Stronger Lower Bounds for Leakage-Resilient Secret Sharing

Charlotte Hoffmann[1]* and Mark Simkin[2]**

[1] Institute of Science and Technology Austria
[2] Ethereum Foundation

Abstract. Threshold secret sharing allows a dealer to split a secret s into n shares, such that any t shares allow for reconstructing s, but no $t-1$ shares reveal any information about s. Leakage-resilient secret sharing requires that the secret remains hidden, even when an adversary additionally obtains a limited amount of leakage from every share. Benhamouda et al. (CRYPTO'18) proved that Shamir's secret sharing scheme is one bit leakage-resilient for reconstruction threshold $t \geq 0.85n$ and conjectured that the same holds for $t = c \cdot n$ for any constant $0 \leq c \leq 1$. Nielsen and Simkin (EUROCRYPT'20) showed that this is the best one can hope for by proving that Shamir's scheme is not secure against one-bit leakage when $t = c \cdot n / \log(n)$.

In this work, we strengthen the lower bound of Nielsen and Simkin. We consider noisy leakage-resilience, where a random subset of leakages is replaced by uniformly random noise. We prove a lower bound for Shamir's secret sharing, similar to that of Nielsen and Simkin, which holds even when a constant fraction of leakages is replaced by random noise. To this end, we first prove a lower bound on the share size of any noisy-leakage-resilient sharing scheme. We then use this lower bound to show that there exist universal constants c_1, c_2, such that for sufficiently large n it holds that Shamir's secret sharing scheme is not noisy-leakage-resilient for $t \leq c_1 \cdot n / \log(n)$, even when a c_2 fraction of leakages are replaced by random noise.

Keywords: Threshold secret sharing · Noisy leakage-resilience · Lower bounds · Shamir's secret sharing scheme.

1 Introduction

Threshold secret sharing was introduced by Shamir [Sha79] and Blakley [Bla79] and allows a dealer to split a secret s into shares $\mathsf{sh}_1, \ldots, \mathsf{sh}_n$, such that any t shares allow for reconstructing s, but no $t-1$ shares reveal anything about s at all in the information-theoretic sense. Since its introduction, this primitive in general, and Shamir's secret sharing scheme in particular, has found countless applications in various fields of cryptography. Naturally, it is important to understand the precise security it provides.

The security definitions for regular threshold secret sharing schemes and variants like robust [RB89] or verifiable secret sharing [CGMA85] all assume that some shares are fully known and some shares are fully hidden from the adversary. As it turns out, these all-or-nothing type of security models do not always precisely reflect the security we want in practice. Real-world implementations of cryptographic primitives are susceptible to different types of side-channel attacks, which may give the adversary limited access to secrets that should ideally be fully hidden from her. Cryptographic primitives have, for example, been successfully attacked through leakages obtained via timing [Koc96] and power consumption [KJJ99] side-channels.

Motivated by the emergence of such side-channel attacks, the security definitions of secret sharing have been strengthened to account for additional leakages from the shares that were previously assumed to be fully hidden. Such schemes require that the secret that is shared remains hidden, when the adversary not only receives $t-1$ shares, but additionally obtains some limited amount of leakage from *all* other shares. Leakage-resilient secret sharing schemes have received significant interest and many constructions have been proposed over the past few years [DP07, BGK14, GK18b, GK18a, ADN+19, KMS19, SV19, CKOS21, CKOS22].

* charlotte.hoffmann@ist.ac.at
** mark.simkin@ethereum.org

Realistically, however, it seems unlikely that Shamir's secret sharing scheme will be replaced by a leakage-resilient alternative any time soon. Shamir's scheme is a cornerstone of many cryptographic constructions and has been implemented and deployed as part of many different projects. Replacing a scheme that is so deeply embedded into so many different projects, seems like a insurmountable challenge. For this reason, it is crucially important to understand the leakage-resilience of Shamir's secret sharing scheme itself.

Benhamouda et al. [BDIR18] studied this question in a setting, where the adversary submits arbitrary leakage functions $\textsc{Leak}_1, \ldots, \textsc{Leak}_n$ and obtains leakages $\textsc{Leak}_i(\mathsf{sh}_i)$ for $i \in [n]$. The only restriction imposed on the leakage functions is that they are having a bounded output length. The authors show that Shamir's scheme provides some leakage-resilience, when $t \geq 0.85n$ and they conjecture that Shamir's scheme is leakage-resilient against one bit leakages for any $t = c \cdot n$, where $0 \leq c \leq 1$ is a constant. Subsequently, Nielsen and Simkin [NS20] showed that Shamir's scheme is not secure against one bit leakages when $t = c \cdot n / \log(n)$, thereby showing that their conjecture is the best one can hope for.

