

Continuously Non-Malleable Secret Sharing for General Access Structures*

Gianluca Brian^{†1}, Antonio Faonio^{‡2}, and Daniele Venturi^{†1}

¹*Department of Computer Science, Sapienza University of Rome, Rome, Italy*

²*IMDEA Software Institute, Madrid, Spain*

September 24, 2019

Abstract

We study leakage-resilient continuously non-malleable secret sharing, as recently introduced by Faonio and Venturi (CRYPTO 2019). In this setting, an attacker can continuously tamper and leak from a target secret sharing of some message, with the goal of producing a modified set of shares that reconstructs to a message related to the originally shared value. Our contributions are two fold.

- In the plain model, assuming one-to-one one-way functions, we show how to obtain noisy-leakage-resilient continuous non-malleability for arbitrary access structures, in case the attacker can continuously leak from and tamper with all of the shares independently.
- In the common reference string model, we show how to obtain a new flavor of security which we dub bounded-leakage-resilient continuous non-malleability under selective k -partitioning. In this model, the attacker is allowed to partition the target n shares into any number of non-overlapping blocks of maximal size k , and then can continuously leak from and tamper with the shares within each block jointly. Our construction works for arbitrary access structures, and assuming (doubly enhanced) trapdoor permutations and collision-resistant hash functions, we achieve a concrete instantiation for $k \in O(\log n)$.

Prior to our work, there was no secret sharing scheme achieving continuous non-malleability against joint tampering, and the only known scheme for independent tampering was tailored to threshold access structures.

Keywords: Secret sharing, Non-malleability, Leakage resilience.

*An abridged version of this paper appears in the proceeding of the 17th Theory of Cryptography Conference (TCC 2019). This is the full version.

[†]Supported in part by the research projects “PRIVacy-preserving, Security, and MACHine-learning techniques for healthcare applications (PRISMA)” and “Protect yourself and your data when using social networks”, both funded by Sapienza University of Rome, and in part by the MIUR under grant “Dipartimenti di eccellenza 2018-2022” of the Computer Science Department of Sapienza University of Rome.

[‡]Supported by the Spanish Government through the projects Datamantium (ref. RTC-2016-4930-7), SCUM (RTI2018-102043-B-I00), and ERC2018-092822, and by the Madrid Regional Government under project BLOQUES (ref. S2018/TCS-4339).

Contents

1	Introduction	1
1.1	Our Contributions	1
1.2	Related Work	2
1.3	Organization	3
2	Standard Definitions	3
2.1	Secret Sharing Schemes	4
2.2	Non-Interactive Commitments	5
2.3	Non-Interactive Zero Knowledge	6
2.4	Randomness Extractors	6
3	Continuous Tampering under Selective Partitioning	7
3.1	The Definition	7
3.2	On Augmented Non-Malleability	9
3.3	Related Notions	9
4	Construction in the CRS Model	10
4.1	Description of the Scheme	10
4.2	Proof Overview	12
4.3	Security Analysis	14
4.4	Concrete Instantiation	23
5	Construction in the Plain Model	24
5.1	Description of the Scheme	24
5.2	Proof Overview	24
5.3	Security Analysis	27
6	Statistical One-Time Non-Malleability with Noisy Leakage	34
6.1	Asymmetric Noisy-Leakage-Resilient Secret Sharing	34
6.2	Construction	35
6.3	Proof Overview	36
6.4	Security Analysis	36
7	Conclusions and Open Problems	41
A	Construction of Asymmetric Secret Sharing	45

1 Introduction

A non-malleable secret sharing for an access structure \mathcal{A} over n parties allows to share a secret message m into n shares $s = (s_1, \dots, s_n)$, in such a way that the following properties are guaranteed.

Privacy: No attacker given the shares belonging to an arbitrary unauthorized subset $\mathcal{U} \notin \mathcal{A}$ of the players can infer any information on m .

Non-malleability: No attacker tampering with all of the shares via some function $f \in \mathcal{F}$ within some family of allowed¹ modifications can generate a mauled secret sharing $\tilde{s} = f(s)$ that reconstructs to $\tilde{m} \neq m$ related to m .

