

Non-Malleable Secret Sharing against Bounded Joint-Tampering Attacks in the Plain Model

Gianluca Brian¹, Antonio Faonio^{2*}, Maciej Obremski^{3 **}, Mark Simkin^{4***},
and Daniele Venturi¹

¹ Sapienza University of Rome

² IMDEA Software Institute

³ National University of Singapore

⁴ Aarhus University

Abstract. Secret sharing enables a dealer to split a secret into a set of shares, in such a way that certain authorized subsets of share holders can reconstruct the secret, whereas all unauthorized subsets cannot. Non-malleable secret sharing (Goyal and Kumar, STOC 2018) additionally requires that, even if the shares have been tampered with, the reconstructed secret is either the original or a completely unrelated one. In this work, we construct non-malleable secret sharing tolerating p -time *joint-tampering* attacks in the plain model (in the computational setting), where the latter means that, for any $p > 0$ fixed *a priori*, the attacker can tamper with the same target secret sharing up to p times. In particular, assuming one-to-one one-way functions, we obtain:

- A secret sharing scheme for threshold access structures which tolerates joint p -time tampering with subsets of the shares of maximal size (*i.e.*, matching the privacy threshold of the scheme). This holds in a model where the attacker commits to a partition of the shares into non-overlapping subsets, and keeps tampering jointly with the shares within such a partition (so-called *selective partitioning*).
- A secret sharing scheme for general access structures which tolerates joint p -time tampering with subsets of the shares of size $O(\sqrt{\log n})$, where n is the number of parties. This holds in a stronger model where the attacker is allowed to adaptively change the partition within each tampering query, under the restriction that once a subset of the shares has been tampered with jointly, that subset is always either tampered jointly or not modified by other tampering queries (so-called *semi-adaptive partitioning*).

* Supported by the Spanish Government under projects SCUM (ref. RTI2018-102043-B-I00), CRYPTOEPIC (ref. EUR2019-103816), and SECURITAS (ref. RED2018-102321-T), by the Madrid Regional Government under project BLOQUES (ref. S2018/TCS-4339).

** Supported by MOE2019-T2-1-145 Foundations of quantum-safe cryptography.

*** Supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO), grant agreement No 803096 (SPEC), Danish Independent Research Council under Grant-ID DFF-6108-00169 (FoCC), and the Concordium Blockchain Research Center.

At the heart of our result for selective partitioning lies a new technique showing that every one-time *statistically* non-malleable secret sharing against joint tampering is in fact *leakage-resilient* non-malleable (*i.e.*, the attacker can leak jointly from the shares prior to tampering). We believe this may be of independent interest, and in fact we show it implies lower bounds on the share size and randomness complexity of statistically non-malleable secret sharing against *independent* tampering.

Keywords: secret sharing, non-malleability, joint tampering

1 Introduction

In the past 40 years, secret sharing [32, 9] became one of the most fundamental cryptographic primitives. Secret sharing schemes allow a trusted dealer to split a message m into shares s_1, \dots, s_n and distribute them among n participants, such that only certain authorized subsets of share holders are allowed to recover m . The collection \mathcal{A} of authorized subsets is called the *access structure*. The most basic security guarantee is that any *unauthorized* subset outside \mathcal{A} collectively has no information about the shared message. Shamir [32] and Blakley [9] showed how to construct secret sharing schemes with information-theoretic security, and Krawczyk [25] presented the first computationally-secure construction with improved efficiency parameters.

Non-malleable secret sharing. A long line of research [31, 12, 23, 24, 8, 2, 21, 33, 26, 11, 14] has focused on different settings with active adversaries that were allowed to tamper with the shares in one or another way. In verifiable secret sharing [31] the dealer is considered to be untrusted and the share holders want to ensure they hold shares of a consistent secret. In robust secret sharing [12] some parties may act maliciously and try to prevent the correct reconstruction of the shared secret by providing incorrect shares. It is well known that robust secret sharing is impossible when more than half of the parties are malicious.

A recent line of works considers an adversary that has some form of restricted access to *all* shares. In non-malleable secret sharing [23] the adversary can partition the shares in disjoint sets and can then independently tamper with each set of shares. Security guarantees that whatever is reconstructed from the tampered shares is either the original secret, or a completely unrelated value. Most previous works have focused on the setting of independent tampering [23, 24, 8, 2, 21, 33, 26, 11], where the adversary is only allowed to tamper with each share *independently*. Only a few papers [23, 24, 14, 11] have considered the stronger setting where the adversary is allowed to tamper with subsets of shares *jointly*.

Continuous non-malleability. The first notions of non-malleability only focused on security against a *single* round of tampering. A natural extension of this setting is to consider adversaries that may perform several rounds of tampering

attacks on a secret sharing scheme. Badrinarayanan and Srinivasan [8] and Aggarwal *et al.* [2] considered p -time tampering attacks in the information-theoretic setting, where p must be *a-priori* bounded. The works of Faonio and Venturi [21] and Brian, Faonio and Venturi [11] considered *continuous*, i.e., poly-many tampering attacks in the computational setting. It is well known that cryptographic assumptions are inherent in the latter case [22, 8, 21].

An important limitation of all works mentioned above is that, with the exception of [11], they only consider the setting of independent tampering. Brian Faonio, and Venturi [11] achieve continuous non-malleability against joint tampering, where each tampering function can tamper with $O(\log n)$ -large sets of shares assuming a trusted setup in the form of a common reference string. This leads to the following question:

Can we obtain continuously non-malleable secret sharing against joint tampering in the plain model?

1.1 Our Contributions

In this work, we make progress towards answering the above question. Our main contribution is a general framework for reducing *computational* p -time non-malleability against joint tampering to *statistical* one-time non-malleability against joint tampering. Our framework encompasses the following models:

- **Selective partitioning.** Here, the adversary has to initially fix any k -sized partition¹ of the n shares, at the beginning of the experiment. Afterwards, the adversary can tamper p times with the shares within each subset in a joint manner. We call this notion k -joint p -time non-malleability under *selective partitioning*.
- **Semi-adaptive partitioning.** In this setting, the adversary can adaptively choose different k -sized partitions for each tampering query. However, once a subset of the shares has been tampered with jointly, that subset is always either tampered jointly or not modified by other tampering queries. We call this notion k -joint p -time non-malleability under *semi-adaptive partitioning*.

Combining known constructions of one-time statistically non-malleable secret sharing schemes against joint tampering [23, 24, 14] with a new secret sharing scheme that we present in this work, we obtain the following result:

Theorem 1 (Main Theorem, Informal). *Assuming the existence of one-to-one one-way functions, there exist:*

- (i) *A τ -out-of- n secret sharing scheme satisfying k -joint p -time non-malleability under selective partitioning,² for any $\tau \leq n$, $k \leq \tau - 1$, and $p > 0$.*

¹ This a sequence of non-overlapping subsets $\mathcal{B}_1, \dots, \mathcal{B}_t$ covering $[n]$, such that each \mathcal{B}_i has size at most k .

² Here, we inherit a few restrictions from [23]. Namely, the attacker is allowed to tamper jointly using a partition of a minimal reconstruction set in subsets of different sizes. We can remove these restrictions relying on the scheme from [24], which however only works for the n -out-of- n access structure.

- (ii) An (n, τ) -ramp³ secret sharing scheme with binary shares satisfying k -joint p -time non-malleability under selective partitioning, for $\tau = n - n^\beta$, $k \leq \tau - 1$, $\beta < 1$, and $p \in O(\sqrt{n})$.
- (iii) A secret sharing scheme satisfying k -joint p -time non-malleability under semi-adaptive partitioning, for $k \in O(\sqrt{\log n})$ and $p > 0$, and for any access structure that can be described by a polynomial-size monotone span program for which authorized sets have size greater than k .

1.2 Technical Overview

Our initial observation is that a slight variant of a transformation by Ostrovsky *et al.* [30] allows to turn a *bounded leakage-resilient*, statistically one-time non-malleable secret sharing Σ into a *bounded-time* non-malleable secret sharing Σ^* against joint tampering. Bounded leakage resilience here means that, prior to tampering, the attacker may also repeatedly leak information jointly from the shares of Σ , as long as the overall leakage is bounded.

In the setting of joint tampering under selective partitioning, the leakage resilience property of Σ has to hold w.r.t. the same partition used for tampering. For joint tampering under semi-adaptive partitioning, we need Σ to be leakage-resilient under a semi-adaptive choice of the partitions too. A nice feature of this transformation is that it only requires perfectly binding commitments, which can be built from injective one-way functions. Moreover, it preserves the access structure of the underlying secret sharing scheme Σ .

Given the above result, we can focus on the simpler task of constructing bounded leakage-resilient, statistically one-time non-malleable secret sharing, instead of directly attempting to construct their multi-time counterparts. We show different ways of doing that for both settings of selective and semi-adaptive partitioning.

Selective partitioning. First, we show that every statistically one-time non-malleable secret sharing scheme Σ is also resilient to *bounded leakage* under selective partitioning. Let ℓ be an upper bound on the total bit-length of the leakage over all shares. We use an argument reminiscent to standard complexity leveraging to prove that every one-time non-malleable secret sharing scheme with statistical security $\epsilon \in [0, 1)$ is also ℓ -bounded leakage-resilient one-time non-malleable under selective partitioning with statistical security $\epsilon/2^\ell$. The proof roughly works as follows. Given an unbounded attacker A breaking the leakage-resilient one-time non-malleability of Σ , we construct an unbounded attacker \hat{A} against one-time non-malleability of Σ (without leakage). The challenge is how \hat{A} can answer the leakage queries done by A . Our strategy is to simply guess the overall leakage Λ by sampling it uniformly at random, and use this guess to answer all of A 's leakage queries.

³ This means privacy holds with threshold τ , but all of the n shares are required to reconstruct the message.

The problem with this approach is that, whenever our guess was incorrect, the attacker A may notice that it is being used in a simulation and start behaving arbitrarily. We solve this issue with the help of \hat{A} 's final tampering query. Recall that in the model of selective partitioning, all leakage queries and the tampering query, act on the same arbitrary but fixed subsets $\mathcal{B}_1, \dots, \mathcal{B}_t$ of a k -sized partition of the shares. Hence, when A outputs its tampering query (f_1, \dots, f_t) , the reduction \hat{A} defines a modified tampering query $(\hat{f}_1, \dots, \hat{f}_t)$ that first checks whether the guessed leakage from each subset \mathcal{B}_i was correct; if not, the tampering function sets⁴ the modified shares within \mathcal{B}_i to \perp , else it acts identically to f_i . This strategy ensures that our reduction either performs a correct simulation or destroys the secret. In turn, destroying the secret whenever we guessed incorrectly implies that the success probability of \hat{A} is exactly that of A times the probability of guessing the leakage correctly, which is $2^{-\ell}$.

By plugging the schemes from [23, Thm. 2], [24, Thm. 6], and [14, Thm. 3], together with our refined analysis of the transformation by Ostrovsky *et al.* [30], the above insights directly imply items **i** and **ii** of Thm. 1.

Semi-adaptive partitioning. Unfortunately, the argument for showing that one-time non-malleability implies bounded leakage resilience breaks in the setting of adaptive (or even semi-adaptive) partitioning. Intuitively, the problem is that the adversary can leak jointly from adaptively chosen partitions, and thus it is unclear how the reduction can check whether the simulated leakage was correct using a single tampering query.

Hence, we take a different approach. We directly construct a bounded leakage-resilient, statistically one-time non-malleable secret sharing scheme for general access structures. Our construction Σ combines a 2-out-of-2 non-malleable secret sharing scheme Σ_2 with two auxiliary leakage-resilient secret sharing schemes Σ_0 and Σ_1 realizing different access structures. When taking Σ_0 to be the secret sharing scheme from [26, Thm. 1], our construction achieves k -joint bounded leakage-resilient statistical one-time non-malleability under semi-adaptive partitioning for $k \in O(\sqrt{\log n})$. This implies item **iii** of Thm. 1. We refer the reader directly to §5 for a thorough description of our new secret sharing scheme and its security analysis.

Lower bounds. Our complexity leveraging argument implies that every statistically one-time non-malleable secret sharing scheme against *independent* tampering with the shares is also statistically bounded leakage resilient against independent leakage (and no tampering).

By invoking a recent result of Nielsen and Simkin [29], we immediately obtain lower bounds on the share size and randomness complexity of any statistically one-time non-malleable secret sharing scheme against *independent* tampering.

⁴ We assume that the reconstruction algorithm outputs \perp whenever one of the input shares is set to \perp . As we will see later, this is without loss of generality.