1.1 Our Contribution

The works of Benhamouda et al. [BDIR18] and Nielsen and Simkin [NS20] assume that the adversary is able to obtain the precise outputs of its leakage functions. In practice, however, side-channel attacks are inherently noisy and there are practical techniques that can amplify this noise [CCD00, CK09, MOP07, CJRR99] to counter potential side-channel attacks. One might hope that it is possible to circumvent the lower bound of Nielsen and Simkin by considering a weaker, but more realistic noisy leakage model, where some random subset of the leakages is replaced by uniformly random noise.

In this work we show that this is *not* the case. We prove a lower bound similar to that of Nielsen and Simkin for Shamir's secret sharing scheme, which holds even when a *constant fraction* of leakages is replaced by random noise. To this end, we first prove a lower bound on the share size of any noisy-leakage-resilient secret sharing scheme. We then use this lower bound to obtain the following theorem:

Theorem 1 (Informal). *There exist universal constants*c_1, c_2*, such that for sufficiently large* n*, it holds that* $(c_1 \cdot n / \log(n))$*-out-of-*n *Shamir secret sharing is not leakage-resilient against one bit leakage, even when a* c_2 *fraction of the leakage function outputs are replaced by random noise.*

To prove this lower bound, we construct a generic adversary $\mathcal{A}$ that can use the noisy leakage to recover the secret shared value, whenever the shares are too small in size. The main idea of this attack is similar to the one in the proof of [NS20, Theorem 2]. We apply a separate uniformly random leakage function to each share. Given the noisy leakage vector, our adversary $\mathcal{A}$ iterates over all possible secret values and all possible secret sharings thereof.[3] Whenever there is a vector of shares that would produce a leakage vector that is consistent enough with the obtained noisy leakage vector, the adversary remembers the corresponding secret value in an initially empty set S. Finally, the adversary hopes that S contains exactly one element in which case she returns that element as her guess for what was the actual secret shared value.

In contrast to the previous lower bound of Nielsen and Simkin, our adversary needs to account for the noise in the leakage vector and thus it needs to add values s to S, even if there was no secret sharing of s that produced a fully consistent vector of leakages. Relaxing the conditions under which values s are added to S needs to be done carefully, since we would like to ensure that we do not add too many elements to the S. In a nutshell, our lower bound shows that the noisy leakage vector and any other noiseless leakage vector belonging to the incorrect secret, will differ in many positions. Making this intuition formal and arguing that our adversary is successful with a sufficiently high probability requires a careful analysis, which is the main contribution of this work.

1.2 Other Related Works

The work by Guruswami and Wootters [GW16] demonstrated that some linear secret sharing schemes, such as Shamir's scheme over certain fields, allow for very communication efficient reconstruction of the secret.

[3] We are only concerned with information-theoretic security in which case the adversary is computationally unbounded.

More precisely, they show that Shamir's scheme over fields of characteristic two, allows for recovering a multi-bit secret from only one bit of leakage from each share.

Inspired by these results, Benhamouda et al. [BDIR18] investigate to what extend natural secret sharing schemes offer leakage-resilience. They prove that Shamir's secret sharing scheme is leakage-resilient against one bit leakages, when the reconstruction thresholds is at least 0.92 times the number of parties. This constant was then improved to 0.8675 [MPSW21], then to 0.85 [BDIR21] and later to 0.78 [MNPCW22].

The currently best known constant is 0.69, which was recently proven by Klein and Komargodski [KK23]. The authors additionally show that whenever the leakage functions are guaranteed to be *balanced*, i.e. approximately half of the domain gives output 1 and the other half gives output -1, then the constant can be reduced to 0.58. Similarly, whenever the leakage functions are guaranteed to be sufficiently *unbalanced*, then Shamir's scheme is leakage resilient as long as the reconstruction threshold is at least 0.01 times the number of parties. This result is the first one that breaks the barrier of 0.5, which was known to be inherent in the proof techniques used in the previous works.

Maji et al. [MNP$^+$21] consider much weaker *physical-bit leakages*, which only allows for a fixed number of bits to be leaked from the binary representation of each secret share. They prove that Shamir's secret sharing scheme with random evaluation points is physical-bit leakage resilient if the order of the field is sufficiently large. Adams et al. [AMN$^+$21] consider *noisy* physical-bit leakage, where each physical-bit leakage is replaced by noise with some fixed probability. They prove a lower bound for the reconstruction threshold of $\log(\lambda)/\log\log(\lambda)$ for Shamir's secret sharing scheme, when the size of the field is 2^λ and the evaluation points can be chosen adversarially. In [MNPC$^+$22] Maji et al. improve their lower bound to $\log(\lambda)$. This bound is interesting in the setting where the size of the field is much larger than the number of parties. In the setting we consider, we have $\lambda \approx \log(n)$, in which case their lower says that the reconstruction threshold needs to be larger than $\log\log(n)$.

