamma}\right) + 2$$
$$\varepsilon' = (2n+1)\varepsilon + 2n\varepsilon^{1/4\delta} + 2n\gamma + 2np.$$

Next, using the connection between ME-noisy and U-noisy leakage with dense leakage established in §4, we obtain the following corollary.

Corollary 4 *Any (u, ℓ', ε')-BLLR-SS is also an:*

(i) *(u, ℓ, ε)-ME-NLLR-SS so long as $\ell = \ell' - 2\log(1/\varepsilon') - \alpha - \log\log(1/\varepsilon') - 1$ and $\varepsilon = (6n+1)\varepsilon'$, and assuming that $(X_1, \ldots, X_n) = \mathsf{Share}(y)$ is such that X_i is α-semi-flat for all $i \in [n]$.*

(ii) *(u, ℓ, ε)-U-NLLR-SS so long as $\ell = \ell' - 2\log(1/\varepsilon') - \log\log(1/\varepsilon') - 1$ and $\varepsilon = (6n+1)\varepsilon'$.*

Proof. The statement follows by choosing $\beta = 2 + 2\log(1/\varepsilon')$ and $\varepsilon = \varepsilon'$ in Corollary 2 and Corollary 3. □

We present concrete instantiations of Corollary 4 in [13].

5.2 Bounded-Collusion Protocols

In this section, we deal with applications related to so-called bounded-collusion protocols (BCPs). These are interactive protocols between n parties where at each round a subset of $k < n$ parties are selected, and the output of a leakage function applied to the input of such parties is appended to the protocol's transcript.

Definition 12 (Bounded-communication BCPs) *An interactive (possibly randomized) protocol π is called an n-party r-round ℓ-bounded communication k-bounded-collusion protocol ((n, r, ℓ, k)-BC-BCP, for short) if:*

(i) *the n parties start the protocol with input $x_1, \ldots, x_n \in \mathcal{X}$, and the transcript τ is empty at the beginning of the protocol;*

(ii) *there is a function $\mathsf{Next} : \{0,1\}^* \to \binom{[n]}{k}$ taking as input a (partial) transcript τ and outputting a set $\mathcal{S} \subset [n]$ with $|\mathcal{S}| = k$ along with a function $f : \mathcal{X}^k \to \{0,1\}^*$;*

(iii) *at each round $j \in [r]$ with current transcript τ, the protocol runs $\mathsf{Next}(\tau)$ obtaining $(\mathcal{S}_j, f_j)$ and appends the message $f_j(X_{\mathcal{S}_j})$ to the current transcript τ;*

(iv) *the final transcript τ consists of at most $\ell \in \mathbb{N}$ bits.*

The above notion, which was introduced by Kumar, Meka, and Sahai [38], interpolates nicely between the well-known *number-in-hand* (NIH) and *number-on-forehead* (NOF) models, which correspond respectively to the extreme cases $k = 1$ and $k = n - 1$. Note that the number of rounds in a BC-BCP is at most $r \leq \ell$.

Below, we generalize the definition of BCPs to settings where the round functions correspond to noisy-leakage (in particular, dense and uniform-noisy leakage) functions on the parties' inputs, and thus there is no restriction on the size of the final transcript.

Definition 13 (Dense-communication BCPs) *An interactive (possibly randomized) protocol π is called an n-party r-round (p, γ, δ)-dense communication k-bounded-collusion protocol ($(n, r, p, \gamma, \delta, k)$-DC-BCP, for short) if it satisfies the same properties as in Definition 12, except that property (iv) is replaced by*

(iv') for each $j \in [r]$, the function $f_j : \mathcal{X}^k \to \{0,1\}^*$ is (p, γ, δ_j)-dense leakage for $X_{\mathcal{S}_j} | \tau_{j-1}$, where τ_j denotes the transcript up to the j-th round and $0 < \delta_j \leq 1$, and where additionally $\prod_{j=1}^{r} \delta_j \geq \delta$.