1.3 Related Works

Non-malleable secret sharing is intimately related to non-malleable codes [19]. The difference between the two lies in the privacy property: While any non-malleable code in the split-state model [19, 28, 17, 15, 22, 3, 5, 1, 13, 27, 6, 30, 20, 7, 16] is also a 2-out-of-2 secret sharing [17], for any $n \geq 3$ there are n -split-state non-malleable codes that are not private.

Continuously non-malleable codes in the n -split-state model are currently known for $n = 8$ [4] (with statistical security), and for $n = 2$ [22, 20, 30, 16] (with computational security).

Non-malleable secret sharing schemes have useful cryptographic applications, such as non-malleable message transmission [23] and continuously non-malleable threshold signatures [2, 21].

1.4 Paper Organization

The rest of this paper is organized as follows. In §2, we recall a few standard definitions. In §3, we define our model of k -joint non-malleability under selective and semi-adaptive partitioning.

In §4 and §5, we describe our constructions of bounded leakage-resilient statistically one-time non-malleable secret sharing schemes under selective and semi-adaptive partitioning. The lower bounds for non-malleable secret sharing, and the compiler for achieving p -time non-malleability against joint tampering are presented in §6. Finally, in §7, we conclude the paper with a list of open problems for further research.

2 Preliminaries

2.1 Standard Notation

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} . We denote by $[n]$ the set $\{1, \dots, n\}$. For a set of indices $\mathcal{I} = (i_1, \dots, i_t)$ and a vector $x = (x_1, \dots, x_n)$, we write $x_{\mathcal{I}}$ to denote the vector $(x_{i_1}, \dots, x_{i_t})$. When x is chosen randomly in \mathcal{X} , we write $x \leftarrow_s \mathcal{X}$. When A is a randomized algorithm, we write $y \leftarrow_s A(x)$ 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 $A(x; r)$ 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 $A(x; r)$ terminates in a polynomial number of steps (in the size of the input).

Negligible functions. We denote with $\lambda \in \mathbb{N}$ the security parameter. A function p is *polynomial* (in the security parameter), denoted $p \in \text{poly}(\lambda)$, if $p(\lambda) \in O(\lambda^c)$ for some constant $c > 0$. A function $\nu : \mathbb{N} \rightarrow [0, 1]$ 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 often write $\nu(\lambda) \in$

$\text{negl}(\lambda)$ to denote that $\nu(\lambda)$ is negligible. Unless stated otherwise, throughout the paper, we implicitly assume that the security parameter is given as input (in unary) to all algorithms.

Random variables. For a random variable \mathbf{X} , we write $\mathbb{P}[\mathbf{X} = x]$ for the probability that \mathbf{X} takes on a particular value $x \in \mathcal{X}$, with \mathcal{X} being the set where \mathbf{X} is defined. The statistical distance between two random variables \mathbf{X} and \mathbf{Y} over the same set \mathcal{X} is defined as

$$\Delta(\mathbf{X}, \mathbf{Y}) := \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathbb{P}[\mathbf{X} = x] - \mathbb{P}[\mathbf{Y} = x]|.$$

Given two ensembles $\mathbf{X} = \{\mathbf{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathbf{Y} = \{\mathbf{Y}_\lambda\}_{\lambda \in \mathbb{N}}$, we write $\mathbf{X} \equiv \mathbf{Y}$ to denote that they are identically distributed, $\mathbf{X} \stackrel{s}{\approx} \mathbf{Y}$ to denote that they are *statistically close*, i.e. $\Delta(\mathbf{X}_\lambda, \mathbf{Y}_\lambda) \in \text{negl}(\lambda)$, and $\mathbf{X} \stackrel{c}{\approx} \mathbf{Y}$ to denote that they are *computationally indistinguishable*, i.e. for all PPT distinguishers D :

$$|\mathbb{P}[D(\mathbf{X}_\lambda) = 1] - \mathbb{P}[D(\mathbf{Y}_\lambda) = 1]| \in \text{negl}(\lambda).$$

Sometimes we explicitly denote by $\mathbf{X} \stackrel{s}{\approx}_\epsilon \mathbf{Y}$ the fact that $\Delta(\mathbf{X}_\lambda, \mathbf{Y}_\lambda) \leq \epsilon$ for a parameter $\epsilon = \epsilon(\lambda)$. We also extend the notion of computational indistinguishability to the case of interactive experiments (a.k.a. games) featuring an adversary A . In particular, let $\mathbf{G}_A(\lambda)$ be the random variable corresponding to the output of A at the end of the experiment, where wlog. we may assume A outputs a decision bit. Given two experiments $\mathbf{G}_A(\lambda, 0)$ and $\mathbf{G}_A(\lambda, 1)$, we write $\{\mathbf{G}_A(\lambda, 0)\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx} \{\mathbf{G}_A(\lambda, 1)\}_{\lambda \in \mathbb{N}}$ as a shorthand for

$$|\mathbb{P}[\mathbf{G}_A(\lambda, 0) = 1] - \mathbb{P}[\mathbf{G}_A(\lambda, 1) = 1]| \in \text{negl}(\lambda).$$

The above naturally generalizes to statistical distance, which we denote by $\Delta(\mathbf{G}_A(\lambda, 0), \mathbf{G}_A(\lambda, 1))$, in case of *unbounded* adversaries.

We recall a lemma from Dziembowski and Pietrzak [18]:

Lemma 1. *Let \mathbf{X} and \mathbf{Y} be two independent random variables, and $\mathcal{O}_{\text{leak}}(\cdot, \cdot)$ be an oracle that upon input arbitrary functions (g_0, g_1) returns $(g_0(\mathbf{X}), g_1(\mathbf{Y}))$. Then, for any adversary A outputting $\mathbf{Z} \leftarrow A^{\mathcal{O}_{\text{leak}}(\cdot, \cdot)}$, it holds that the random variables $\mathbf{X}|\mathbf{Z}$ and $\mathbf{Y}|\mathbf{Z}$ are independent.*

2.2 Secret Sharing Schemes

An n -party secret sharing scheme Σ consists of polynomial-time algorithms (Share, Rec) specified as follows. The randomized sharing algorithm Share takes a message $m \in \mathcal{M}$ as input and outputs n shares s_1, \dots, s_n , where each $s_i \in \mathcal{S}_i$. The deterministic algorithm Rec takes some number of shares as input and outputs a value in $\mathcal{M} \cup \{\perp\}$. We define $\mu := \log |\mathcal{M}|$ and $\sigma_i := \log |\mathcal{S}_i|$ respectively, to be the bit length of the message and of the i th share.

Which subsets of shares are authorized to reconstruct the secret and which are not is defined via an *access structure*, which is the set of all authorized subsets.

Definition 1 (Access structure). We say that \mathcal{A} is an access structure for n parties if \mathcal{A} is a monotone class of subsets of $[n]$, i.e., if $\mathcal{I}_1 \in \mathcal{A}$ and $\mathcal{I}_1 \subseteq \mathcal{I}_2$, then $\mathcal{I}_2 \in \mathcal{A}$. We call authorized or qualified any set $\mathcal{I} \in \mathcal{A}$, and unauthorized or unqualified any other set. We say that an authorized set $\mathcal{I} \in \mathcal{A}$ is minimal if any proper subset of \mathcal{I} is unauthorized, i.e., if $\mathcal{U} \subsetneq \mathcal{I}$, then $\mathcal{U} \notin \mathcal{A}$.

Intuitively, a perfectly secure secret sharing scheme must be such that all qualified subsets of players can efficiently reconstruct the secret, whereas all unqualified subsets have no information (possibly in a computational sense) about the secret.

Definition 2 (Secret sharing scheme). Let $n \in \mathbb{N}$ and \mathcal{A} be an access structure for n parties. We say that $\Sigma = (\text{Share}, \text{Rec})$ is a secret sharing scheme realizing access structure \mathcal{A} with message space \mathcal{M} and share space $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ if it is an n -party secret sharing with the following properties.

- (i) **Correctness:** For all $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}$ and all authorized subsets $\mathcal{I} \in \mathcal{A}$, we have that $\text{Rec}((\text{Share}(m))_{\mathcal{I}}) = m$ with overwhelming probability over the randomness of the sharing algorithm.
- (ii) **Privacy:** For all PPT adversaries \mathbf{A} , all pairs of messages $m_0, m_1 \in \mathcal{M}$ and all unauthorized subsets $\mathcal{U} \notin \mathcal{A}$, we have that

$$\{(\text{Share}(1^\lambda, m_0))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx} \{(\text{Share}(1^\lambda, m_1))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}}.$$

If the above ensembles are statistically close (resp. identically distributed), we speak of statistical (resp. perfect) privacy.

2.3 Non-Interactive Commitments

A non-interactive commitment scheme Commit is a randomized algorithm taking as input a message $m \in \mathcal{M}$ and outputting a value $c = \text{Commit}(m; r)$ called commitment, using random coins $r \in \mathcal{R}$. The pair (m, r) is called the opening.

Intuitively, a secure commitment satisfies two properties called binding and hiding. The first property says that it is hard to open a commitment in two different ways. The second property says that a commitment hides the underlying message. The formal definition follows.

Definition 3 (Binding). We say that a non-interactive commitment scheme Commit is computationally binding if for all PPT adversaries \mathbf{A} , all messages $m \in \mathcal{M}$, and all random coins $r \in \mathcal{R}$, the following probability is negligible:

$$\mathbb{P}[m' \neq m \wedge \text{Commit}(m'; r') = \text{Commit}(m; r) : (m', r') \leftarrow_{\$} \mathbf{A}(m, r)].$$

If the above holds even in the case of unbounded adversaries, we say that Commit is statistically binding. Finally, if the above probability is exactly 0 for all adversaries (i.e., each commitment can be opened to at most a single message), then we say that Commit is perfectly binding.

Definition 4 (Hiding). We say that a non-interactive commitment scheme Commit is computationally hiding if, for all $m_0, m_1 \in \mathcal{M}$, it holds that

$$\{\text{Commit}(1^\lambda; m_0)\}_{\lambda \in \mathbb{N}} \stackrel{\text{c}}{\approx} \{\text{Commit}(1^\lambda; m_1)\}_{\lambda \in \mathbb{N}}.$$

In case the above ensembles are statistically close (resp. identically distributed), we say that Commit is statistically (resp. perfectly) hiding.

3 Our Leakage and Tampering Model

In this section we define various notions of non-malleability against joint tampering and leakage for secret sharing. Very roughly, in our model the attacker is allowed to partition the set of share holders into t (non-overlapping) blocks with size at most k , covering the entire set $[n]$. This is formalized through the notion of a k -sized partition.

Definition 5 (k -sized partition). Let $n, k, t \in \mathbb{N}$. We call $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_t)$ a k -sized partition of $[n]$ when: (i) $\bigcup_{i=1}^t \mathcal{B}_i = [n]$; (ii) $\forall i_1, i_2 \in [t]$ such that $i_1 \neq i_2$, $\mathcal{B}_{i_1} \cap \mathcal{B}_{i_2} = \emptyset$; (iii) $\forall i \in [t]$, $|\mathcal{B}_i| \leq k$.

Let $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_t)$ be a k -sized partition of $[n]$. To define non-malleability, we consider an adversary \mathbf{A} interacting with a target secret sharing $s = (s_1, \dots, s_n)$ via the following queries:

- **Leakage queries.** For each $i \in [t]$, the attacker can leak jointly from the shares $s_{\mathcal{B}_i}$. This can be done repeatedly and in an adaptive⁵ fashion, as long as the total number of bits that the adversary leaks from each share does not exceed $\ell \in \mathbb{N}$.
- **Tampering queries.** For each $i \in [t]$, the attacker can tamper jointly with the shares $s_{\mathcal{B}_i}$. Each such query yields mauled shares $(\tilde{s}_1, \dots, \tilde{s}_n)$, for which the adversary is allowed to see the corresponding reconstructed message w.r.t. a reconstruction set $\mathcal{T} \in \mathcal{A}$ of his choice. This can be done for at most $p \in \mathbb{N}$ times, and in an adaptive fashion.

Depending on the partition \mathcal{B} being fixed, or chosen adaptively with each leakage/tampering query, we obtain two different flavors of non-malleability, as defined in the following subsections.

3.1 Selective Partitioning

Here, we restrict the adversary to jointly leak from and tamper with subsets of shares belonging to a fixed partition of $[n]$.

⁵ This means that the choice of the next leakage query depends on the overall leakage so far.