Sometimes, non-malleability is considered together with leakage resilience. This means that the attacker can additionally leak partial information $g(s)$ from all of the shares (via functions $g \in \mathcal{G}$) before launching a tampering attack. Leakage resilience typically comes in one of two flavors: *bounded leakage* (i.e., there is a fixed upper bound on the maximum amount of information retrieved from the shares) or *noisy leakage* (i.e., the length of the retrieved information is arbitrary as long as it does not decrease the entropy of the shares by too much).

In this work we focus on leakage-resilient *continuous* non-malleability with *adaptive concurrent reconstruction*, as recently introduced by Faonio and Venturi [FV19].² Here, the attacker can (leak from and) tamper poly-many times with a target secret sharing using functions $f^{(q)} \in \mathcal{F}$ as above, and for each tampering query q it can also choose adaptively the reconstruction set $\mathcal{T}^{(q)} \in \mathcal{A}$ used to determine the reconstructed message. There are only two limitations: First, the attacker is computationally bounded; second, the experiment stops (we say it “self-destructs”) after the first tampering query yielding an invalid set of shares. Both limitations are *inherent* for continuous non-malleability [FMNV14, BS19, FV19].

The only known scheme achieving such a strong flavor of non-malleability is the one by Faonio and Venturi, which tolerates the families \mathcal{F} and \mathcal{G} of *independent tampering/leakage*, i.e. for each query q we have $f^{(q)} = (f_1^{(q)}, \dots, f_n^{(q)}) \in \mathcal{F}$ where $f_i^{(q)}$ gets as input the i -th share (and similarly $g^{(q)} = (g_1^{(q)}, \dots, g_n^{(q)}) \in \mathcal{G}$). The access structure \mathcal{A} supported by their construction is the τ -threshold access structure—i.e., any subset of at most τ players has no information about the message—with the caveat that reconstruction works with at least $\tau + 2$ shares, namely a *ramp* secret sharing, thus leaving a minimal gap between the reconstruction and privacy threshold. The following natural question arise:

Problem 1. Can we obtain leakage-resilient continuously non-malleable secret sharing against independent leakage/tampering, for general access structures?

Another open question is whether leakage-resilient continuous non-malleability is achievable for stronger tampering and leakage families \mathcal{F}, \mathcal{G} , e.g. in case the attacker can leak from and manipulate subsets of the shares jointly.

Problem 2. Can we obtain leakage-resilient continuously non-malleable secret sharing against joint leakage/tampering?

1.1 Our Contributions

We make significant progress towards solving the above problems. In particular, our first contribution is a positive answer to Problem 1:

¹It is easy to see that non-malleability is impossible for arbitrary (polynomial-time) tampering.

²From now on, we omit to explicitly mention the feature of adaptive concurrent reconstruction and simply talk about continuous non-malleability.

Theorem 1 (Informal). *Assuming one-to-one one-way functions, for any access structure \mathcal{A} over n parties there exists a noisy-leakage-resilient continuously non-malleable secret sharing scheme realizing \mathcal{A} against independent leakage and tampering, in the plain model.*

Our second contribution is a positive answer to Problem 2 assuming trusted setup, in the form of a common reference string (CRS). More in details, we put forward a new security notion for secret sharing dubbed continuous non-malleability under *selective k -partitioning*. This roughly means that the attacker, after seeing the CRS, must commit to a partition of the set $[n]$ into β (non-overlapping) blocks $(\mathcal{B}_1, \dots, \mathcal{B}_\beta)$ of size at most k ; hence, the adversary can jointly, and continuously, tamper with and leak from each collection $s_{\mathcal{B}_i}$ of the shares.³

Theorem 2 (Informal). *Assuming (doubly-enhanced) trapdoor permutations and collision-resistant hash functions, for any access structure \mathcal{A} over n parties there exists a bounded-leakage-resilient continuously non-malleable secret sharing scheme realizing \mathcal{A} against selective $O(\log n)$ -joint leakage and tampering in the CRS model.*

Prior to our work, we had secret sharing schemes unconditionally achieving security either against joint leakage [KMS18] or joint tampering [GK18a, GK18b], but nothing was known for both even in the much simpler case of one-time non-malleability.