Definition 14 (Noisy-communication BCPs) *An interactive (possibly randomized) protocol π is called an n-party r-round ℓ-noisy communication k-bounded-collusion protocol ((n, r, ℓ, k)-NC-BCP, for short) if it satisfies the same properties as in Definition 12, except that property (iv) is replaced by*

(iv") for each $j \in [r]$, the function $f_j : \mathcal{X}^k \to \{0,1\}^*$ is ℓ_j-U-noisy leakage for $X_{\mathcal{S}_j}$, where $\ell_j \geq 0$ and additionally $\left\lceil \sum_{j=1}^{r} \ell_j \right\rceil \leq \ell$.

Observe that the number of rounds in a DC-BCP or NC-BCP is unbounded. Also, note that property (iv") in Definition 14 implicitly implies that the overall leakage drops the min-entropy of the uniform distribution over any subset of k inputs by at most ℓ. More formally, the final transcript τ is such that[16] for all subsets $\mathcal{S} \in \binom{[n]}{k}$ we have

$$\widetilde{\mathbb{H}}_\infty(U_{\mathcal{S}} | \pi(U_1, \ldots, U_n)) \geq \mathbb{H}_\infty(U_{\mathcal{S}}) - \ell, \tag{4}$$

where $U = (U_1, \ldots, U_n)$ is uniform over $\mathcal{X}^n$ and $\pi(U_1, \ldots, U_n)$ denotes the distribution of the transcript τ at the end of the protocol.

Clearly, any BC-BCP is also a NC-BCP with the same leakage parameter. Below, we show that the converse is also true, in the sense that the transcript of any NC-BCP π can be simulated using the transcript of a related BC-BCP π', up to a small statistical distance. In fact, the latter statement holds true for the more general case of DC-BCPs.

[16] This is because, by [20, Lemma L.3], any sequence of (adaptively chosen) functions $f_1, \ldots, f_r$ on a random variable X, such that each function f_j is ℓ_j-ME-noisy leakage for some $\ell_j \geq 0$ and where $\sum_{j=1}^{r} \ell_j \leq \ell$, satisfies $\widetilde{\mathbb{H}}_\infty(X | f_1(X), \ldots, f_r(X)) \geq \mathbb{H}_\infty(X) - \ell$. Furthermore, for the case of NC-BCPs, in the worst case all the leakage happens on the same subset $\mathcal{S}$ of inputs.

Theorem 11 *Let π be an $(n, r, p, \gamma, \delta, k)$-DC-BCP. There exists an (n, r, ℓ', k)-BC-BCP π' such that, for any $\varepsilon > 0$, a transcript of π can be simulated within statistical distance $r \cdot (\varepsilon + \varepsilon^{1/4} + \gamma + p)$ given a transcript of π' with length $\ell' = \log(1/\delta) + r \cdot (\log\log(1/\varepsilon) + 2\log(1/(1-\gamma)) + 2)$.*

Proof. We start by describing protocol π' acting on a random variable $X = (X_1, \ldots, X_n)$. Consider the simulator Sim_f guaranteed by Theorem 3.

- Let τ' be initially empty, and sample r independent random tapes $\rho_1, \ldots, \rho_r$ for Sim.
- At each round $j \in [r]$, the function Next' takes as input the current transcript $\tau' = z'_1 || \ldots || z'_{j-1}$ and runs $\mathsf{Next}(\tilde{\tau})$, where

$$\tilde{\tau} = \mathsf{Sim}_{f_1}(z'_1; \rho_1) || \ldots || \mathsf{Sim}_{f_{j-1}}(z'_{j-1}; \rho_{j-1}).$$

- Let $(f_j, \mathcal{S}_j)$ be the j-th output of Next. Then, Next' runs Sim_{f_j} on $X_{\mathcal{S}_j} | \tilde{\tau}$ (with fixed random tape ρ_j), obtaining a leakage function $f'_j : \mathcal{X}^k \to \{0,1\}^{\ell'_j}$, and outputs $(f'_j, \mathcal{S}_j)$.

Next, we claim that protocol π' has ℓ'-bounded communication for ℓ' as in the statement of the theorem. Recall that, for each $j \in [r]$, the function f_j output by Next is (p, γ, δ_j)-dense leakage for $X_{\mathcal{S}_j} | \tilde{\tau}$, with $0 < \delta_j \leq 1$. Then, by applying Theorem 3, for any $\varepsilon > 0$ we get that $\ell'_j = \log(1/\delta_j) + \log\log(1/\varepsilon) + 2\log(1/(1-\gamma)) + 2$. Hence, the final transcript τ' has size at most $\ell' = \sum_{j=1}^{r} \ell'_j = \log(1/\delta) + r \cdot (\log\log(1/\varepsilon) + 2\log(1/(1-\gamma)) + 2)$, which is the bound in the statement of the theorem.