Definition 6 (Selective bounded-leakage and tampering admissible adversary). Let $n, k, t, \ell, p \in \mathbb{N}$, and fix an arbitrary message space \mathcal{M} , sharing space $\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$, and access structure \mathcal{A} for n parties. We say that a (possibly unbounded) adversary \mathbf{A} is selective k -joint ℓ -bounded leakage p -tampering admissible (selective (k, ℓ, p) -BLTA for short) if, for every fixed k -sized partition $(\mathcal{B}_1, \dots, \mathcal{B}_t)$ of $[n]$, \mathbf{A} satisfies the following conditions:

- \mathbf{A} outputs a sequence of poly-many leakage queries $(g_1^{(q)}, \dots, g_t^{(q)})$, such that for all $q \in \text{poly}(\lambda)$ and all $i \in [t]$,

$$g_i^{(q)} : \prod_{j \in \mathcal{B}_i} \mathcal{S}_j \rightarrow \{0, 1\}^{\ell_i^{(q)}},$$

where $\ell_i^{(q)}$ is the length of the output $\Lambda_i^{(q)}$ of $g_i^{(q)}$. The only restriction is that $|\Lambda| \leq \ell$, where Λ is the string containing the total leakage performed (over all queries).

- \mathbf{A} outputs a sequence of tampering queries $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_t^{(q)}))$, such that, for all $q \in [p]$, and for all $i \in [t]$, it holds that

$$f_i^{(q)} : \prod_{j \in \mathcal{B}_i} \mathcal{S}_j \rightarrow \prod_{j \in \mathcal{B}_i} \mathcal{S}_j \quad \text{and} \quad \mathcal{T}^{(q)} \cap \mathcal{B}_i \neq \emptyset,$$

and moreover $\mathcal{T}^{(q)} \in \mathcal{A}$ is a minimal authorized subset.

- All queries performed by \mathbf{A} are chosen adaptively, i.e. each query may depend on the information obtained from all the previous queries.
- If $p > 0$, the last query performed by \mathbf{A} is a tampering query.

Note that \mathbf{A} can choose a different reconstruction set $\mathcal{T}^{(q)}$ with each tampering query, in a fully adaptive manner. This feature is known as *adaptive reconstruction* [21]. However, we consider the following two restrictions (that were not present in previous works): (i) Each set $\mathcal{T}^{(q)}$ must be minimal and contain at least one mauled share from each subset \mathcal{B}_i ; (ii) The last query asked by \mathbf{A} is a tampering query. Looking ahead, these technical conditions are needed for the complexity leveraging argument used in Thm. 3. Note that the above restrictions are still meaningful, as they allow, *e.g.*, to capture the setting in which the attacker first leaks from all the shares and then tampers with the shares in a minimal authorized subset.

3.2 Semi-Adaptive Partitioning

Next, we generalize the above definition to the stronger setting in which the adversary is allowed to change the k -sized partition with each leakage and tampering query. Here, we do not consider the restriction (i) mentioned above as it is not needed for the analysis of our secret sharing scheme in §5; yet we still consider the restriction (ii), and we will need to restrict the way in which the attacker specifies the partitions corresponding to each leakage and tampering query. For this reason, we refer to our model as *semi-adaptive partitioning*.

Definition 7 (Semi-adaptive bounded-leakage and tampering admissible adversary). Let $n, k, \ell, p \in \mathbb{N}$ and $\mathcal{M}, \mathcal{S}, \mathcal{A}$ as in Def. 6. We say that a (possibly unbounded) adversary \mathbf{A} is semi-adaptive k -joint ℓ -bounded leakage p -tampering admissible (semi-adaptive (k, ℓ, p) -BLTA for short) if it satisfies the following conditions:

- \mathbf{A} outputs a sequence of poly-many leakage queries $(\mathcal{B}^{(q)}, (g_1^{(q)}, \dots, g_{t^{(q)}}^{(q)}))$, chosen adaptively, such that, for all $q \in \text{poly}(\lambda)$, and for all $i \in [t^{(q)}]$, it holds that $\mathcal{B}^{(q)} = (\mathcal{B}_1^{(q)}, \dots, \mathcal{B}_{t^{(q)}}^{(q)})$ is a k -sized partition of $[n]$ and

$$g_i^{(q)} : \prod_{j \in \mathcal{B}_i^{(q)}} \mathcal{S}_j \rightarrow \{0, 1\}^{\ell_i^{(q)}},$$

where $\ell_i^{(q)}$ is the length of the output. The only restriction is that $|\Lambda| \leq \ell$, where $\Lambda = (\Lambda^{(1)}, \Lambda^{(2)}, \dots)$ is the total leakage (over all queries).

- \mathbf{A} outputs a sequence of p tampering queries $(\mathcal{B}^{(q)}, \mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_{t^{(q)}}^{(q)}))$, chosen adaptively, such that, for all $q \in [p]$, and for all $i \in [t^{(q)}]$, it holds that $\mathcal{B}^{(q)}$ is a k -sized partition of $[n]$ and

$$f_i^{(q)} : \prod_{j \in \mathcal{B}_i^{(q)}} \mathcal{S}_j \rightarrow \prod_{j \in \mathcal{B}_i^{(q)}} \mathcal{S}_j.$$

- All queries performed by \mathbf{A} are chosen adaptively, i.e. each query may depend on the information obtained from all the previous queries.
- If $p > 0$, the last query performed by \mathbf{A} is a tampering query.
- Given a tampering query $(\mathcal{B}, \mathcal{T}, f)$, let $\mathcal{T} = \{\beta_1, \dots, \beta_\tau\}$ for $\tau \in \mathbb{N}$. We write $\xi(i)$ for the index such that $\beta_i \in \mathcal{B}_{\xi(i)}$; namely, the i -th share used in the reconstruction is tampered by the $\xi(i)$ -th tampering function. Then:
 - (i) For all leakage queries (\mathcal{B}, g) and all tampering queries $(\mathcal{B}', \mathcal{T}', f')$, where $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_t)$ and $\mathcal{B}' = (\mathcal{B}'_1, \dots, \mathcal{B}'_{t'})$, the following holds: for all indices $i \in [t]$, either there exists $j \in \mathcal{T}'$ such that $\mathcal{B}_i \subseteq \mathcal{B}'_{\xi(j)}$, or for all $j \in \mathcal{T}'$ we have $\mathcal{B}_i \cap \mathcal{B}'_{\xi(j)} = \emptyset$.
 - (ii) For any pair of tampering queries $(\mathcal{B}', \mathcal{T}', f')$ and $(\mathcal{B}'', \mathcal{T}'', f'')$, where $\mathcal{B}' = \{\mathcal{B}'_1, \dots, \mathcal{B}'_{t'}\}$ and $\mathcal{B}'' = \{\mathcal{B}''_1, \dots, \mathcal{B}''_{t''}\}$, the following holds: for all $i \in \mathcal{T}'$, either there exists $j \in \mathcal{T}''$ such that $\mathcal{B}'_{\xi(i)} \subseteq \mathcal{B}''_{\xi(j)}$, or for all $j \in \mathcal{T}''$ we have $\mathcal{B}'_{\xi(i)} \cap \mathcal{B}''_{\xi(j)} = \emptyset$.

Intuitively, condition (i) means that whenever the attacker leaks jointly from the shares within a subset \mathcal{B}_i , then for any tampering query the adversary must either tamper jointly with the shares within \mathcal{B}_i , or do not modify those shares at all. Condition (ii) is the same translated to the partitions corresponding to different tampering queries. Looking ahead, condition (i) is needed for the proof in §5.3, whereas condition (ii) is needed for the proof in §6.2. Note that the above restrictions are still meaningful, as they allow, e.g., to capture the setting

<p>JSTamper$_{\Sigma, \mathcal{A}}^{\mathcal{B}, m_0, m_1}(\lambda, b)$:</p> <p>$s := (s_1, \dots, s_n) \leftarrow \text{Share}(m_b)$</p> <p>stop \leftarrow false</p> <p>Return $\mathbf{A}^{\mathcal{O}_{\text{nmss}}(s, \mathcal{B}, \cdot, \cdot), \mathcal{O}_{\text{leak}}(s, \mathcal{B}, \cdot)}(1^\lambda)$</p>	<p>Oracle $\mathcal{O}_{\text{nmss}}(s, \mathcal{B}, \mathcal{T}, (f_1, \dots, f_t))$:</p> <p>If stop = true</p> <p>Return \perp</p> <p>Else</p> <p>$\forall i \in [t] : \tilde{s}_{\mathcal{B}_i} := f_i(s_{\mathcal{B}_i})$</p> <p>$\tilde{s} = (\tilde{s}_1, \dots, \tilde{s}_n)$</p> <p>$\tilde{m} = \text{Rec}(\tilde{s}_{\mathcal{T}})$</p> <p>If $\tilde{m} \in \{m_0, m_1\}$</p> <p>Return \diamond</p> <p>If $\tilde{m} = \perp$</p> <p>Return \perp</p> <p>stop \leftarrow true</p> <p>Else return \tilde{m}</p>
<p>JATamper$_{\Sigma, \mathcal{A}}^{m_0, m_1}(\lambda, b)$:</p> <p>$s := (s_1, \dots, s_n) \leftarrow \text{Share}(m_b)$</p> <p>stop \leftarrow false</p> <p>Return $\mathbf{A}^{\mathcal{O}_{\text{nmss}}(s, \cdot, \cdot, \cdot), \mathcal{O}_{\text{leak}}(s, \cdot, \cdot)}(1^\lambda)$</p>	<p>Oracle $\mathcal{O}_{\text{leak}}(s, \mathcal{B}, (g_1, \dots, g_t))$:</p> <p>Return $g_1(s_{\mathcal{B}_1}), \dots, g_t(s_{\mathcal{B}_t})$</p>

Fig. 1. Experiments defining selective (**JSTamper**) and adaptive (**JATamper**) joint leakage-resilient (continuously) non-malleable secret sharing. The oracle $\mathcal{O}_{\text{nmss}}$ is implicitly parameterized by the flag **stop**.

in which the attacker defines two non-overlapping⁶ subsets of $[n]$ and then performs joint leakage under adaptive partitioning within the first subset and joint leakage/tampering under selective partitioning within the second subset.

3.3 The Definition

Very roughly, leakage-resilient non-malleability states that no admissible adversary, as defined above, can distinguish whether it is interacting with a secret sharing of m_0 or of m_1 .

Definition 8 (Leakage-resilient non-malleability). *Let $n, k, \ell, p \in \mathbb{N}$ and $\epsilon \in [0, 1]$ be parameters, and \mathcal{A} be an access structure for n parties. We say that $\Sigma = (\text{Share}, \text{Rec})$ is a k -joint ℓ -bounded leakage-resilient p -time ϵ -non-malleable secret sharing scheme realizing \mathcal{A} , shortened (k, ℓ, p, ϵ) -BLR-NMSS, if it is an n -party secret sharing scheme realizing \mathcal{A} , and additionally, for all pairs of messages $m_0, m_1 \in \mathcal{M}$, we have one of the following:*

- For all selective (k, ℓ, p) -BLTA adversaries \mathbf{A} , and for all k -sized partitions \mathcal{B} of $[n]$,

$$\left\{ \mathbf{JSTamper}_{\Sigma, \mathcal{A}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) \right\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_{\epsilon} \left\{ \mathbf{JSTamper}_{\Sigma, \mathcal{A}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) \right\}_{\lambda \in \mathbb{N}}. \quad (1)$$

In this case, we speak of (k, ℓ, p, ϵ) -BLR-NMSS under selective partitioning.

- For all semi-adaptive (k, ℓ, p) -BLTA adversaries \mathbf{A} ,

$$\left\{ \mathbf{JATamper}_{\Sigma, \mathcal{A}}^{m_0, m_1}(\lambda, 0) \right\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_{\epsilon} \left\{ \mathbf{JATamper}_{\Sigma, \mathcal{A}}^{m_0, m_1}(\lambda, 1) \right\}_{\lambda \in \mathbb{N}}. \quad (2)$$

⁶ In fact, the two subsets do not need to be fixed a priori.

In this case, we speak of (k, ℓ, p, ϵ) -BLR-NMSS under semi-adaptive partitioning.

Experiments $\mathbf{JSTamper}_{\Sigma, \mathbf{A}}^{\mathcal{B}, m_0, m_1}(\lambda, b)$ and $\mathbf{JATamper}_{\Sigma, \mathbf{A}}^{m_0, m_1}(\lambda, b)$, for $b \in \{0, 1\}$, are depicted in Fig. 1.

In case there exists $\epsilon = \epsilon(\lambda) \in \text{negl}(\lambda)$ such that indistinguishability still holds computationally in the above definitions for any $p = p(\lambda) \in \text{poly}(\lambda)$, and any PPT adversaries \mathbf{A} , we call Σ bounded leakage-resilient continuously non-malleable, shortened (k, ℓ) -BLR-CNMSS, under selective/semi-adaptive partitioning.