1.2 Related Work

Non-malleable secret sharing was introduced by Goyal and Kumar [GK18a]. For any $\tau \leq n$, they showed how to realize τ -threshold access structures, against one-time tampering with either all of the shares independently, or jointly after partitioning the players into two non-overlapping blocks of size at most⁴ $\tau - 1$. In a subsequent work [GK18b], the same authors show how to extend the result for independent tampering to the case of arbitrary access structures; additionally, for the case of joint tampering, they provide a new scheme realizing the n -threshold access structure (i.e., an n -out-of- n secret sharing) in a stronger model where the attacker can partition the players into two possibly overlapping blocks of size at most $n - 1$. Srinivasan and Vasudevan [SV19] built the first non-malleable secret sharing schemes for general access structures against independent tampering, with non-zero rate⁵ (in fact, even constant rate in case of threshold access structures). Chattopadhyay *et al.* [CL18] construct non-malleable secret sharing for threshold access structures, against affine tampering composed with joint split-state tampering. Lin *et al.* [LCG⁺19] consider non-malleability against affine tampering in an adaptive setting where the adversary gets to see an unauthorized subset of the shares before launching a single tampering attack.

Badrinarayanan and Srinivasan [BS19] generalize non-malleability to p -time tampering attacks, where p is an a-priori upper bound on the number of tampering queries the adversary can ask. For each attempt, however, the reconstruction set \mathcal{T} must be chosen in advance at the beginning of the experiment. In this model, they show how to realize arbitrary access structures against independent tampering with all of the shares. Aggarwal *et al.* [ADN⁺19] were the first to consider p -time non-malleability under *non-adaptive* concurrent reconstruction, i.e. the attacker now can specify a different reconstruction set $\mathcal{T}^{(q)}$ during the q -th tampering query, although the sequence of sets $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(p)}$ must be chosen non-adaptively. Kumar, Meka, and

³The only restriction is that no block in the partition can contain an authorized set of players, otherwise trivial attacks are possible.

⁴An additional (artificial) requirement is that the size of the two blocks must be different in order for their technique to work.

⁵The rate refers to the asymptotic ratio between the maximal length of a share and that of the message.

Sahai [KMS18] pioneered bounded-leakage-resilient one-time non-malleable secret sharing for general access structures, against independent leakage and tampering with all of the shares.

In the special case of 2-threshold access structures over $n = 2$ parties, the notion of (leakage-resilient) non-malleable secret sharing collapses to that of split-state (leakage-resilient) non-malleable codes [DPW10, LL12, DKO13, ADL14, CG14, FMNV14, ADKO15b, ADKO15a, AAG⁺16, Li17, FNSV18, OPVV18, CFV19].

1.3 Organization

All of our constructions rely on standard cryptographic primitives, which we recall in §2 (together with some basic notation). The new model of continuous tampering under selective partitioning is presented in §3.

Our main constructions appear in §4 (for joint tampering in the CRS model) and §5–§6 (for independent tampering in the plain model), respectively; there, we also explain how to instantiate these constructions with concrete building blocks, thus establishing Thm. 1 and Thm. 2. Finally, in §7, we conclude the paper with a list of open problems and interesting directions for further research.

2 Standard Definitions

Basic notation. For a string x , we denote its length by $|x|$; if \mathcal{X} is a set, $|\mathcal{X}|$ represents the number of elements in \mathcal{X} . When x is chosen randomly in \mathcal{X} , we write $x \leftarrow_s \mathcal{X}$. When \mathbf{A} is a randomized algorithm, we write $y \leftarrow_s \mathbf{A}(x)$ to denote a run of \mathbf{A} on input x (and implicit random coins r) and output y ; the value y is a random variable, and $\mathbf{A}(x; r)$ denotes a run of \mathbf{A} on input x and randomness r . An algorithm \mathbf{A} is *probabilistic polynomial-time* (PPT) if \mathbf{A} is randomized and for any input $x, r \in \{0, 1\}^*$ the computation of $\mathbf{A}(x; r)$ 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 a polynomial, denoted $p(\lambda) \in \text{poly}(\lambda)$, 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 (or simply negligible) 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{X}' defined over the same set \mathcal{X} is defined as $\mathbb{SD}(\mathbf{X}; \mathbf{X}') = \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathbb{P}[\mathbf{X} = x] - \mathbb{P}[\mathbf{X}' = x]|$.