It remains to prove that we can simulate a transcript of π given a transcript of π'. Consider the simulator that, after running π' with random tapes $\rho_1, \ldots, \rho_r$, obtains the transcript $\tau' = z'_1 || \ldots || z'_r$ and simply outputs the simulated transcript $\tilde{\tau} = \mathsf{Sim}_{f_1}(z'_1; \rho_1) || \ldots || \mathsf{Sim}_{f_r}(z'_r; \rho_r)$. By a hybrid argument, Theorem 3 implies that the transcript $\tilde{\tau}$ is within statistical distance at most $r \cdot (\varepsilon + \gamma + p) + \sum_{j=1}^{r} \varepsilon^{1/4\delta_j} \leq r \cdot (\varepsilon + \varepsilon^{1/4} + \gamma + \delta)$ from the transcript τ obtained by running π. This finishes the proof. □

Theorem 12 *Let π be an (n, r, ℓ, k)-NC-BCP. There exists an (n, r, ℓ', k)-BC-BCP π' such that, for any $0 < \delta < 1$, a transcript of π can be simulated within statistical distance $r \cdot 3\delta$ given a transcript of π' with length $\ell' \leq \ell + r \cdot (6 + 2\log(1/\delta) + \log\log(1/\delta))$.*

Next, we show that Theorem 11 and Theorem 12 have applications to communication complexity lower bounds, and to constructing cryptographic primitives with adaptive noisy-leakage resilience (i.e., where leakage resilience is modeled either as a NC-BCP or as a DC-BCP).

Communication complexity lower bounds We say that an (n, r, ℓ, k)-BCP π (with either bounded or noisy communication) ε-computes a (deterministic) Boolean function $\phi : \mathcal{X}^n \to \{0,1\}$, if there exists an unbounded predictor P that, after running a BCP protocol π on the parties' inputs yielding a final transcript τ, outputs $\phi(X_1, \ldots, X_n)$ with probability at least $1/2 + \varepsilon$ (over the randomness of $(X_i)_{i\in[n]}$, π and P). The theorem below says that for any NC-BCP π that computes a Boolean function ϕ there is a BC-BCP π' that computes the same function with roughly the same probability, where the size ℓ' of a transcript of π' is related to the leakage parameter ℓ of π.

Corollary 5 *Let π be any (n, r, ℓ, k)-NC-BCP that ε-computes a Boolean function ϕ. Then, there exists an (n, r, ℓ', k)-BC-BCP π' that ε'-computes ϕ so long as $\ell' \leq \ell + r \cdot (6 + 2\log(6r/\varepsilon) + \log\log(6r/\varepsilon))$ and $\varepsilon' = \varepsilon/2$.*

The above corollary can be used to translate known lower bounds in communication complexity for BC-BCPs to the more general setting of NC-BCPs.[17] Note that a lower bound on the communication complexity of BC-BCPs does not necessarily imply a lower bound on the noisiness of NC-BCPs, as the fact that the transcript must consist of *at least* ℓ bits does not say anything about how each round function reveals on the players' inputs. We argue how the result from Corollary 5 can be used to lift lower bounds on bounded communication needed to compute certain functions ϕ to more general lower bounds on *noisy* communication in [13].

Remark 6 (On lower bounds on the leakage parameters of NC-BCPs). It may seem that a lower bound on the parameter ℓ of NC-BCPs does not necessarily mean that any protocol must reveal a lot of information on the parties' inputs, as the actual min-entropy drop in Eq. (4) could be much smaller[18] than ℓ. Nevertheless, we observe that the definition of NC-BCP implies that there must exist an index $j^* \in [r]$ such that, say, $\ell_{j^*} \geq \frac{\ell-1}{r}$. This is because, if $\ell_j < \frac{\ell-1}{r}$ for all $j \in [r]$, then $\lceil \sum_{j=1}^{r} \ell_j \rceil \leq \ell - 1$. In this light, the corollaries below still say that, for certain Boolean functions, a transcript must necessarily reveal enough information about the inputs so long as the number of rounds is not too large.