Non-malleable secret sharing. When no leakage is allowed (*i.e.*, $\ell = 0$), we obtain the notion of non-malleable secret sharing as a special case. In particular, an adversary is k -joint p -time tampering admissible, shortened (k, p) -TA, if it is $(k, 0, p)$ -BLTA. Furthermore, we say that Σ is a k -joint p -time ϵ -non-malleable secret sharing, shortened (k, p, ϵ) -NMSS, if Σ is a $(k, 0, p, \epsilon)$ -BLR-NMSS scheme.

Leakage-resilient secret sharing. When no tampering is allowed (*i.e.*, $p = 0$), we obtain the notion of leakage-resilient secret sharing as a special case. In particular, an adversary is k -joint ℓ -bounded leakage admissible, shortened (k, ℓ) -BLA, if it is $(k, \ell, 0)$ -BLTA. Furthermore, we say that Σ is a k -joint ℓ -bounded ϵ -leakage-resilient secret sharing, shortened (k, ℓ, ϵ) -BLRSS, if Σ is a $(k, \ell, 0, \epsilon)$ -BLR-NMSS scheme.

Finally, we denote by $\mathbf{JSLeak}_{\Sigma, \mathbf{A}}^{\mathcal{B}, m_0, m_1}(\lambda, b)$ and $\mathbf{JALeak}_{\Sigma, \mathbf{A}}^{m_0, m_1}(\lambda, b)$ the experiments in Def. 8 defining leakage resilience against selective and semi-adaptive partitioning respectively. However, note that when no tampering happens the conditions (i) and (ii) of Def. 7 are irrelevant, and thus we simply speak of (k, ℓ, ϵ) -BLRSS under adaptive partitioning.

Augmented leakage resilience. We also define a seemingly stronger variant of leakage-resilient secret sharing, in which \mathbf{A} is allowed to obtain the shares within a subset of the partition \mathcal{B} (in the case of selective partitioning, or any unauthorized subset of at most k shares in the case of adaptive partitioning) at the end of the experiment. In particular, in the case of selective partitioning, an *augmented* admissible adversary is an attacker $\mathbf{A}^+ = (\mathbf{A}_1^+, \mathbf{A}_2^+)$ such that:

- \mathbf{A}_1^+ is an admissible adversary in the sense of Def. 6, the only difference being that \mathbf{A}_1^+ outputs a tuple (α, i^*) , where α is an auxiliary state, and $i^* \in [t]$;
- \mathbf{A}_2^+ takes as input α and all the shares $s_{\mathcal{B}_{i^*}}$, and outputs a decision bit.

In case of adaptive partitioning, the definition changes as follows: the adversary \mathbf{A}_1^+ is admissible in the sense of Def. 7 and outputs an unauthorized subset $\mathcal{U} \notin \mathcal{A}$ of size at most k instead of the index i^* , and \mathbf{A}_2^+ takes as input the shares $s_{\mathcal{U}}$ instead of the shares $s_{\mathcal{B}_{i^*}}$.

This flavor of security is called *augmented leakage resilience*. The theorem below, which was established by [11, 26] for the case of independent leakage, shows that any joint LRSS is also an augmented LRSS at the cost of an extra bit of leakage.

Theorem 2. *Let Σ be a $(k, \ell + 1, \epsilon)$ -BLRSS realizing access structure \mathcal{A} under selective/adaptive partitioning. Then, Σ is an augmented (k, ℓ, ϵ) -BLRSS realizing \mathcal{A} under selective/adaptive partitioning.*

Proof. By reduction to non-augmented leakage resilience. Let $A^+ = (A_1^+, A_2^+)$ be a (k, ℓ, ϵ) -BLA adversary violating augmented leakage-resilience; we construct an adversary A breaking the non-augmented variant of leakage resilience. Fix $m_0, m_1 \in \mathcal{M}$ and a k -sized partition $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_t)$. Attacker A works as follows.

- Run A_1^+ and, upon input a leakage query (g_1, \dots, g_t) , forward the same query to the target leakage oracle and return the answer to A_1^+ .
- Let (α, i^*) be the final output of A_1^+ . Define the leakage function $\hat{g}_{i^*}^{\alpha, A_2^+}$ which hard-wires α and a description of A_2^+ , takes as input the shares $s_{\mathcal{B}_{i^*}}$ and returns the decision bit $b' \leftarrow A_2^+(\alpha, s_{\mathcal{B}_{i^*}})$.
- Forward $(\varepsilon, \dots, \varepsilon, \hat{g}_{i^*}^{\alpha, A_2^+}, \varepsilon, \dots, \varepsilon)$ to the target leakage oracle, obtaining a bit b' .
- Output b' .

The statement follows by observing that A 's simulation to A^+ 's leakage queries is perfect, thus A and A^+ have the same advantage, and moreover A leaks a total of at most $\ell + 1$ bits. \square

4 Selective Partitioning

In this section, we construct bounded leakage-resilient, statistically one-time non-malleable secret sharing under selective partitioning. We achieve this in two steps. First, in §4.1, we prove that every statistically one-time non-malleable secret sharing is in fact bounded leakage-resilient, statistically one-time non-malleable under selective partitioning at the price of a security loss exponential in the size of the leakage. Then, in §4.2, we provide concrete instantiations using known results from the literature.

4.1 Non-Malleability Implies Bounded Leakage Resilience

Theorem 3. *Let $\Sigma = (\text{Share}, \text{Rec})$ be a $(k, 1, \epsilon/2^\ell)$ -NMSS realizing \mathcal{A} . Then, Σ is also a $(k, \ell, 1, \epsilon)$ -BLR-NMSS realizing \mathcal{A} under selective partitioning.*

Proof. By contradiction, assume that there exist a pair of messages $m_0, m_1 \in \mathcal{M}$, a k -partition $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_t)$ of $[n]$, and a $(k, \ell, 1)$ -BLTA unbounded adversary A such that

$$\left| \mathbb{P} \left[\text{JSTamper}_{\Sigma, \mathcal{A}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) = 1 \right] - \mathbb{P} \left[\text{JSTamper}_{\Sigma, \mathcal{A}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) = 1 \right] \right| > \epsilon.$$

Consider the following unbounded reduction \hat{A} trying to break $(k, 0, 1, \epsilon/2^\ell)$ -non-malleability using the same partition \mathcal{B} , and the same messages m_0, m_1 .

1. Run $\mathbf{A}(1^\lambda)$.
2. Upon input the q -th leakage query $g^{(q)} = (g_1^{(q)}, \dots, g_t^{(q)})$, generate a uniformly random string $\Lambda^{(q)} = (\Lambda_1^{(q)}, \dots, \Lambda_t^{(q)})$ compatible with the range of $g^{(q)}$, and output $\Lambda^{(q)}$ to \mathbf{A} .
3. Upon input the final tampering query $f = (f_1, \dots, f_t)$, construct the following tampering function $\hat{f} = (\hat{f}_1, \dots, \hat{f}_t)$:
 - The function hard-wires (a description of) all the leakage functions $g^{(q)}$, the tampering query f , and the guess on the leakage $\Lambda = \Lambda^{(1)} \parallel \Lambda^{(2)} \parallel \dots$.
 - Upon input the shares $(s_j)_{j \in \mathcal{B}_i}$, the function \hat{f}_i checks that the guess on the leakage was correct, *i.e.* $g_i^{(q)}((s_j)_{j \in \mathcal{B}_i}) = \Lambda_i^{(q)}$ for all q . If the guess was correct, compute and output $f_i((s_j)_{j \in \mathcal{B}_i})$; else, output \perp .
4. Send \hat{f} to the tampering oracle and pass the answer $\tilde{m} \in \mathcal{M} \cup \{\diamond, \perp\}$ to \mathbf{A} .
5. Output the same guessing bit as \mathbf{A} .

For the analysis, we now compute the distinguishing advantage of $\hat{\mathbf{A}}$. In particular, call \mathbf{Miss}_b the event in which the guess on the leakage was wrong in experiment $\mathbf{JTamper}_{\Sigma, \mathbf{A}}^{\mathcal{B}, m_0, m_1}(\lambda, b)$, *i.e.* there exists $i \in [t]$ such that \hat{f}_i outputs \perp in step 3, and call \mathbf{Hit}_b its complementary event. We notice that the probability of \mathbf{Hit}_0 is equal to the probability of \mathbf{Hit}_1 , since the strings $\Lambda^{(q)}$ are sampled uniformly at random:

$$\mathbb{P}[\mathbf{Hit}_b] = \sum_{\Lambda \in \{0,1\}^\ell} \mathbb{P}[\mathbf{U}_\ell = \Lambda \wedge g(\mathbf{S}^b) = \Lambda] = 2^{-\ell} \sum_{\Lambda \in \{0,1\}^\ell} \mathbb{P}[g(\mathbf{S}^b) = \Lambda] = 2^{-\ell},$$

where \mathbf{S}^b is the random variable corresponding to $\text{Share}(m_b)$, \mathbf{U}_ℓ is the uniform distribution over $\{0,1\}^\ell$, and g is the concatenation of all the leakage functions. Then, we can write:

$$\begin{aligned} & \left| \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) = 1 \right] - \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) = 1 \right] \right| \\ &= \left| \mathbb{P}[\mathbf{Hit}_0] \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) = 1 \mid \mathbf{Hit}_0 \right] \right. \end{aligned} \quad (3)$$

$$\begin{aligned} & \quad - \mathbb{P}[\mathbf{Hit}_1] \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) = 1 \mid \mathbf{Hit}_1 \right] \\ & \quad + \mathbb{P}[\mathbf{Miss}_0] \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) = 1 \mid \mathbf{Miss}_0 \right] \\ & \quad - \mathbb{P}[\mathbf{Miss}_1] \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) = 1 \mid \mathbf{Miss}_1 \right] \\ &= 2^{-\ell} \left| \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) = 1 \mid \mathbf{Hit}_0 \right] \right. \end{aligned} \quad (4)$$

$$\begin{aligned} & \quad - \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \hat{\mathbf{A}}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) = 1 \mid \mathbf{Hit}_1 \right] \\ &= 2^{-\ell} \left| \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \mathbf{A}}^{\mathcal{B}, m_0, m_1}(\lambda, 0) = 1 \right] \right. \end{aligned} \quad (5)$$

$$\begin{aligned} & \quad - \mathbb{P} \left[\mathbf{JTamper}_{\Sigma, \mathbf{A}}^{\mathcal{B}, m_0, m_1}(\lambda, 1) = 1 \right] \Big| > \frac{\epsilon}{2^\ell}, \end{aligned} \quad (6)$$

In the above derivation, Eq. (3) follows from the law of total probability, Eq. (4) comes from the fact that, when \mathbf{Miss} happens, the view of \mathbf{A} (*i.e.* the leakage

A and the output of the tampering query) is independent⁷ of the target secret sharing, and thus its distinguishing advantage is zero, and Eq. (5) follows because $\mathbb{P}[\mathbf{Hit}] = 2^{-\ell}$ and moreover, when \mathbf{Hit} happens, the view of A is perfectly simulated and thus \hat{A} has the same distinguishing advantage of A , which is at least ϵ by assumption.

Therefore, \hat{A} has a distinguishing advantage of at least $\epsilon/2^\ell$. Finally, note that \hat{A} performs no leakage and uses only one tampering query, and thus \hat{A} is $(k, 1)$ -TA. The lemma follows. \square

4.2 Instantiations

Using known constructions of one-time non-malleable secret sharing schemes against joint tampering, we obtain the following:

Corollary 1. *For every $\lambda, \ell, n \geq 0$, and every $k, \tau \geq 0$ such that $k < \tau \leq n$, there exists a τ -out-of- n secret sharing Σ that is a $(k, \ell, 1, 2^{-\lambda})$ -BLR-NMSS under selective partitioning.*

Proof. Follows by combining Thm. 3 with the secret sharing scheme⁸ of [23, Thm. 4], using security parameter $\lambda' + \ell$ and choosing $\lambda \geq (\lambda' + \ell)^{\Omega(1)} - \ell$ in order to obtain

$$\epsilon = 2^\ell \cdot 2^{-(\lambda' + \ell)^{\Omega(1)}} \leq 2^{-\lambda}.$$

\square

Corollary 2. *For every $\ell, n \geq 0$, any $\beta < 1$, and every $k, \tau \geq 0$ such that $k < \tau \leq n$, there exists an (n, τ) -ramp secret sharing Σ that is a $(k, \ell, 1, 2^\ell \cdot 2^{-n^{\Omega(1)}})$ -BLR-NMSS under selective partitioning with binary shares.*

Proof. Follows by combining Thm. 3 with the secret sharing scheme of [14, Thm. 4.1].