In another work, Maji et al. [MNP$^+$22] consider *global leakage* functions with bounded output length that can compute arbitrary functions over all shares simultaneously. Generally, this would allow the leakage functions to just reconstruct the secret, which is an attack that cannot be prevented. For this reason, the authors artificially restrict their leakage functions to not depend on some of the random choices made by the secret sharing scheme. For the case of Shamir secret sharing with random evaluation points, the authors show that one obtains some leakage-resilience properties, if the leakage functions are not allowed to depend on the evaluation points.

2 Preliminaries

Notation. We write $[n]$ to denote the set $\{1, \ldots, n\}$. For a set X, we write $x \leftarrow X$ to denote the process of sampling a uniformly random element x from the set X. For a vector $v = (v_1, \ldots, v_n)$ and a vector $w = (w_1, \ldots, w_t) \in [n]^t$, we define $v_w := (v_{w_1}, \ldots, v_{w_t})$. We will sometimes abuse notation and write v_w, where w is a set, rather than a vector. In this case the elements can be ordered arbitrarily in the vector. We denote by $\textsc{Noise}(v, \ell, p)$ the algorithm that takes vector $v = (v_1, \ldots, v_n) \in \left(\{0,1\}^\ell\right)^n$, $\ell \in \mathbb{N}$, and $0 \leq p \leq 1$ as input and returns a new vector $(\tilde{v}_1, \ldots, \tilde{v}_n)$, where for $i \in [n]$ each $\tilde{v}_i = v_i$ with probability $1 - p$ and $\tilde{v}_i \leftarrow \{0,1\}^\ell$ with probability p. That means that $\textsc{Noise}(v, \ell, 1)$ returns a uniformly random vector and $\textsc{Noise}(v, \ell, 0)$ returns v.

2.1 Leakage-Resilient Secret Sharing

We define threshold secret sharing schemes similarly to how it was done by Nielsen and Simkin [NS20]. The full reconstruction parameter $\hat{t}$ defines how many shares are needed to reconstruct all shares of a particular secret sharing. Intuitively, $\hat{t}$ corresponds to a crude measure of how much entropy the vector of shares contains.

Definition 1 (Threshold Secret Sharing Scheme). *A t-out-of-n threshold secret sharing scheme is a pair* $(\textsc{Share}, \textsc{Rec})$ *of efficient algorithms. The randomized sharing algorithm* $\textsc{Share} : \{0,1\}^k \rightarrow (\{0,1\}^p)^n$

takes a k-bit secret as input and returns a vector of n secret shares, each p-bits long. The deterministic reconstruction algorithm $\textsc{Rec} : (\{0,1\}^p)^t \to \{0,1\}^k$ *takes t of the shares as input and returns a k-bit string. We require a secret sharing scheme to satisfy the following properties:*

Perfect Correctness: *For $t, n \in \mathbb{N}$ with $t \le n$, any $T \subseteq [n]$ with $|T| = t$ and any $x \in \{0,1\}^k$, it holds that*

$$\Pr[\textsc{Rec}(\textsc{Share}(x)_T) = x] = 1,$$

where the probability is taken over the random coins of $\textsc{Share}$.

Full Reconstruction: $(\textsc{Share}, \textsc{Rec})$ *has $\hat{t}$-full-reconstruction, if for any x, the vector* $\textsc{Share}(x)$ *can be computed from any subvector* $\textsc{Share}(x)_T$ *with $|T| \ge \hat{t}$.*

We assume for simplicity that all shares are of the same size p but the proof of our lower bound can easily be adapted to schemes with shares of different sizes.

Definition 2 (Leakage Functions). *Let* $(\textsc{Share}, \textsc{Rec})$ *with* $\textsc{Share} : \{0,1\}^k \to \bigtimes_{i=1}^{n} \{0,1\}^p$ *be a secret sharing scheme and for $i \in [n]$, let* $\textsc{Leak}_i : \{0,1\}^p \to \{0,1\}^\ell$. *We call* $\textsc{Leak} = (\textsc{Leak}_1, \ldots, \textsc{Leak}_n)$ *an ℓ-leakage function for* $(\textsc{Share}, \textsc{Rec})$. *We define* $\textsc{Leak}(\mathsf{sh}_1, \ldots, \mathsf{sh}_n) := (\textsc{Leak}_1(\mathsf{sh}_1), \ldots, \textsc{Leak}_n(\mathsf{sh}_n))$.

We now define the privacy notion, which is a direct extension of the (noiseless) weak one-way local leakage resilience notion of Nielsen and Simkin [NS20, Definition 5], for which we will prove our lower bounds. The adversary $\mathcal{A}$ obtains a noisy leakage vector and it knows the probability η with which each leakage is replaced by noise. She does, however, not know *which* leakage outputs are replaced by random noise. Our privacy notion requires that $\mathcal{A}$ is not able to learn the secret with probability greater than $1/2$.