BCP Leakage Resilience Finally, we show how to lift bounded-leakage resilience to dense-leakage and uniform-noisy-leakage resilience in applications where the leakage itself is modelled as a BCP protocol. For concreteness, we focus again on secret sharing schemes and refer the reader to [13] for additional examples.

Let $\Sigma = (\mathsf{Share}, \mathsf{Rec})$ be a secret sharing scheme. The definition below captures security of Σ in the presence of an adversary leaking information jointly from subsets of the shares of size $k < n$, where both the leakage functions and the subsets of shares are chosen adaptively. For simplicity, we focus on threshold secret sharing but our treatment can be generalized to arbitrary access structures.

[17] In fact, using Theorem 11, we could also derive lower bounds on DC-BCPs. However, we stick to the setting of NC-BCPs for simplicity.

[18] For instance, take $k = 1$ and consider the functions $f_1, \ldots, f_n$ that always reveal the first bit of X_1. Then, $\ell = \sum_{j=1}^{n} \ell_j = n$, but $\widetilde{\mathbb{H}}_\infty(U_1|\pi(U_1, \ldots, U_n)) = \mathbb{H}_\infty(U_1) - 1$.

Definition 15 (Secret sharing with BCP leakage resilience) *Let $t, n, \ell \in \mathbb{N}, \varepsilon \in [0,1]$ be parameters. A t-out-of-n secret sharing scheme* $(\mathsf{Share}, \mathsf{Rec})$ *is a k-joint r-adaptive (p, γ, δ)-dense ε-leakage-resilient secret sharing scheme, $(k, r, p, \gamma, \delta, \varepsilon)$-JA-DLR-SS for short, if for all messages $y_0, y_1 \in \{0,1\}^m$ and all $(n, r, p, \gamma, \delta, k)$-DC-BCP π we have $\pi(X_1^{(0)}, \ldots, X_n^{(0)}) \approx_\varepsilon \pi(X_1^{(1)}, \ldots, X_n^{(1)})$, where $(X_1^{(b)}, \ldots, X_n^{(b)}) = \mathsf{Share}(y_b)$ is the distribution of the shares of message $y_b \in \{0,1\}^m$ for all $b \in \{0,1\}$.*

Moreover, in case the protocol π in the above definition is an:

- (n, r, ℓ, k)-NC-BCP, we say that Σ is k-joint r-adaptive ℓ-noisy ε-leakage-resilient (or $(k, r, \ell, \varepsilon)$-JA-NLR-SS);
- (n, r, ℓ, k)-BC-BCP, we say that Σ is k-joint r-adaptive ℓ-bounded ε-leakage-resilient (or $(k, r, \ell, \varepsilon)$-JA-BLR-SS).

Corollary 6 *Every $(k, r, \ell, \varepsilon)$-JA-BLR-SS scheme Σ is also a $(k, r, p, \gamma, \delta, \varepsilon')$-JA-DLR-SS so long as $\ell = \log(1/\delta) + r \cdot (\log\log(1/\varepsilon) + 2\log(1/(1-\gamma)) + 2)$ and $\varepsilon' = \varepsilon + 2r \cdot (\varepsilon + \varepsilon^{1/4} + \gamma + p)$.*

Corollary 7 *Every $(k, r, \ell', \varepsilon')$-JA-BLR-SS scheme Σ is also a $(k, r, \ell, \varepsilon)$-JA-NLR-SS scheme so long as $\ell' = \ell + r \cdot O(\log(r/\varepsilon))$ and $\varepsilon = 3\varepsilon'$.*

Explicit constructions of secret sharing schemes with BCP leakage resilience in the bounded leakage setting can be built for any leakage bound ℓ and any $\varepsilon > 0$ from n-party functions with large NOF complexity with collusion bound $k = O(\log(n))$ [38] (for arbitrary access structures) and $k = O(t/\log(t))$ [39] (for threshold access structures). By the above corollaries, these schemes are also directly secure in the settings of dense and U-noisy leakage.