5 Semi-Adaptive Partitioning

As mentioned in the introduction, the proof of Thm. 3 breaks in the setting of semi-adaptive partitioning. To overcome this issue, in §5.1, we give a direct construction of a bounded leakage-resilient, one-time statistically non-malleable secret sharing (for general access structures) under semi-adaptive partitioning. We explain the main intuition behind our design in §5.2, and formally prove security in §5.3. Finally, in §5.4, we explain how to instantiate our construction using known results from the literature.

⁷ Here is where we use the restriction that the reconstruction set \mathcal{T} must be minimal and contain at least one share from each subset \mathcal{B}_i ; otherwise, we cannot argue that the output of the tampering query is \perp , and thus independent of the target.

⁸ The construction in [23, Thm. 4] actually only achieves security against joint tampering within a partition \mathcal{B} of the reconstruction set \mathcal{T} (rather than the entire set $[n]$). Accordingly, in this case we can only tolerate joint leakage from the shares within the same partition \mathcal{B} .

5.1 Our New Secret Sharing Scheme

Let Σ_0 be a secret sharing realizing access structure \mathcal{A} , let Σ_1 be a k_1 -out-of- n secret sharing, and let Σ_2 be a 2-out-of-2 secret sharing. Consider the following scheme $\Sigma = (\text{Share}, \text{Rec})$:

- **Algorithm Share:** Upon input m , first compute $(s_0, s_1) \leftarrow \text{Share}_2(m)$, $(s_{0,1}, \dots, s_{0,n}) \leftarrow \text{Share}_0(s_0)$, and $(s_{1,1}, \dots, s_{1,n}) \leftarrow \text{Share}_1(s_1)$. Then set $s_i := (s_{0,i}, s_{1,i})$ for all $i \in [n]$, and output (s_1, \dots, s_n) .
- **Algorithm Rec:** Upon input $(s_i)_{i \in \mathcal{I}}$, parse $s_i = (s_{0,i}, s_{1,i})$ and $\mathcal{I} = \{i_1, \dots, i_{|\mathcal{I}|}\}$, and define $\mathcal{I}_{k_1} := \{i_1, \dots, i_{k_1}\}$; compute $s_1 = \text{Rec}_1((s_{1,i})_{i \in \mathcal{I}_{k_1}})$ and $s_0 = \text{Rec}_0((s_{0,i})_{i \in \mathcal{I}})$, and finally output $m' = \text{Rec}_2((s_0, s_1))$.

With the above defined scheme, we achieve the following:

Theorem 4. *Let $n, k(\lambda), \ell(\lambda), \sigma_0(\lambda) \in \mathbb{N}$ and $\epsilon_0, \epsilon_1, \epsilon_2 \in [0, 1]$ be parameters, and set $k_1 := \sqrt{k}$, $\ell_0 := \ell + 1$ and $\ell_1 := \ell + n \cdot \sigma_0$. Let \mathcal{A} be an arbitrary access structure for n parties, where for any $\mathcal{I} \in \mathcal{A}$ we have $|\mathcal{I}| > k_1$. Assume that:*

1. Σ_0 is a (k, ℓ_0, ϵ_0) -BLRSS realizing \mathcal{A} under adaptive partitioning, with share space such that $\log |\mathcal{S}_{0,i}| \leq \sigma_0$ (for any $i \in [n]$);
2. Σ_1 is a $(k_1 - 1, \ell_1, \epsilon_1)$ -BLRSS realizing the k_1 -out-of- n threshold access structure under adaptive partitioning;
3. Σ_2 is a one-time ϵ_2 -non-malleable 2-out-of-2 secret sharing (i.e. a $(1, 1, \epsilon_2)$ -NMSS).

Then, the above defined Σ is a $(k_1 - 1, \ell, 1, 2(\epsilon_0 + \epsilon_1) + \epsilon_2)$ -BLR-NMSS realizing \mathcal{A} under semi-adaptive partitioning.

5.2 Proof Overview

In order to prove Thm. 4, we first make some considerations on the tampering query $(\mathcal{T}, \mathcal{B}, f)$. In particular, we construct two disjoint sets \mathcal{T}_0^* and \mathcal{T}_1^* that are the union of subsets from the partition \mathcal{B} , in such a way that (i) $\mathcal{T}_0^* \cap \mathcal{T}$ contains at least k_1 elements (so that it can be used as a reconstruction set for Rec_1); and (ii) each subset \mathcal{B}_i of the partition \mathcal{B} intersects at most one of $\mathcal{T}_0^*, \mathcal{T}_1^*$ (so that both leakage and tampering queries can be computed on \mathcal{T}_0^* and on \mathcal{T}_1^* independently). Hence, we define four hybrid experiments as described below.

First Hybrid: In the first hybrid experiment, we change how the tampering query is answered. Namely, after the last leakage query, we replace all the left shares $(s_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ with new shares $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ that are valid shares of s_0 and consistent with the leakage obtained by the adversary and with the shares $(s_{0,\beta})_{\beta \in \mathcal{T}_0^*}$. Here, we note that due to the fact that we only consider *semi-adaptive* partitioning,⁹ the shares $(s_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ and $(s_{1,\beta})_{\beta \in \mathcal{T}_0^*}$ are independent

⁹ We thank Ashutosh Kumar for pointing out to us that independence given the leakage does not necessarily hold in the case of fully adaptive (rather than semi-adaptive) partitioning.

even given the leakage. In particular, the above shares are independent before the leakage occurs, and furthermore condition (i) in Def. 7 ensures that the adversary never leaks jointly from shares in \mathcal{T}_0^* and in \mathcal{T}_1^* . Thus, since the old and the new shares are sampled from the same distribution, this change does not affect the view of the adversary and does not modify its advantage.

Second Hybrid: In the second hybrid experiment, we change the distribution of the left shares. Namely, we discard the original ones and we replace them with left shares of some unrelated message \hat{s}_0 , where $(\hat{s}_0, \hat{s}_1) \leftarrow^* \text{Share}_2(0)$. In order to prove that this hybrid experiment is ϵ_0 -close to the previous one, we construct an admissible reduction to leakage resilience of Σ_0 , thus proving that, if some admissible adversary is able to notice the difference between the old and the new experiment with advantage more than ϵ_0 , then our reduction can distinguish between a secret sharing of s_0 and a secret sharing of \hat{s}_0 with exactly the same advantage.

The key idea here is to forward leakage queries to the target oracle and, once the adversary outputs its tampering query, obtain all the shares in \mathcal{T}_0^* from the challenger, using the augmented property ensured by Thm. 2; the reduction remains admissible because Σ_0 has security against adaptive k -partitioning and $|\mathcal{T}_0^*| \leq k$. After receiving such shares, the reduction can sample the shares $(s_{0,\beta}^*)_{\mathcal{T}_1^*}$ as in the first hybrid experiment and compute the tampering on both s_0 (using the shares in \mathcal{T}_0^* and the sampled shares in \mathcal{T}_1^*) and s_1 (only using the shares in \mathcal{T}_0^*), which allows to simulate the tampering query.

Third Hybrid: In the third hybrid experiment, we change how the tampering query is answered. Similarly to the modification introduced in the first hybrid experiment, after the last leakage query, we replace all the right shares $(s_{1,\beta})_{\beta \in \mathcal{T}_0^*}$ with new shares $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ that are valid shares of s_1 and consistent with the leakage obtained by the adversary. However, we now further require that this change does not affect the outcome of the tampering query on the left shares; in particular, if the tampering function applied to $(\hat{s}_{0,\beta}, s_{1,\beta})$ leads to $(\tilde{s}_{0,\beta}, *)$, the same tampering function applied to $(\hat{s}_{0,\beta}, s_{1,\beta}^*)$ must lead to $(\tilde{s}_{0,\beta}, *)$. This is required in order to keep consistency with the modifications introduced in the second hybrid experiment. As before, since the old and the new shares are sampled from the same distribution, this change does not modify the advantage of the adversary.

Fourth Hybrid: In the fourth hybrid experiment, we change the distribution of the right shares. Similarly to the modification introduced in the third hybrid experiment, we discard the original shares and replace them with the right shares of the previously computed unrelated message, *i.e.* \hat{s}_1 . In order to prove that this hybrid experiment is ϵ_1 -close to the previous one, we construct an admissible reduction to leakage resilience of Σ_1 .

The key idea here is to simulate the tampering query with a leakage query that yields the result of the tampering on all the left shares $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$, where $\mathcal{T}^* = \mathcal{T}_0^* \cup \mathcal{T}_1^*$. This is allowed because of the restriction on the shares of Σ_0 being at most σ_0 bits long, so that the total performed leakage is bounded by $\ell + n\sigma_0$. In particular, after sampling the fake shares

$(\hat{s}_{0,1}, \dots, \hat{s}_{0,n})$, forwarding the leakage queries to the target oracle and receiving the tampering query, the reduction samples the shares $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ as in the second hybrid experiment and hard-wires them, along with the shares $(\hat{s}_{0,1}, \dots, \hat{s}_{0,n})$, inside a leakage function that computes $(\tilde{s}_{0,\beta}, \tilde{s}_{1,\beta})_{\beta \in \mathcal{T}^*}$ and outputs $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$. After receiving the mauled shares, the reduction samples the shares $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ as in the third hybrid and computes the corresponding tampered shares $(\tilde{s}_{1,\beta})_{\beta \in \mathcal{T}_0^*}$. Given the mauled shares $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$ and $(\tilde{s}_{1,\beta})_{\beta \in \mathcal{T}_0^*}$, the reduction can then simulate the tampering query correctly.

Since the above defined hybrid experiments are all statistically close, it only remains to show that no adversary can distinguish between the last hybrid experiment with bit $b = 0$ and the same experiment with $b = 1$ with an advantage more than ϵ_2 , thus proving the security of our scheme. Here, we once again construct a reduction, this time to one-time ϵ_2 -non-malleability, that achieves the same advantage of an adversary distinguishing between the two experiments.

The key idea is to use s_0 to sample the shares $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ and s_1 to sample the shares $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$. In particular, all the missing shares needed for the computation are the one sampled from (\hat{s}_0, \hat{s}_1) and, since $\mathcal{T}_0^* \cap \mathcal{T}_1^* = \emptyset$, there is no overlap and the tampering can be split between two functions f_0, f_1 that hard-wire the sampled values. These two functions take as input s_0 and s_1 , respectively, and can thus compute the mauled values \tilde{s}_0 and \tilde{s}_1 , which in turn allows the reduction to simulate the tampering query.

5.3 Security Analysis

Before proceeding with the analysis, we introduce some useful notation. We will define a sequence of hybrid experiments $\mathbf{H}_i(\lambda, b)$ for $i \in \mathbb{N}$ and $b \in \{0, 1\}$, starting with $\mathbf{H}_0(\lambda, b)$ which is identical to the $\mathbf{JATamper}_{\Sigma, \mathcal{A}}(\lambda, b)$ experiment. Recall that, after the leakage phase, the adversary sends a single tampering query $(\mathcal{T}, \mathcal{B}, f)$.

- Let $\tau \in \mathbb{N}$ and $\mathcal{T} = \{\beta_1, \dots, \beta_\tau\}$, and write $\xi(i)$ for the index such that $\beta_i \in \mathcal{B}_{\xi(i)}$ (i.e., the i -th share of the reconstruction is tampered by the $\xi(i)$ -th tampering function).
- We define some subsets starting from \mathcal{T} . Call

$$\mathcal{T}_0^* = \bigcup_{\beta \in \mathcal{T}_{1k_1}} \mathcal{B}_{\xi(\beta)} \quad \text{and} \quad \mathcal{T}_0 = \mathcal{T}_0^* \cap \mathcal{T}.$$

Then, use the above to define

$$\mathcal{T}_1 = \mathcal{T} \setminus \mathcal{T}_0 \quad \text{and} \quad \mathcal{T}_1^* = \bigcup_{\beta \in \mathcal{T}_1} \mathcal{B}_{\xi(\beta)}.$$

and let $\mathcal{T}^* = \mathcal{T}_0^* \cup \mathcal{T}_1^*$.

Note that, with the above notation, we can write:

$$\bigcup_{\beta \in \mathcal{T}_{1_{k_1}}} \mathcal{B}_{\xi(\beta)} = \bigcup_{\beta \in \mathcal{T}_0} \mathcal{B}_{\xi(\beta)}.$$

Moreover, \mathcal{T}_0 and \mathcal{T}_1 are defined in such a way that $|\mathcal{T}_0| \geq k_1$ and, if $\mathcal{B}_i \cap \mathcal{T} \neq \emptyset$, then either $\mathcal{B}_i \cap \mathcal{T}_0 \neq \emptyset$ or $\mathcal{B}_i \cap \mathcal{T}_1 \neq \emptyset$, but not both. In this way, we also obtain that $\mathcal{T}_0^* \cap \mathcal{T}_1^* = \emptyset$.