Definition 3 (Weak One-Way Noisy Local Leakage-Resilience). *We say a secret sharing scheme* $(\textsc{Share}, \textsc{Rec})$ *is (ℓ, η)-weakly one-way noisy local leakage-resilient ((ℓ,η)-*$\mathsf{WOW\text{-}NLLR}$*), if for any ℓ-leakage function* $\textsc{Leak}$ *and any adversary $\mathcal{A}$, it holds that*

$$\Pr\left[\begin{array}{r} x \leftarrow \{0,1\}^k \\ (\mathsf{sh}_1, \ldots, \mathsf{sh}_n) \leftarrow \textsc{Share}(x) \\ (\textsc{Leak}_1, \ldots, \textsc{Leak}_n) \leftarrow \mathcal{A}(n) \\ (\tilde{b}_1, \ldots, \tilde{b}_n) \leftarrow \textsc{Leak}(\mathsf{sh}_1, \ldots, \mathsf{sh}_n) \\ (b_1, \ldots, b_n) \leftarrow \textsc{Noise}((\tilde{b}_1, \ldots, \tilde{b}_n), \ell, \eta) \\ x' \leftarrow \mathcal{A}(b_1, \ldots, b_n) \end{array} : x' = x\right] \le \frac{1}{2},$$

where the probability is taken over the random coins of $\textsc{Share}$, $\textsc{Noise}$ *and $\mathcal{A}$.*

We note that this is a very weak privacy notion. We only require a form of one-wayness that prevents the adversary from fully recovering the secret shared value and we only require the adversary to be successful with a probability less than $1/2$. Notably, this notion is even weaker than a standard indistinguishability type of notion. Since we are proving a *lower bound*, working with a weaker privacy notion only *strengthens* our lower bounds.

3 Lower Bound

In this section we prove our lower bound on the share size of any threshold secret sharing scheme that satisfies (ℓ, η)-$\mathsf{WOW\text{-}NLLR}$.

Theorem 2. *Let $\mathcal{S}$ =* $(\textsc{Share}, \textsc{Rec})$ *be a t-out-of-n secret sharing scheme with $\hat{t}$-full-reconstruction and shares consisting of p bits each. Let $\ell \ge 1$ and let $0 < \eta \le (n-t)/4n$. If $\mathcal{S}$ is (ℓ, η)-*$\mathsf{WOW\text{-}NLLR}$*, then*

$$p \ge \frac{\ell(n-t)}{\hat{t}} - \frac{4n\eta(\ell + \log(1/\eta)) + 1}{\hat{t}}.$$

Remark 1. We note that the theorem requires $\eta \leq (n-t)/4n$. In principle, our lower bound could be tightened to only require, for instance, $\eta \leq (n-t)/1.1n$ by replacing a single Markov inequality in the proofs with a stronger tail bound. We opted for clarity instead of optimizing the constants in our exposition. Next, we note that $\eta \leq (n-t)/n$ is a sensible restriction. If $\eta > (n-t)/n$ would hold, then with high probability $n-t+1$ leakages would be replaced by random noise. In this case, our adversary could not hope to recover the secret, even if the leakage functions would leak the full shares.

Remark 2. It can be interesting to compare our lower bound to the one of Nielsen and Simkin [NS20]. Their work shows that any secret sharing scheme that satisfies $(\ell, 0)$-$\mathsf{WOW\text{-}NLLR}$, needs to satisfy

$$p \geq \frac{\ell(n-t)}{\hat{t}}.$$

As η approaches 0, our work effectively proves the same lower bound.

Proof (of Theorem 2). Towards proving the theorem statement, we provide a generic attacker that successfully wins the (ℓ, η)-$\mathsf{WOW\text{-}NLLR}$ game against any secret sharing scheme that does not satisfy the constraints on the share size p that are stated in the theorem statement. This adversary works as follows. It picks $\textsc{Leak} = (\textsc{Leak}_1, \ldots, \textsc{Leak}_n)$ by picking each $\textsc{Leak}_i : \{0,1\}^p \to \{0,1\}^\ell$ for $i \in [n]$ uniformly and independently at random. The challenger picks a uniformly random secret s and computes $(\mathsf{sh}_1, \ldots, \mathsf{sh}_n) \leftarrow \textsc{Share}(s)$. Adversary $\mathcal{A}$ submits the ℓ-leakage function $\textsc{Leak}$ to the challenger, who responds with $(b_1, \ldots, b_n)$, where each b_i is either $\textsc{Leak}_i(\mathsf{sh}_i)$ with probability $1-\eta$ or a uniformly random value from $\{0,1\}^\ell$ with probability η. Let N be the number of components that were replaced by uniformly random noise values by the challenger and let $S = \emptyset$. The adversary now iterates over all possible secrets s' and random coins r' to compute

$$(\mathsf{sh}'_1, \ldots, \mathsf{sh}'_n) \leftarrow \textsc{Share}(s'; r')$$

and

$$(b'_1, \ldots, b'_n) \leftarrow \textsc{Leak}(\mathsf{sh}'_1, \ldots, \mathsf{sh}'_n).$$

If $|\{i \in [n] \mid b'_i = b_i\}| \geq n(1-4\eta)$ for some r', then add s' to S. Finally, once $\mathcal{A}$ iterated over all possible secret sharings, if $|S| = 1$, then it outputs that one element in S and in any other case it returns $\perp$.