6 Conclusions and Open Problems

We have shown that a single query of *noisy* leakage can be simulated in the information-theoretic setting using a single query of *bounded* leakage, up to a small statistical distance and at the price of a slight loss in the leakage parameter. The latter holds true for a fairly general class of noisy leakage (which we introduce) dubbed *dense* leakage. Importantly, dense leakage captures many already existing noisy-leakage models including those where the noisiness of the leakage is measured using the conditional average min-entropy [20, 44], the statistical distance [24], or the mutual information [46]. For some of these models, our simulation theorems require additional assumptions on the input distribution or only hold for certain range of parameters, but in each case we show this is the best one can hope for.

The above result has applications to leakage-resilient cryptography, where we can reduce noisy-leakage resilience to bounded-leakage resilience in a black-box way. Interestingly, for some applications, the latter holds true even in the *computational* setting. Additionally, we have shown that our simulation theorems yield new lower bounds in communication complexity.

Several interesting open questions remain. We list some of them below:

- Can we prove that other families of noisy leakage (e.g., hard-to-invert leakage [21]) fall within the class of dense leakage (or directly admit simulation theorems with good parameters from bounded leakage)?
- Can we make the simulator efficient for certain families of noisy leakage? The latter would allow to lift bounded-leakage resilience to noisy-leakage resilience for *all* computationally-secure applications, and for statistically-secure applications with simulation-based security in which the running time of the simulator needs to be polynomial in the running time of the adversary (such as leakage-tolerant MPC [10]).
- Can we generalize Theorem 12 to a more general setting where the leakage parameter ℓ of NC-BCPs measures the worst-case average min-entropy drop w.r.t. the final transcript of the protocol (instead of being the summation over the worst-case min-entropy drops of each round function in isolation)? The latter would allow to strengthen the lower bounds in §5.2, as well as the security of the applications in §5.2 and [13].

Acknowledgments

We thank François-Xavier Standaert for bringing [25, 31] to our attention. Maciej Obremski was funded by the Singapore Ministry of Education and the National Research Foundation under grant R-710-000-012-135. Daniele Venturi was partially supported by the research project SPECTRA funded by Sapienza University of Rome.

References

1. Aggarwal, D., Damgård, I., Nielsen, J.B., Obremski, M., Purwanto, E., Ribeiro, J., Simkin, M.: Stronger leakage-resilient and non-malleable secret sharing schemes for general access structures. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 510–539. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_18
2. Aggarwal, D., Dziembowski, S., Kazana, T., Obremski, M.: Leakage-resilient non-malleable codes. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 398–426. Springer, Heidelberg (Mar 2015). https://doi.org/10.1007/978-3-662-46494-6_17
3. Aggarwal, D., Kazana, T., Obremski, M.: Leakage-resilient algebraic manipulation detection codes with optimal parameters. In: IEEE International Symposium on Information Theory. pp. 1131–1135 (2018)
4. Agrawal, D., Archambeault, B., Rao, J.R., Rohatgi, P.: The EM side-channel(s). In: Kaliski Jr., B.S., Koç, Çetin Kaya., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 29–45. Springer, Heidelberg (Aug 2003). https://doi.org/10.1007/3-540-36400-5_4
5. Agrawal, S., Dodis, Y., Vaikuntanathan, V., Wichs, D.: On continual leakage of discrete log representations. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 401–420. Springer, Heidelberg (Dec 2013). https://doi.org/10.1007/978-3-642-42045-0_21