Finally recall that the adversary sends leakage queries $(\mathcal{B}^{(1)}, g^{(1)}), \dots, (\mathcal{B}^{(q)}, g^{(q)})$, for $q \in \text{poly}(\lambda)$, and by condition (i) in the definition of semi-adaptive admissibility (cf. Def. 7) we have that for all $\mathcal{B}^* \in \bigcup_{i \in [q]} \mathcal{B}^{(i)}$ either (1) $\exists j \in \mathcal{T} : \mathcal{B}^* \subseteq \mathcal{B}_{\xi(j)}$, or (2) $\forall j \in \mathcal{T} : \mathcal{B}^* \cap \mathcal{B}_{\xi(j)} = \emptyset$.

Hybrid 1. Let $\mathbf{H}_1(\lambda, b)$ be the same as $\mathbf{H}_0(\lambda, b)$ except for the shares of s_0 being re-sampled at the end of the leakage phase. Namely, in $\mathbf{H}_1(\lambda, b)$ we sample $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ such that $(s_{0,\beta})_{\beta \in \mathcal{T}_0^*}, (s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ are valid shares of s_0 and consistent with the leakage. Then, we answer to A 's queries as follows:

- upon receiving a leakage query, use $(s_{0,1}, s_{1,1}), \dots, (s_{0,n}, s_{1,n})$ to compute the answer;
- upon receiving the tampering query, use $(s_{0,\beta}, s_{1,\beta})_{\beta \in \mathcal{T}_0^*}, (s_{0,\beta}^*, s_{1,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ to compute the answer.

Lemma 2. For $b \in \{0, 1\}$, $\Delta(\mathbf{H}_0(\lambda, b), \mathbf{H}_1(\lambda, b)) = 0$.

Proof. Let $(\mathbf{S}_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ and $(\mathbf{S}_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ be the random variables for the values $(s_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ and $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ in experiments \mathbf{H}_0 and \mathbf{H}_1 . More in details, the random variable $(\mathbf{S}_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ comes from the distribution of the shares $(s_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ conditioned on the fixed values $(s_{0,\beta})_{\beta \in \mathcal{T}_0^*}$ and the overall leakage Λ . We claim that $(\mathbf{S}_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ and $(\mathbf{S}_{1,\beta})_{\beta \in \mathcal{T}_0^*}$ are independent conditioned on the leakage Λ . This is because the random variables $(\mathbf{S}_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ and $(\mathbf{S}_{1,\beta})_{\beta \in \mathcal{T}_0^*}$ are independent in isolation, and, by condition (i) in the definition of semi-adaptive admissibility, none of the leakage functions leaks simultaneously from a share in \mathcal{T}_0^* and a share in \mathcal{T}_1^* . The latter holds as otherwise there would exist $\mathcal{B}^* \in \bigcup_{i \in [q]} \mathcal{B}^{(i)}$ such that $\mathcal{T}_1^* \cap \mathcal{B}^* \neq \emptyset$ and $\mathcal{T}_0^* \cap \mathcal{B}^* \neq \emptyset$, and therefore: (1) $\forall j \in \mathcal{T} : \mathcal{B}^* \not\subseteq \mathcal{B}_{\xi(j)}$, and (2) $\exists j \in \mathcal{T} : \mathcal{B}^* \cap \mathcal{B}_{\xi(j)} \neq \emptyset$. Finally, by Lemma 1, we can conclude that the two random variables are independent even conditioned on the leakage.

For any string \bar{s} , let $\mathbf{B}_0^{\bar{s}}$ and $\mathbf{B}_1^{\bar{s}}$ be, respectively, the event that $(\mathbf{S}_{0,\beta})_{\beta \in \mathcal{T}_1^*} = \bar{s}$ and $(\mathbf{S}_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*} = \bar{s}$. Then:

$$\begin{aligned} & \mathbb{P}[\mathbf{H}_0(\lambda, b) = 1] - \mathbb{P}[\mathbf{H}_1(\lambda, b) = 1] \\ &= \sum_{\bar{s}} \mathbb{P}[\mathbf{B}_0^{\bar{s}}] \mathbb{P}[\mathbf{H}_0(\lambda, b) = 1 | \mathbf{B}_0^{\bar{s}}] - \sum_{\bar{s}} \mathbb{P}[\mathbf{B}_1^{\bar{s}}] \mathbb{P}[\mathbf{H}_1(\lambda, b) = 1 | \mathbf{B}_1^{\bar{s}}] \\ &= \sum_{\bar{s}} \mathbb{P}[\mathbf{B}_0^{\bar{s}}] (\mathbb{P}[\mathbf{H}_0(\lambda, b) = 1 | \mathbf{B}_0^{\bar{s}}] - \mathbb{P}[\mathbf{H}_1(\lambda, b) = 1 | \mathbf{B}_1^{\bar{s}}]) \end{aligned} \quad (7)$$

$$= 0, \quad (8)$$

where Eq. (7) holds because of $(\mathbf{S}_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ is re-sampled from the distribution of the $(s_{0,\beta})_{\beta \in \mathcal{T}_1^*}$ conditioned on the measured leakage Λ and fixed $(s_{0,\beta})_{\beta \in \mathcal{T}_0^*}$ and moreover it is independent of $(s_{1,\beta})_{\beta \in \mathcal{T}_0^*}$ thus is distributed exactly as the conditional distribution of the $(\mathbf{S}_{0,\beta})_{\beta \in \mathcal{T}_1^*}$. The Eq. (8) holds because, once fixed the value of \bar{s} , if both $\mathbf{B}_0^{\bar{s}}$ and $\mathbf{B}_1^{\bar{s}}$ happen, then $(\mathbf{S}_{0,\beta})_{\beta \in \mathcal{T}_1^*} = \bar{s} = (\mathbf{S}_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ and the two hybrids are the same. \square

Hybrid 2. Let $\mathbf{H}_2(\lambda, b)$ be the same as $\mathbf{H}_1(\lambda, b)$ except for the leakage being performed on fake shares of s_0 . Namely, compute $(\hat{s}_0, \hat{s}_1) \leftarrow \text{Share}_2(0)$, let $\hat{s}_i = (\hat{s}_{0,i}, s_{1,i})$ where $(\hat{s}_{0,1}, \dots, \hat{s}_{0,n}) \leftarrow \text{Share}_0(\hat{s}_0)$, and sample the shares $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ of \mathbf{H}_1 such that $(\hat{s}_{0,\beta})_{\beta \in \mathcal{T}_0^*}, (s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ are valid shares of s_0 and consistent with the leakage. Then:

- upon receiving a leakage query, use $(\hat{s}_{0,1}, s_{1,1}), \dots, (\hat{s}_{0,n}, s_{1,n})$ to compute the answer;
- upon receiving the tampering query, use $(\hat{s}_{0,\beta}, s_{1,\beta})_{\beta \in \mathcal{T}_0^*}, (s_{0,\beta}^*, s_{1,\beta})_{\beta \in \mathcal{T}_1^*}$ to compute the answer.

Lemma 3. For $b \in \{0, 1\}$, $\Delta((\mathbf{H}_1(\lambda, b), \mathbf{H}_2(\lambda, b))) \leq \epsilon_0(\lambda)$.

Proof. By reduction to leakage resilience of Σ_0 . Suppose towards contradiction that there exist $b \in \{0, 1\}$, messages m_0, m_1 , and an adversary \mathbf{A} able to tell apart $\mathbf{H}_1(\lambda, b)$ and $\mathbf{H}_2(\lambda, b)$ with advantage more than $\epsilon_0(\lambda)$. Let (s_0, s_1) and (\hat{s}_0, \hat{s}_1) be, respectively, a secret sharing of m_b and of the all-zero string under Σ_2 . Consider the following reduction trying to distinguish a secret sharing of s_0 and a secret sharing of \hat{s}_0 under Σ_0 , where we call s_0^{target} the target secret sharing in the leakage oracle.

Adversary $\hat{\mathbf{A}}^{\mathcal{O}_{\text{leak}}((s_{0,i}^{\text{target}})_{i \in [n]}, \cdot)}(1^\lambda)$:

1. Sample $(s_{1,1}, \dots, s_{1,n}) \leftarrow \text{Share}_1(s_1)$ and run the experiment as in \mathbf{H}_1 with the adversary \mathbf{A} ; upon receiving each leakage function, hard-code into it the shares of s_1 and forward it to the leakage oracle.
2. Eventually, the adversary sends its tampering query. Obtain from the challenger the shares $(s_{0,\beta}^{\text{target}})_{\beta \in \mathcal{T}_0^*}$ (using the augmented property from Thm. 2).
3. For all $\beta \in \mathcal{T}_0$, compute $(\tilde{s}_{0,j}, \tilde{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}} = f_{\xi(\beta)}((s_{0,j}^{\text{target}}, s_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}})$ and compute $\tilde{s}_1 = \text{Rec}_1((\tilde{s}_{1,\beta})_{\beta \in \mathcal{T}_{k_1}})$.
4. Sample $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ as described in \mathbf{H}_2 and compute \tilde{s}_0 as follows: for all $\beta \in \mathcal{T}_1$, let $(\tilde{s}_{0,j}, \tilde{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}} = f_{\xi(\beta)}((s_{0,j}^*, s_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}})$ and $\tilde{s}_0 = \text{Rec}_0((\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}})$.
5. Compute the value $\tilde{m} = \text{Rec}_2(\tilde{s}_0, \tilde{s}_1)$. In case $\tilde{m} \in \{m_0, m_1\}$ return \diamond to \mathbf{A} , and else return \tilde{m} .
6. Output the same as \mathbf{A} .

For the analysis, note that the reduction is perfect. In particular, the reduction perfectly simulates \mathbf{H}_1 when $(s_{0,i}^{\text{target}})_{i \in [n]}$ is a secret sharing of s_0 and perfectly simulates \mathbf{H}_2 when $(s_{0,i}^{\text{target}})_{i \in [n]}$ is a secret sharing of \hat{s}_0 . Moreover, the leakage

requested by \mathbf{A} is forwarded to the leakage oracle of $\hat{\mathbf{A}}$ and perfectly simulated by it. Finally, the reduction gets in full $(s_{0,\beta}^{\text{target}})_{\beta \in \mathcal{T}_0^*}$, which allows it to compute \tilde{s}_1 , and computes \tilde{s}_0 by sampling the values $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ as in \mathbf{H}_1 .

Let us now analyze the admissibility of $\hat{\mathbf{A}}$. The only leakage performed by $\hat{\mathbf{A}}$ is the one requested by \mathbf{A} , and augmented leakage resilience can be obtained with 1 extra bit of leakage by Thm. 2. Finally, since $|\mathcal{T}_0^*| \leq k_1(k_1 - 1) \leq k$, it follows that if \mathbf{A} is $(k_1 - 1, \ell, 1)$ -BLTA, $\hat{\mathbf{A}}$ is $(k, \ell + 1)$ -BLA. \square

Hybrid 3. Let $\mathbf{H}_3(\lambda, b)$ be the same as $\mathbf{H}_2(\lambda, b)$ except for the shares of s_1 being re-sampled at the end of the leakage phase. Namely, in $\mathbf{H}_3(\lambda, b)$ we sample $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ such that (1) the shares $(s_{1,\beta})_{\beta \in \mathcal{T}_0^*}$ and $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ agree with the same leakage and the same reconstructed secret s_1 , and (2) for all $\beta \in \mathcal{T}_0$, applying the tampering function $f_{\xi(\beta)}$ to $(\hat{s}_{0,j}, s_{1,j}^*)_{j \in \mathcal{B}_{\xi(\beta)}}$ or to $(\hat{s}_{0,j}, s_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}}$ leads to the same values $(\tilde{s}_{0,j})_{j \in \mathcal{B}_{\xi(\beta)}}$. Then, we answer to \mathbf{A} 's queries as follows:

- upon receiving a leakage query, use $(\hat{s}_{0,1}, s_{1,1}), \dots, (\hat{s}_{0,n}, s_{1,n})$ to compute the answer;
- upon receiving the tampering query, use $(\hat{s}_{0,\beta}, s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}, (s_{0,\beta}^*, s_{1,\beta})_{\beta \in \mathcal{T}_1^*}$ to compute the answer.