Let us now analyze the success probability of $\mathcal{A}$. We observe that if the challenger replaced at most $4n\eta$ coordinates by uniformly random noise, i.e. if $N \leq 4n\eta$, then $s \in S$. Since in expectation N is equal to $n\eta$, it holds by Markov's inequality that

$$\begin{aligned}
\Pr[s \in S] = &\Pr[s \in S \mid N \leq 4n\eta] \cdot \Pr[N \leq 4n\eta] \\
&+ \Pr[s \in S \mid N > 4n\eta] \cdot \Pr[N > 4n\eta] \\
= &\Pr[N \leq 4n\eta] + \Pr[s \in S \mid N > 4n\eta] \cdot \Pr[N > 4n\eta] \\
\geq &\Pr[N \leq 4n\eta] \geq 3/4.
\end{aligned}$$

Our adversary is successful, if and only if s is the *only* element in S. Let $E_{s'}$ be the event that $s' \in S$. Then

$$\begin{aligned}
\Pr[S = \{s\}] = \Pr\left[\left(\bigwedge_{s' \neq s} \neg E_{s'}\right) \wedge E_s\right] &\geq \Pr\left[\left(\bigwedge_{s' \neq s} \neg E_{s'}\right) \wedge N \leq 4n\eta\right] \\
&\geq \Pr\left[\left(\bigwedge_{s' \neq s} \neg E_{s'}\right) \mid N \leq 4n\eta\right] \cdot 3/4 \\
&= \left(1 - \Pr\left[\bigvee_{s' \neq s} E_{s'} \mid N \leq 4n\eta\right]\right) \cdot 3/4.
\end{aligned}$$

To prove the theorem statement, we need to show that the adversary's attack is successful with a sufficiently high probability, i.e. we need to show that $\Pr\left[S = \{s\}\right] \geq 1/2$ and thus by the above it suffices to show that

$$\Pr\left[\bigvee_{s' \neq s} E_{s'} \mid N \leq 4n\eta\right] \leq 1/3.$$

By the union bound[4] we have that

$$\Pr\left[\bigvee_{s' \neq s} E_{s'} \mid N \leq 4n\eta\right] \leq \sum_{s' \neq s} \Pr\left[E_{s'} \mid N \leq 4n\eta\right].$$

Let us now fix an arbitrary $s' \neq s$, fix random coins r', and let $(\mathsf{sh}'_1, \ldots, \mathsf{sh}'_n) \leftarrow \textsc{Share}(s'; r')$. Let $E_{s',r'}$ be the event that the adversary includes s' into S based on the leakage from $(\mathsf{sh}'_1, \ldots, \mathsf{sh}'_n)$, i.e. the event that $|\{i \in [n] \mid b'_i = b_i\}| \geq n(1 - 4\eta)$, where $b'_i \leftarrow \textsc{Leak}_i(\mathsf{sh}'_i)$. Let us bound the probability of $E_{s',r'}$ conditioned on $N \leq 4n\eta$. From the perfect correctness of the secret sharing scheme and since $s \neq s'$, we know that there exists a set of indices $I \subseteq [n]$ with $|I| \geq n - t + 1$, such that for all $i \in I$, it holds that $\mathsf{sh}_i \neq \mathsf{sh}'_i$. For each $i \in I$, there are two cases. Either the leakage b_i is the real leakage or it is a uniformly random element from $\{0,1\}^\ell$. In either case, it holds that $b_i = b'_i$ with probability $2^{-\ell}$, since the corresponding shares are different and the leakage function $\textsc{Leak}_i$ is chosen uniformly random and independently of its inputs.

Let $\mathcal{T}$ be the set of subsets of I of size $n-t+1-4n\eta$. Note that $n-t+1-4n\eta > 0$, since $\eta \leq (n-t)/4n$ by assumption. For $T \in \mathcal{T}$, let $E_{s',r',T}$ be the event that the noisy leakage vector $(b_1, \ldots, b_n)$ and the noiseless vector $(b'_1, \ldots, b'_n)$ agree on all coordinates in T. Note that by the union bound

$$\Pr\left[E_{s',r'} \mid N \leq 4n\eta\right] \leq \sum_{T \in \mathcal{T}} \Pr\left[E_{s',r',T} \mid N \leq 4n\eta\right].$$