6. Ahmadi, H., Safavi-Naini, R.: Detection of algebraic manipulation in the presence of leakage. In: Padró, C. (ed.) ICITS 13. LNCS, vol. 8317, pp. 238–258. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-319-04268-8_14
7. Ananth, P., Jain, A., Sahai, A.: Indistinguishability obfuscation for turing machines: Constant overhead and amortization. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 252–279. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63715-0_9
8. Babai, L., Nisan, N., Szegedy, M.: Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. J. Comput. Syst. Sci. **45**(2), 204–232 (1992)
9. Benhamouda, F., Degwekar, A., Ishai, Y., Rabin, T.: On the local leakage resilience of linear secret sharing schemes. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 531–561. Springer, Heidelberg (Aug 2018). https://doi.org/10.1007/978-3-319-96884-1_18
10. Bitansky, N., Canetti, R., Halevi, S.: Leakage-tolerant interactive protocols. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 266–284. Springer, Heidelberg (Mar 2012). https://doi.org/10.1007/978-3-642-28914-9_15
11. Bourgain, J.: More on the sum-product phenomenon in prime fields and its applications. International Journal of Number Theory **1**(1), 1–32 (2005)
12. Brakerski, Z., Kalai, Y.T., Katz, J., Vaikuntanathan, V.: Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In: 51st FOCS. pp. 501–510. IEEE Computer Society Press (Oct 2010). https://doi.org/10.1109/FOCS.2010.55
13. Brian, G., Faonio, A., Obremski, M., Ribeiro, J., Simkin, M., Skórski, M., Venturi, D.: The mother of all leakages: How to simulate noisy leakages via bounded leakage (almost) for free. Cryptology ePrint Archive, Report 2020/1246 (2020), https://eprint.iacr.org/2020/1246
14. Brian, G., Faonio, A., Obremski, M., Simkin, M., Venturi, D.: Non-malleable secret sharing against bounded joint-tampering attacks in the plain model. Cryptology ePrint Archive, Report 2020/725 (2020), https://eprint.iacr.org/2020/725
15. Chandra, A.K., Furst, M.L., Lipton, R.J.: Multi-party protocols. In: 15th ACM STOC. pp. 94–99. ACM Press (Apr 1983). https://doi.org/10.1145/800061.808737
16. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 398–412. Springer, Heidelberg (Aug 1999). https://doi.org/10.1007/3-540-48405-1_26
17. Chattopadhyay, E., Goodman, J., Goyal, V., Li, X.: Leakage-resilient extractors and secret-sharing against bounded collusion protocols. Cryptology ePrint Archive, Report 2020/478 (2020), https://eprint.iacr.org/2020/478
18. Chung, F.R.K.: Quasi-random classes of hypergraphs. Random Struct. Algorithms **1**(4), 363–382 (1990)
19. Davì, F., Dziembowski, S., Venturi, D.: Leakage-resilient storage. In: Garay, J.A., Prisco, R.D. (eds.) SCN 10. LNCS, vol. 6280, pp. 121–137. Springer, Heidelberg (Sep 2010). https://doi.org/10.1007/978-3-642-15317-4_9
20. Dodis, Y., Haralambiev, K., López-Alt, A., Wichs, D.: Cryptography against continuous memory attacks. In: 51st FOCS. pp. 511–520. IEEE Computer Society Press (Oct 2010). https://doi.org/10.1109/FOCS.2010.56
21. Dodis, Y., Kalai, Y.T., Lovett, S.: On cryptography with auxiliary input. In: Mitzenmacher, M. (ed.) 41st ACM STOC. pp. 621–630. ACM Press (May / Jun 2009). https://doi.org/10.1145/1536414.1536498

22. Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.D.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. SIAM J. Comput. **38**(1), 97–139 (2008)
23. Duc, A., Dziembowski, S., Faust, S.: Unifying leakage models: From probing attacks to noisy leakage. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 423–440. Springer, Heidelberg (May 2014). https://doi.org/10.1007/978-3-642-55220-5_24
24. Duc, A., Dziembowski, S., Faust, S.: Unifying leakage models: From probing attacks to noisy leakage. Journal of Cryptology **32**(1), 151–177 (Jan 2019). https://doi.org/10.1007/s00145-018-9284-1
25. Duc, A., Faust, S., Standaert, F.X.: Making masking security proofs concrete (or how to evaluate the security of any leaking device), extended version. Journal of Cryptology **32**(4), 1263–1297 (Oct 2019). https://doi.org/10.1007/s00145-018-9277-0
26. Dziembowski, S.: Intrusion-resilience via the bounded-storage model. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 207–224. Springer, Heidelberg (Mar 2006). https://doi.org/10.1007/11681878_11
27. Dziembowski, S., Faust, S., Skorski, M.: Noisy leakage revisited. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 159–188. Springer, Heidelberg (Apr 2015). https://doi.org/10.1007/978-3-662-46803-6_6
28. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 49th FOCS. pp. 293–302. IEEE Computer Society Press (Oct 2008). https://doi.org/10.1109/FOCS.2008.56
29. Faust, S., Rabin, T., Reyzin, L., Tromer, E., Vaikuntanathan, V.: Protecting circuits from leakage: the computationally-bounded and noisy cases. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 135–156. Springer, Heidelberg (May / Jun 2010). https://doi.org/10.1007/978-3-642-13190-5_7
30. Faust, S., Rabin, T., Reyzin, L., Tromer, E., Vaikuntanathan, V.: Protecting circuits from computationally bounded and noisy leakage. SIAM J. Comput. **43**(5), 1564–1614 (2014)
31. Fuller, B., Hamlin, A.: Unifying leakage classes: Simulatable leakage and pseudoentropy. In: Lehmann, A., Wolf, S. (eds.) Information Theoretic Security. pp. 69–86. Springer International Publishing, Cham (2015)
32. Goyal, V., Ishai, Y., Maji, H.K., Sahai, A., Sherstov, A.A.: Bounded-communication leakage resilience via parity-resilient circuits. In: Dinur, I. (ed.) 57th FOCS. pp. 1–10. IEEE Computer Society Press (Oct 2016). https://doi.org/10.1109/FOCS.2016.10
33. Harsha, P., Ishai, Y., Kilian, J., Nissim, K., Venkatesh, S.: Communication versus computation. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 745–756. Springer, Heidelberg (Jul 2004). https://doi.org/10.1007/978-3-540-27836-8_63
34. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_27
35. Kalai, Y.T., Reyzin, L.: A survey of leakage-resilient cryptography. In: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, pp. 727–794. ACM (2019)
36. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) CRYPTO'96. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (Aug 1996). https://doi.org/10.1007/3-540-68697-5_9