Lemma 4. For $b \in \{0, 1\}$, $\Delta(\mathbf{H}_2(\lambda, b), \mathbf{H}_3(\lambda, b)) = 0$.

Proof. The proof is similar to that of Lemma 2, and thus omitted.

Hybrid 4. Let $\mathbf{H}_4(\lambda, b)$ be the same as $\mathbf{H}_3(\lambda, b)$ except for the leakage being performed on fake shares of s_1 . Namely, let $(\hat{s}_{1,i})_{i \in [n]} \leftarrow \text{Share}_1(\hat{s}_1)$, where \hat{s}_1 comes from $\text{Share}_2(0)$ as in \mathbf{H}_2 . Then:

- upon receiving a leakage query, use $(\hat{s}_{0,1}, \hat{s}_{1,1}), \dots, (\hat{s}_{0,n}, \hat{s}_{1,n})$ to compute the answer;
- upon receiving the tampering query, use $(\hat{s}_{0,\beta}, s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}, (s_{0,\beta}^*, \hat{s}_{1,\beta})_{\beta \in \mathcal{T}_1^*}$ to compute the answer.

Lemma 5. For $b \in \{0, 1\}$, $\Delta(\mathbf{H}_3(\lambda, b), \mathbf{H}_4(\lambda, b)) \leq \epsilon_1(\lambda)$.

Proof. By reduction to the leakage resilience of Σ_1 . Suppose towards contradiction that there exist $b \in \{0, 1\}$, messages m_0, m_1 , and an adversary \mathbf{A} able to tell apart $\mathbf{H}_3(\lambda, b)$ and $\mathbf{H}_4(\lambda, b)$ with advantage more than $\epsilon_1(\lambda)$. Let (s_0, s_1) and (\hat{s}_0, \hat{s}_1) be, respectively, a secret sharing of m_b and of the all-zero string under Σ_2 . Consider the following reduction trying to distinguish a secret sharing of s_1 and a secret sharing of \hat{s}_1 under Σ_1 , where we call s_1^{target} the target secret sharing in the leakage oracle.

Adversary $\hat{\mathbf{A}}^{\mathcal{O}_{\text{leak}}((s_{1,i}^{\text{target}})_{i \in [n]}, \cdot)}(1^\lambda)$:

1. Sample $(\hat{s}_{0,1}, \dots, \hat{s}_{0,n}) \leftarrow \text{Share}_0(\hat{s}_0)$ and run the experiment as in \mathbf{H}_3 with the adversary \mathbf{A} ; upon receiving each leakage function, hard-code into it the shares of \hat{s}_0 and forward it to the leakage oracle.
2. Eventually, the adversary sends its tampering query $(\mathcal{T}, \mathcal{B}, f)$.

3. Sample $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ as in **H₂**. In particular, recall that we can sample these shares as a function of just the shares $(s_{0,\beta})_{\beta \in \mathcal{T}_0^*}$ and the leakage. Then, set

$$s'_{0,\beta} := \begin{cases} \hat{s}_{0,\beta} & \text{if } \beta \in \mathcal{T}_0^*, \\ s_{0,\beta}^* & \text{if } \beta \in \mathcal{T}_1^*. \end{cases}$$

Note that this is well defined since $\mathcal{T}_0^* \cap \mathcal{T}_1^* = \emptyset$.

4. For all $i \in [t]$, construct the leakage function g_i that, given as input $(s_{1,\beta}^{\text{target}})_{\beta \in \mathcal{B}_i}$, computes $(\tilde{s}_{0,\beta}, \tilde{s}_{1,\beta})_{\beta \in \mathcal{B}_i} = f_j((s'_{0,\beta}, s_{1,\beta}^{\text{target}})_{\beta \in \mathcal{B}_i})$ and outputs $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{B}_i}$. Send $(\mathcal{B}, (g_1, \dots, g_t))$ to the leakage oracle obtaining values $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$.
5. Sample the values $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ as in **H₃** using $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$ and the leakage.
6. For all $j \in \mathcal{T}_0$, compute $(\tilde{s}_{0,\beta}, \tilde{s}_{1,\beta})_{\beta \in \mathcal{B}_{\xi(j)}} = f_j((s'_{0,\beta}, s_{1,\beta}^*)_{\beta \in \mathcal{B}_{\xi(j)}})$; then, compute $s_0 = \text{Rec}_0((\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}})$ and $s_1 = \text{Rec}_1((\tilde{s}_{1,\beta})_{\beta \in \mathcal{T}_{1k_1}})$ and let $\tilde{m} = \text{Rec}_2(s_0, s_1)$. In case $\tilde{m} \in \{m_0, m_1\}$ return \diamond to **A**, and else return \tilde{m} to **A**.
7. Output the same as **A**.

For the analysis, note that the reduction is perfect. In particular, the reduction perfectly simulates **H₃** when $(s_{1,i}^{\text{target}})_{i \in [n]}$ is a secret sharing of s_1 and perfectly simulates **H₄** when $(s_{1,i}^{\text{target}})_{i \in [n]}$ is a secret sharing of \hat{s}_1 . Moreover, the leakage requested by the adversary **A** is forwarded to the leakage oracle of $\hat{\mathbf{A}}$ and perfectly simulated by it. Finally, the reduction obtains all the shares $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$, and thus it is able to both compute \tilde{s}_0 and sample the values $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$.

Let us now analyze the admissibility of $\hat{\mathbf{A}}$. The only leakage performed by $\hat{\mathbf{A}}$ is the one requested by **A** in step 1 plus the one needed in order to get the values $(\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}^*}$ in step 4; summing up, the overall leakage performed by $\hat{\mathbf{A}}$ is:

$$\ell + \sum_{\beta \in \mathcal{T}^*} \log |\mathcal{S}_{0,\beta}| \leq \ell + \sum_{i \in [n]} \log |\mathcal{S}_{0,i}| \leq \ell + n\sigma_0,$$

where the last inequality follows by the fact that $\log |\mathcal{S}_{0,i}| \leq \sigma_0$ for all $i \in [n]$. Therefore, we can conclude that $\hat{\mathbf{A}}$ is $(k_1 - 1, \ell + n\sigma_0)$ -BLA. \square

Final step. Finally, we show:

Lemma 6. $\Delta(\mathbf{H}_4(\lambda, 0), \mathbf{H}_4(\lambda, 1)) \leq \epsilon_2(\lambda)$.

Proof. By reduction to non-malleability of Σ_2 . Suppose by contradiction that there exist messages m_0, m_1 and an adversary **A** telling apart $\mathbf{H}_4(\lambda, 0)$ and $\mathbf{H}_4(\lambda, 1)$ with advantage more than $\epsilon_2(\lambda)$. Fix values $(\hat{s}_i)_{i \in [n]} = ((\hat{s}_{0,i}, \hat{s}_{1,i})_{i \in [n]})$ and (s_0, s_1) being either a (2-out-of-2) secret sharing of m_0 or of m_1 . Consider the following reduction:

$$\text{Adversary } \hat{\mathbf{A}}_{\mathcal{O}_{\text{nms}}((s_0^{\text{target}}, s_1^{\text{target}}), \cdot)}(1^\lambda):$$

1. Run the experiment as in \mathbf{H}_4 with the adversary A ; upon receiving each leakage function, answer using the values $(\hat{s}_i)_{i \in [n]}$.
2. Upon input the tampering query $(\mathcal{T}, \mathcal{B}, f)$, construct the following two tampering functions:
 - Function f_0 , upon input s_0 , samples $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ as in \mathbf{H}_2 ; notice that the reduction knows all the information needed to re-sample the shares, as in particular it samples $(s_{0,\beta})_{\beta \in [n]}$ and simulates the leakage. Then, f_0 computes $(\tilde{s}_{0,j}, \tilde{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}} = f_{\xi(\beta)}((\hat{s}_{0,j}, \hat{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}})$ for all $\beta \in \mathcal{T}_0$ and $(\tilde{s}_{0,j}, \tilde{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}} = f_{\xi(\beta)}((s_{0,j}^*, \hat{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}})$ for all $\beta \in \mathcal{T}_1$ and outputs $\tilde{s}_0 = \text{Rec}_0((\tilde{s}_{0,\beta})_{\beta \in \mathcal{T}})$.
 - Function f_1 , upon input s_1 , samples $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ as in \mathbf{H}_3 . Then, f_1 computes $(\tilde{s}_{0,j}, \tilde{s}_{1,j})_{j \in \mathcal{B}_{\xi(\beta)}} = f_{\xi(\beta)}((\hat{s}_{0,j}, s_{1,j}^*)_{j \in \mathcal{B}_{\xi(\beta)}})$ for all $\beta \in \mathcal{T}_0$ and outputs $\tilde{s}_1 = \text{Rec}_1((\tilde{s}_{1,\beta})_{\beta \in \mathcal{T}})$.
3. Send (f_0, f_1) to the tampering oracle, receiving an answer \tilde{m} .
4. Return \tilde{m} to A and output the same as A .

For the analysis, note that the reduction is perfect. In particular, shares $(s_{0,\beta}^*)_{\beta \in \mathcal{T}_1^*}$ and $(s_{1,\beta}^*)_{\beta \in \mathcal{T}_0^*}$ are computed using s_0 and s_1 respectively; moreover, both \tilde{s}_0 and \tilde{s}_1 are computed as in experiment \mathbf{H}_4 and thus the tampering query is perfectly simulated. Finally, the leakage is computed using the fake shares $(\hat{s}_i)_{i \in [n]}$ as in \mathbf{H}_4 and thus, once again, perfectly simulated. The lemma follows.

Proof (Theorem 4). Follows by the above lemmas and the triangular inequality:

$$\begin{aligned}
& \Delta(\mathbf{H}_0(\lambda, 0), \mathbf{H}_0(\lambda, 1)) \\
& \leq \sum_{b \in \{0,1\}} \sum_{i \in [4]} \Delta(\mathbf{H}_{i-1}(\lambda, b), \mathbf{H}_i(\lambda, b)) + \Delta(\mathbf{H}_4(\lambda, 0), \mathbf{H}_4(\lambda, 1)) \\
& \leq 2(\Delta(\mathbf{H}_1(\lambda, b), \mathbf{H}_2(\lambda, b)) + \Delta(\mathbf{H}_3(\lambda, b), \mathbf{H}_4(\lambda, b))) + \Delta(\mathbf{H}_4(\lambda, 0), \mathbf{H}_4(\lambda, 1)) \\
& \leq 2(\epsilon_0 + \epsilon_1) + \epsilon_2.
\end{aligned}$$

□

5.4 Instantiation

Using a previous construction of bounded leakage-resilient secret sharing scheme against joint leakage under adaptive partitioning, we obtain the following:

Corollary 3. *For every $\ell, n, \lambda \geq 0$, every $k \in O(\sqrt{\log n})$, and every access structure \mathcal{A} over n parties that can be described by a polynomial-size monotone span program for which authorized sets have size greater than k , there exists a $(k, \ell, 1, 2^{-\Omega(\lambda/\log(\lambda))})$ -BLR-NMSS with message length $\Omega(\lambda/\log(\lambda))$ realizing \mathcal{A} under semi-adaptive partitioning.*

Proof. By Thm. 4, we need to instantiate Σ_0 , Σ_1 , and Σ_2 . Using [26, Thm. 1] and [26, Cor. 2], we can take $\epsilon_0 = \epsilon_1 = 2^{-\Omega(\lambda/\log(\lambda))}$, $k \in O(\log n)$, and thus $k_1 \in O(\sqrt{\log n})$, $\sigma_0 = \text{poly}(\lambda)$ and any $\ell_0, \ell_1 > 0$. As for Σ_2 , we can take the split-state non-malleable code in [27, Thm. 1.12], which achieves error $2^{-\Omega(\lambda/\log(\lambda))}$. □

6 Applications

6.1 Lower Bounds for Non-Malleable Secret Sharing

Combining our result from Thm. 3 with the lower bound of Nielsen and Simkin [29], we obtain a lower bound on the share size and randomness complexity of non-malleable secret sharing schemes. In particular, we obtain the following:

Corollary 4. *Any τ -out-of- n $(1, 1, \epsilon)$ -NMSS must satisfy*

$$\sigma \geq \frac{(\log(1/\epsilon) - 1)(1 - \tau/n)}{\hat{\tau}},$$

where $\hat{\tau}$ is the number of shares needed to reconstruct the full vector of shares and σ is the bit-length of each share.

Observe that $\hat{\tau}$ is a simplified notion of entropy. If $\tau = \hat{\tau}$, then any authorized set can reconstruct all remaining shares, meaning that those shares have no entropy left.