To see this, observe that even if $(\mathsf{sh}_1, \ldots, \mathsf{sh}_n)$ and $(\mathsf{sh}'_1, \ldots, \mathsf{sh}'_n)$ agree on $t-1$ coordinates, then there must still exist at least $n-t+1-4n\eta$ distinct indices $i \in I$ for which it holds that $b_i = b'_i$ to satisfy the condition $\left|\{j \in [n] \mid b'_j = b_j\}\right| \geq n(1-4\eta)$. It is easy to see that

$$\Pr\left[E_{s',r',T} \mid N \leq 4n\eta\right] \leq 2^{-(n-t+1-4n\eta)\ell}$$

and thus, it holds that

$$\begin{aligned}
\Pr\left[E_{s',r'} \mid N \leq 4n\eta\right] \leq & |\mathcal{T}| \cdot 2^{-(n-t+1-4n\eta)\ell} \\
= & \binom{n-t+1}{n-t+1-4n\eta} \cdot 2^{-(n-t+1-4n\eta)\ell} \\
= & \binom{n-t+1}{4n\eta} \cdot 2^{-(n-t+1-4n\eta)\ell} \\
\leq & \left(\frac{e(n-t+1)}{4n\eta}\right)^{4n\eta} \cdot 2^{-(n-t+1-4n\eta)\ell} \\
\leq & \left(\frac{n-t+1}{n\eta}\right)^{4n\eta} \cdot 2^{-(n-t+1-4n\eta)\ell} \\
\leq & \left(\frac{1}{\eta}\right)^{4n\eta} \cdot 2^{-(n-t+1-4n\eta)\ell}
\end{aligned}$$

[4] The union bound also holds for conditional probabilities, meaning that $\Pr[A \vee B \mid C] = \Pr[(A \vee B) \wedge C] / \Pr[C] \leq (\Pr[A \wedge C] + \Pr[B \wedge C]) / \Pr[C] = \Pr[A \mid C] + \Pr[B \mid C]$.

At this point, recall that each share is p-bits long and that $\hat{t}$ is the full reconstruction threshold, i.e. that any $\hat{t}$ shares are enough to uniquely determine all remaining shares of a specific secret sharing. Thus there are at most $2^{p\hat{t}}$ different secret sharings in total and therefore

$$\sum_{s'\neq s} \Pr\left[E_{s'} \mid N \leq 4n\eta\right] \leq \left(\frac{1}{\eta}\right)^{4n\eta} \cdot 2^{p\hat{t}-(n-t+1-4n\eta)\ell}.$$

As discussed before, the adversary we constructed is successful, if

$$\begin{aligned} &\left(\frac{1}{\eta}\right)^{4n\eta} \cdot 2^{p\hat{t}-(n-t+1-4n\eta)\ell} \leq 1/3 \\ \iff& \log(1/\eta)4n\eta + p\hat{t} - (n-t+1-4n\eta)\ell \leq -\log 3 \\ \iff& p\hat{t} \leq (n-t+1-4n\eta)\ell - \log(1/\eta)4n\eta - \log 3 \\ \iff& p\hat{t} \leq (n-t+1)\ell - \log 3 - 4n\eta(\ell + \log(1/\eta)). \end{aligned}$$

From here it follows that

$$p \geq \frac{(n-t)\ell - 1 - 4n\eta(\ell+\log(1/\eta))}{\hat{t}}$$

must hold, if the secret sharing scheme wants to prevent the attack described above. □

The bound in Theorem 2 can be a little unwieldy and for this reason we also provide a slightly weaker, but simpler to state lower bound in the following corollary.

Corollary 3. *Let $t \leq n/2$. Let $\mathcal{S} = (\textsc{Share}, \textsc{Rec})$ be a t-out-of-n secret sharing scheme with $\hat{t}$-full-reconstruction and shares consisting of p bits each. Let $\ell \geq 1$ and let $0 < \eta \leq 1/64$. If $\mathcal{S}$ is (ℓ, η)-*WOW-NLLR*, then*

$$p \geq \frac{\ell(n-2t)}{2\hat{t}} - 1.$$

Proof. For Theorem 2 to be applicable, it must hold that $0 < \eta \leq (n-t)/4n$, which is always satisfied, when $0 \leq \eta \leq 1/64$, since $t \leq n/2$. Furthermore, it holds that

$$\begin{aligned} \frac{4n\eta(\ell+\log(1/\eta))+1}{\hat{t}} &\leq \frac{n\ell}{16\hat{t}} + \frac{4n\eta\log(1/\eta)}{\hat{t}} + 1 \\ &\leq \frac{n\ell}{16\hat{t}} + \frac{3n}{8\hat{t}} + 1 \leq \frac{7n\ell}{16\hat{t}} + 1 \leq \frac{n\ell}{2\hat{t}} + 1 \end{aligned}$$

From Theorem 2, we know that it must hold that

$$p \geq \frac{\ell(n-t)}{\hat{t}} - \frac{4n\eta(\ell+\log(1/\eta))+1}{\hat{t}}.$$

Thus it must at least hold that

$$\begin{aligned} &p \geq \frac{\ell(n-t)}{\hat{t}} - \frac{n\ell}{2\hat{t}} - 1 \\ \iff& p \geq \frac{\ell(n-2t)}{2\hat{t}} - 1. \end{aligned}$$

□

4 Leakage-Resilience of Shamir's Secret Sharing

In this section we apply our result to Shamir's secret sharing scheme.