37. Kocher, P.C., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (Aug 1999). https://doi.org/10.1007/3-540-48405-1_25
38. Kumar, A., Meka, R., Sahai, A.: Leakage-resilient secret sharing against colluding parties. In: Zuckerman, D. (ed.) 60th FOCS. pp. 636–660. IEEE Computer Society Press (Nov 2019). https://doi.org/10.1109/FOCS.2019.00045
39. Kumar, A., Meka, R., Zuckerman, D.: Bounded collusion protocols, cylinder-intersection extractors and leakage-resilient secret sharing. Cryptology ePrint Archive, Report 2020/473 (2020), https://eprint.iacr.org/2020/473
40. Li, X., Ma, F., Quach, W., Wichs, D.: Leakage-resilient key exchange and two-seed extractors. Cryptology ePrint Archive, Report 2020/771 (2020), https://eprint.iacr.org/2020/771
41. Lin, F., Cheraghchi, M., Guruswami, V., Safavi-Naini, R., Wang, H.: Leakage-resilient secret sharing in non-compartmentalized models. In: Kalai, Y.T., Smith, A.D., Wichs, D. (eds.) 1st Conference on Information-Theoretic Cryptography (ITC 2020). Leibniz International Proceedings in Informatics (LIPIcs), vol. 163, pp. 7:1–7:24. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2020). https://doi.org/10.4230/LIPIcs.ITC.2020.7
42. Lin, F., Safavi-Naini, R., Wang, P.: Detecting algebraic manipulation in leaky storage systems. In: Nascimento, A.C.A., Barreto, P. (eds.) ICITS 16. LNCS, vol. 10015, pp. 129–150. Springer, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-319-49175-2_7
43. Micali, S., Reyzin, L.: Physically observable cryptography (extended abstract). In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Heidelberg (Feb 2004). https://doi.org/10.1007/978-3-540-24638-1_16
44. Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. SIAM J. Comput. **41**(4), 772–814 (2012)
45. Phillips, J.M., Verbin, E., Zhang, Q.: Lower bounds for number-in-hand multiparty communication complexity, made easy. In: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012. pp. 486–501 (2012)
46. Prest, T., Goudarzi, D., Martinelli, A., Passelègue, A.: Unifying leakage models on a Rényi day. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 683–712. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26948-7_24
47. Prouff, E., Rivain, M.: Masking against side-channel attacks: A formal security proof. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 142–159. Springer, Heidelberg (May 2013). https://doi.org/10.1007/978-3-642-38348-9_9
48. Srinivasan, A., Vasudevan, P.N.: Leakage resilient secret sharing and applications. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 480–509. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_17
49. Yao, A.C.: Some complexity questions related to distributive computing (preliminary report). In: Proceedings of the 11h Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA. pp. 209–213 (1979)