6.2 Bounded-Time Non-Malleability

Here, we revisit the compiler from Ostrovsky *et al.* [30] in the setting of non-malleable secret sharing against joint tampering.

The basic idea is as follows. First, we commit to the message m using random coins r , thus obtaining a cryptographic commitment c . Then, we secret share the string $m||r$ using an auxiliary secret sharing scheme Σ , thus obtaining shares s_1, \dots, s_n . The final share of the i -th party is set to be $s_i^* = (c, s_i)$. Given an authorized set \mathcal{I} , the reconstruction first checks that all commitments in $s_{\mathcal{I}}^*$ are equal, and then uses $s_{\mathcal{I}}$ to recover $m||r$, and verifies consistency of the commitments. If any of these checks fails, it outputs \perp ; else, it returns m .

The original analysis by Ostrovsky *et al.* shows that if Σ is a 2-out-of-2 secret sharing that is bounded leakage-resilient, statistically one-time non-malleable, and further satisfies additional non-standard properties, then Σ^* is continuously non-malleable. In a follow up work, Brian *et al.* [11] proved that the additional properties on Σ can be avoided if one assumes that Σ satisfies a stronger form of leakage resilience known as *noisy* leakage resilience, and further extended the original analysis to any value $n \geq 2$ and for arbitrary access structures.

Both the proofs in [30, 11] are for the setting of independent tampering. The theorem below says that the same construction works also in the case of joint p -time tampering under selective/semi-adaptive partitioning as long as Σ tolerates joint bounded leakage resilience, where there is a natural trade off between the leakage bound and the number of tampering queries. The main idea behind the proof is to reduce the security of Σ^* to that of Σ , where the bounded leakage is used to simulate multiple tampering queries. The main difference with the original proof is that we need a small leakage for each tampering query, and thus the analysis only works in case the number of tampering queries is *a priori*

Let **Commit** be a non-interactive commitment scheme with message space \mathcal{M} , randomness space \mathcal{R} and commitment space \mathcal{C} . Let $\Sigma = (\text{Share}, \text{Rec})$ be an auxiliary secret sharing scheme realizing access structure \mathcal{A} with message space $\mathcal{M} \times \mathcal{R}$ and share space $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$. Define the following secret sharing scheme $\Sigma^* = (\text{Share}^*, \text{Rec}^*)$ with message space \mathcal{M} and share space $\mathcal{S}^* = \mathcal{S}_1^* \times \dots \times \mathcal{S}_n^*$, where, for each $i \in [n]$, we have $\mathcal{S}_i^* = \mathcal{C} \times \mathcal{S}_i$

Sharing algorithm Share*: Upon input a value $m \in \mathcal{M}$, sample random coins $r \leftarrow \mathcal{R}$ and compute $c = \text{Commit}(m; r)$ and $(s_1, \dots, s_n) \leftarrow \text{Share}(m||r)$. Return the shares $s^* = (s_1^*, \dots, s_n^*)$ where, for each $i \in [n]$, $s_i^* = (c, s_i)$.

Reconstruction algorithm Rec*: Upon input shares $(s_i^*)_{i \in \mathcal{I}}$, parse $s_i^* = (c_i, s_i)$ for each $i \in \mathcal{I}$. Hence, proceed as follows.

1. If $\exists i_1, i_2 \in \mathcal{I}$ for which $c_{i_1} \neq c_{i_2}$, return \perp ; else, let the input shares be $s_i^* = (c, s_i)$.
2. Run $m||r = \text{Rec}((s_i)_{i \in \mathcal{I}})$; if the outcome equals \perp , return \perp .
3. If $c = \text{Commit}(m; r)$, return m ; else, return \perp .

Fig. 2. Compiler for obtaining bounded-time non-malleability against joint tampering.

bounded. Moreover, in the case of semi-adaptive partitioning, we need to make sure that the leakage performed by the reduction does not violate condition (i) in the definition of semi-adaptive admissibility (cf. Def. 7); intuitively, the latter holds thanks to the fact that the tampering queries chosen by the attacker must satisfy condition (ii) in Def. 7. We refer to the full version of the paper for the details [10].

Theorem 5. *Let $n \in \mathbb{N}$ and let \mathcal{A} be an arbitrary access structure for n parties without singletons. Assume that:*

1. *Commit is a perfectly binding and computationally hiding non-interactive commitment;*
2. *Σ is a n -party k -joint ℓ -bounded leakage-resilient one-time non-malleable secret sharing scheme realizing access structure \mathcal{A} against joint semi-adaptive (resp., selective) partitioning with information-theoretic security and with message space \mathcal{M} such that $|\mathcal{M}| \in \omega(\log(\lambda))$.*

Then, the secret sharing scheme Σ^ described in Fig. 2 is a n -party k -joint p -time non-malleable secret sharing scheme realizing access structure \mathcal{A} against joint semi-adaptive (resp., selective) partitioning with computational security, as long as $\ell = p \cdot (\gamma + n) + 1$, where $\gamma = \log |\mathcal{C}|$ is the size of a commitment.*

7 Conclusions

We presented new constructions of non-malleable secret sharing schemes against joint tampering with the shares, both in the setting of selective and adaptive partitioning.

Our constructions for selective partitioning are for threshold access structures and tolerate joint tampering with maximal subsets of unauthorized parties, *i.e.*, of size equal to the privacy threshold. Our construction for adaptive partitioning is for general access structures, but tolerates joint tampering with much smaller subsets of size $k \in O(\sqrt{\log n})$ (where n is the number of parties) and under some restrictions on the way the partitions are determined by the attacker. Removing the latter limitation is an intriguing open question.

The above results hold for any *a priori* fixed bound $p > 0$ on the number of tampering queries, and under computational assumptions. We leave it as an open problem to design *continuously* non-malleable (*i.e.*, for $p = p(\lambda)$ being an arbitrary polynomial in the security parameter) secret sharing schemes tolerating joint tampering under selective/adaptive partitioning.

Another interesting question would be to improve the rate, *i.e.*, the ratio between message size and maximal size of a share, for non-malleable secret sharing against joint tampering. Note that, in the computational setting, it is always possible to boost the rate as follows: First, share the secret key $\kappa \in \{0, 1\}^\lambda$ of an authenticated symmetric encryption using a secret sharing scheme with poor rate, obtaining shares s_1, \dots, s_n ; hence, encrypt the message m using κ , obtaining a ciphertext c , and define the final i -th share to be $s_i^* = (c, s_i)$. Such a rate-optimizing compiler was originally analyzed in the setting of non-malleable codes [19, 1, 16], and more recently in the setting of non-malleable secret sharing against independent tampering [21]. While this transformation may be proven secure even in the setting of joint tampering with the shares, it yields a rate asymptotically approaching one, which is still far from the optimal share size of $O(\mu/n)$ [25] (where μ is the message size).

Acknowledgments

We thank Ashutosh Kumar for clarifications on the tampering model in [23] and for pointing out an issue in a previous version of the proof of Thm. 4 (leading to the restriction of semi-adaptive partitioning).

References

1. Aggarwal, D., Agrawal, S., Gupta, D., Maji, H.K., Pandey, O., Prabhakaran, M.: Optimal computational split-state non-malleable codes. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016-A, Part II. LNCS, vol. 9563, pp. 393–417. Springer, Heidelberg (Jan 2016)
2. 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)
3. Aggarwal, D., Dodis, Y., Kazana, T., Obremski, M.: Non-malleable reductions and applications. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 459–468. ACM Press (Jun 2015)

4. Aggarwal, D., Döttling, N., Nielsen, J.B., Obremski, M., Purwanto, E.: Continuous non-malleable codes in the 8-split-state model. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 531–561. Springer, Heidelberg (May 2019)
5. 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)
6. Aggarwal, D., Kazana, T., Obremski, M.: Inception makes non-malleable codes stronger. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part II. LNCS, vol. 10678, pp. 319–343. Springer, Heidelberg (Nov 2017)
7. Aggarwal, D., Obremski, M.: A constant-rate non-malleable code in the split-state model. Cryptology ePrint Archive, Report 2019/1299 (2019), <https://eprint.iacr.org/2019/1299>
8. Badrinarayanan, S., Srinivasan, A.: Revisiting non-malleable secret sharing. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 593–622. Springer, Heidelberg (May 2019)
9. Blakley, G.R.: Safeguarding cryptographic keys. Proceedings of AFIPS 1979 National Computer Conference 48, 313–317 (1979)
10. 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>
11. Brian, G., Faonio, A., Venturi, D.: Continuously non-malleable secret sharing for general access structures. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019, Part II. LNCS, vol. 11892, pp. 211–232. Springer, Heidelberg (Dec 2019)
12. Carpentieri, M., De Santis, A., Vaccaro, U.: Size of shares and probability of cheating in threshold schemes. In: Hellese, T. (ed.) EUROCRYPT’93. LNCS, vol. 765, pp. 118–125. Springer, Heidelberg (May 1994)
13. Chattopadhyay, E., Goyal, V., Li, X.: Non-malleable extractors and codes, with their many tampered extensions. In: Wichs, D., Mansour, Y. (eds.) 48th ACM STOC. pp. 285–298. ACM Press (Jun 2016)
14. Chattopadhyay, E., Li, X.: Non-malleable extractors and codes for composition of tampering, interleaved tampering and more. Cryptology ePrint Archive, Report 2018/1069 (2018), <https://eprint.iacr.org/2018/1069>
15. Cheraghchi, M., Guruswami, V.: Non-malleable coding against bit-wise and split-state tampering. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 440–464. Springer, Heidelberg (Feb 2014)
16. Coretti, S., Faonio, A., Venturi, D.: Rate-optimizing compilers for continuously non-malleable codes. In: Deng, R.H., Gauthier-Umaña, V., Ochoa, M., Yung, M. (eds.) ACNS 19. LNCS, vol. 11464, pp. 3–23. Springer, Heidelberg (Jun 2019)
17. Dziembowski, S., Kazana, T., Obremski, M.: Non-malleable codes from two-source extractors. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 239–257. Springer, Heidelberg (Aug 2013)
18. Dziembowski, S., Pietrzak, K.: Intrusion-resilient secret sharing. In: 48th FOCS. pp. 227–237. IEEE Computer Society Press (Oct 2007)
19. Dziembowski, S., Pietrzak, K., Wichs, D.: Non-malleable codes. In: Yao, A.C.C. (ed.) ICS 2010. pp. 434–452. Tsinghua University Press (Jan 2010)
20. Faonio, A., Nielsen, J.B., Simkin, M., Venturi, D.: Continuously non-malleable codes with split-state refresh. In: Preneel, B., Vercauteren, F. (eds.) ACNS 18. LNCS, vol. 10892, pp. 121–139. Springer, Heidelberg (Jul 2018)

21. Faonio, A., Venturi, D.: Non-malleable secret sharing in the computational setting: Adaptive tampering, noisy-leakage resilience, and improved rate. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 448–479. Springer, Heidelberg (Aug 2019)
22. Faust, S., Mukherjee, P., Nielsen, J.B., Venturi, D.: Continuous non-malleable codes. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 465–488. Springer, Heidelberg (Feb 2014)
23. Goyal, V., Kumar, A.: Non-malleable secret sharing. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) 50th ACM STOC. pp. 685–698. ACM Press (Jun 2018)
24. Goyal, V., Kumar, A.: Non-malleable secret sharing for general access structures. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 501–530. Springer, Heidelberg (Aug 2018)
25. Krawczyk, H.: Secret sharing made short. In: Stinson, D.R. (ed.) CRYPTO’93. LNCS, vol. 773, pp. 136–146. Springer, Heidelberg (Aug 1994)
26. 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)
27. Li, X.: Improved non-malleable extractors, non-malleable codes and independent source extractors. In: Hatami, H., McKenzie, P., King, V. (eds.) 49th ACM STOC. pp. 1144–1156. ACM Press (Jun 2017)
28. Liu, F.H., Lysyanskaya, A.: Tamper and leakage resilience in the split-state model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 517–532. Springer, Heidelberg (Aug 2012)
29. Nielsen, J.B., Simkin, M.: Lower bounds for leakage-resilient secret sharing. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part I. LNCS, vol. 12105, pp. 556–577. Springer, Heidelberg (May 2020)
30. Ostrovsky, R., Persiano, G., Venturi, D., Visconti, I.: Continuously non-malleable codes in the split-state model from minimal assumptions. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 608–639. Springer, Heidelberg (Aug 2018)
31. Rabin, T., Ben-Or, M.: Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In: 21st ACM STOC. pp. 73–85. ACM Press (May 1989)
32. Shamir, A.: How to share a secret. *Communications of the Association for Computing Machinery* 22(11), 612–613 (Nov 1979)
33. 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)