4.1 Shamir's Secret Sharing Scheme

In t-out-of-n Shamir secret sharing [Sha79], the secrets are elements of a field $\mathbb{F}_q$ for some prime q, which is chosen as the smallest prime larger than n. To distribute a secret s, the dealer picks a uniformly random polynomial f of degree $t-1$ from $\mathbb{F}_q[X]$ and defines $\mathsf{sh}_i = f(i)$ for $i \in [n]$. Reconstruction of the secret from a subset of t shares is performed via polynomial interpolation, as any polynomial of degree $t-1$ is uniquely defined by t evaluation points.

4.2 Noisy Leakage-Resilience

Benhamouda et al. [BDIR18] conjecture that Shamir's scheme is leakage-resilient against one-bit leakage for any $t = c \cdot n$, where $0 \leq c \leq 1$ is a constant. Nielsen and Simkin [NS20] showed that this is the best one can hope for by proving that the scheme is not secure against one-bit leakage when $t = cn/\log(n)$. In Theorem 4 we show that this lower bound holds even if a constant fraction of leakages is replaced by noise.

Theorem 4. *There exist universal constants c_1, c_2, such that for sufficiently large n, it holds that $(c_1 \cdot n/\log(n))$-out-of-n Shamir secret sharing is not $(1, c_2)$-*$\mathsf{WOW\text{-}NLLR}$.

Proof. Let $c_1 < 1/3$ be arbitrary but fixed and let $c_2 = 1/64$. By Corollary 3, we know that

$$p \geq \frac{n-2t}{2\hat{t}} - 1$$

has to hold for the secret sharing scheme to be $(1, c_2)$-$\mathsf{WOW\text{-}NLLR}$. We note that the full reconstruction threshold $\hat{t} = t$ for Shamir secret sharing, since any t shares allow interpolating any other share. Now plugging in the concrete parameters, we get that

$$\begin{aligned} & p \geq \frac{n - 2c_1 n/\log(n)}{2c_1 n/\log(n)} - 1 \\ \iff & p \geq \frac{\log(n)}{2c_1} - 2 \\ \iff & p \geq \frac{3\log(n)}{2} - 2 \end{aligned}$$

has to hold.

Let q be the first prime larger than n and note that $p = \log(q)$. By the Bertrand-Chebyshev Theorem, we know that $n < q \leq 2n$ and thus it must hold that

$$\begin{aligned} & \log(2n) \geq \frac{3\log(n)}{2} - 2 \\ \iff & 4 \geq 3\log(n) - 2\log(2n) \\ \iff & 4 \geq \log\left(\frac{n^3}{4n^2}\right), \end{aligned}$$

which is clearly not true once n is large enough. □

Conclusion. In this work, we strengthened the lower bounds on the share size of leakage-resilient secret sharing schemes of Nielsen and Simkin [NS20] by showing that similar bounds hold, even if we considerably weaken the security notion we aim for. We show that Shamir secret sharing is not noisy leakage-resilient, if $t \leq c_1 \cdot n/\log(n)$, where c_1 and c_2 are constants, where t is the reconstruction threshold and n is the number of shares. We leave the reader with an interesting open question. Our lower bound crucially relies on an adversary running in exponential time in n. A natural question to consider is whether one can either improve the running time of the adversary to make the attacks more practical or whether one can prove a form of computational leakage-resilience for Shamir secret sharing under an appropriate computational assumption.

References

ADN$^+$19. Divesh Aggarwal, Ivan Damgård, Jesper Buus Nielsen, Maciej Obremski, Erick Purwanto, João Ribeiro, and Mark Simkin. Stronger leakage-resilient and non-malleable secret sharing schemes for general access structures. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 510–539, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany. 1

AMN$^+$21. Donald Q Adams, Hemanta K Maji, Hai H Nguyen, Minh L Nguyen, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Lower bounds for leakage-resilient secret-sharing schemes against probing attacks. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 976–981. IEEE, 2021. 1.2

BDIR18. Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 531–561, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany. 1, 1.1, 1.2, 4.2

BDIR21. Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. *Journal of Cryptology*, 34(2):10, April 2021. 1.2

BGK14. Elette Boyle, Shafi Goldwasser, and Yael Tauman Kalai. Leakage-resilient coin tossing. *Distributed computing*, 27:147–164, 2014. 1

Bla79. G. R. Blakley. Safeguarding cryptographic keys. *Proceedings of AFIPS 1979 National Computer Conference*, 48:313–317, 1979. 1

CCD00. Christophe Clavier, Jean-Sébastien Coron, and Nora Dabbous. Differential power analysis in the presence of hardware countermeasures. In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems – CHES 2000*, volume 1965 of *Lecture Notes in Computer Science*, pages 252–263, Worcester, Massachusetts, USA, August 17–18, 2000. Springer, Heidelberg, Germany. 1.1

CGMA85. Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In *26th Annual Symposium on Foundations of Computer Science*, pages 383–395, Portland, Oregon, October 21–23, 1985. IEEE Computer Society Press. 1

CJRR99. Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 398–412, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany. 1.1

CK09. Jean-Sébastien Coron and Ilya Kizhvatov. An efficient method for random delay generation in embedded software. In Christophe Clavier and Kris Gaj, editors, *Cryptographic Hardware and Embedded Systems – CHES 2009*, volume 5747 of *Lecture Notes in Computer Science*, pages 156–170, Lausanne, Switzerland, September 6–9, 2009. Springer, Heidelberg, Germany. 1.1

CKOS21. Nishanth Chandran, Bhavana Kanukurthi, Sai Lakshmi Bhavana Obbattu, and Sruthi Sekar. Adaptive extractors and their application to leakage resilient secret sharing. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 595–624, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany. 1

CKOS22. Nishanth Chandran, Bhavana Kanukurthi, Sai Lakshmi Bhavana Obbattu, and Sruthi Sekar. Short leakage resilient and non-malleable secret sharing schemes. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 178–207, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany. 1

DP07. Stefan Dziembowski and Krzysztof Pietrzak. Intrusion-resilient secret sharing. In *48th Annual Symposium on Foundations of Computer Science*, pages 227–237, Providence, RI, USA, October 20–23, 2007. IEEE Computer Society Press. 1

GK18a. Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th Annual ACM Symposium on Theory of Computing*, pages 685–698, Los Angeles, CA, USA, June 25–29, 2018. ACM Press. 1

GK18b. Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing for general access structures. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 501–530, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany. 1

GW16. Venkatesan Guruswami and Mary Wootters. Repairing reed-solomon codes. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 216–226, New York, NY, USA, 2016. Association for Computing Machinery. 1.2

KJJ99. Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany. 1

KK23. Ohad Klein and Ilan Komargodski. New bounds on the local leakage resilience of shamir's secret sharing scheme. Cryptology ePrint Archive, Paper 2023/805, 2023. https://eprint.iacr.org/2023/805. 1.2

KMS19. Ashutosh Kumar, Raghu Meka, and Amit Sahai. Leakage-resilient secret sharing against colluding parties. In David Zuckerman, editor, *60th Annual Symposium on Foundations of Computer Science*, pages 636–660, Baltimore, MD, USA, November 9–12, 2019. IEEE Computer Society Press. 1

Koc96. Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113, Santa Barbara, CA, USA, August 18–22, 1996. Springer, Heidelberg, Germany. 1

MNP^{+}21. Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Leakage-resilience of the shamir secret-sharing scheme against physical-bit leakages. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 344–374, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany. 1.2

MNP^{+}22. Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, Mingyuan Wang, Xiuyu Ye, and Albert Yu. Leakage-resilient linear secret-sharing against arbitrary bounded-size leakage family. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022: 20th Theory of Cryptography Conference, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 355–383, Chicago, IL, USA, November 7–10, 2022. Springer, Heidelberg, Germany. 1.2

MNPC^{+}22. Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, Mingyuan Wang, Xiuyu Ye, and Albert Yu. Tight Estimate of the Local Leakage Resilience of the Additive Secret-Sharing Scheme & Its Consequences. In Dana Dachman-Soled, editor, *3rd Conference on Information-Theoretic Cryptography (ITC 2022)*, volume 230 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:19, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 1.2

MNPCW22. Hemanta K Maji, Hai H Nguyen, Anat Paskin-Cherniavsky, and Mingyuan Wang. Improved bound on the local leakage-resilience of shamir's secret sharing. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pages 2678–2683. IEEE, 2022. 1.2

MOP07. Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. SpringerVerlag New York, Inc., Secaucus, NJ, USA, 01 2007. 1.1

MPSW21. Hemanta K. Maji, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Constructing locally leakage-resilient linear secret-sharing schemes. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 779–808, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany. 1.2

NS20. Jesper Buus Nielsen and Mark Simkin. Lower bounds for leakage-resilient secret sharing. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 556–577, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany. 1, 1.1, 1.1, 2.1, 2.1, 2, 4.2, 4.2

RB89. Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *21st Annual ACM Symposium on Theory of Computing*, pages 73–85, Seattle, WA, USA, May 15–17, 1989. ACM Press. 1

Sha79. Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979. 1, 4.1

SV19. Akshayaram Srinivasan and Prashant Nalini Vasudevan. Leakage resilient secret sharing and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 480–509, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany. 1