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, $\mathbf{X} \approx_s \mathbf{Y}$ to denote that they are statistically close, i.e. $\mathbb{SD}(\mathbf{X}_\lambda; \mathbf{Y}_\lambda) \in \text{negl}(\lambda)$, and $\mathbf{X} \approx_c \mathbf{Y}$ to denote that they are computationally indistinguishable, i.e., for all PPT distinguishers \mathbf{D} :

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

We extend the notion of computational indistinguishability to the case of interactive experiments (a.k.a. games) featuring an adversary \mathbf{A} . In particular, let $\mathbf{G}_\mathbf{A}(\lambda)$ be the random

variable corresponding to the output of \mathbf{A} at the end of the experiment, where wlog. we may assume \mathbf{A} outputs a decision bit. Given two experiments $\mathbf{G}_A(\lambda, 0)$ and $\mathbf{G}_A(\lambda, 1)$, we write $\{\mathbf{G}_A(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{G}_A(\lambda, 1)\}_{\lambda \in \mathbb{N}}$ 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 (in case of unbounded adversaries). We recall a useful lemma from [DP07, DDV10].

Lemma 1 ([DDV10], Lemma 4). *Let $\mathcal{O}_{\text{leak}}(x, g)$ be an oracle that upon input a value x and a function g outputs $g(x)$, and let \mathbf{X} and \mathbf{Y} be two independently distributed random variables. For any adversary \mathbf{A} , and for any value z , the distributions $(\mathbf{X}|z = \mathbf{A}^{\mathcal{O}_{\text{leak}}(\mathbf{X}, \cdot), \mathcal{O}_{\text{leak}}(\mathbf{Y}, \cdot)})$ and $(\mathbf{Y}|z = \mathbf{A}^{\mathcal{O}_{\text{leak}}(\mathbf{X}, \cdot), \mathcal{O}_{\text{leak}}(\mathbf{Y}, \cdot)})$ are independently distributed.*

Average min-entropy. The min-entropy of a random variable \mathbf{X} with domain \mathcal{X} is $\mathbb{H}_\infty(\mathbf{X}) := -\log \max_{x \in \mathcal{X}} \mathbb{P}[\mathbf{X} = x]$, and intuitively it measures the best chance to predict \mathbf{X} (by a computationally unbounded algorithm). For conditional distributions, unpredictability is measured by the conditional average min-entropy [DORS08]: $\tilde{\mathbb{H}}_\infty(\mathbf{X}|\mathbf{Y}) := -\log \mathbb{E}_y[2^{-\mathbb{H}_\infty(\mathbf{X}|\mathbf{Y}=y)}]$. The lemma below is sometimes known as the “chain rule” for conditional average min-entropy.

Lemma 2 ([DORS08], Lemma 2.2). *Let $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ be random variables. If \mathbf{Y} has at most 2^ℓ possible values, then $\tilde{\mathbb{H}}_\infty(\mathbf{X}|\mathbf{Y}, \mathbf{Z}) \geq \tilde{\mathbb{H}}_\infty(\mathbf{X}, \mathbf{Y}|\mathbf{Z}) - \ell \geq \tilde{\mathbb{H}}_\infty(\mathbf{X}|\mathbf{Z}) - \ell$. In particular, $\tilde{\mathbb{H}}_\infty(\mathbf{X}|\mathbf{Y}) \geq \tilde{\mathbb{H}}_\infty(\mathbf{X}, \mathbf{Y}) - \ell \geq \tilde{\mathbb{H}}_\infty(\mathbf{X}) - \ell$.*

2.1 Secret Sharing Schemes

An n -party secret sharing scheme Σ in the common reference string (CRS) model consists of polynomial-time algorithms (Init, Share, Rec) specified as follows: (i) The randomized initialization algorithm **Init** takes as input the security parameter 1^λ , and outputs a CRS $\omega \in \{0, 1\}^*$; (ii) The randomized sharing algorithm **Share** takes as input a CRS $\omega \in \{0, 1\}^*$ and a message $m \in \mathcal{M}$, and outputs n shares s_1, \dots, s_n where each $s_i \in \mathcal{S}_i$; (iii) The deterministic algorithm **Rec** takes as input a CRS $\omega \in \{0, 1\}^*$ and a certain number of candidate shares, and outputs a value in $\mathcal{M} \cup \{\perp\}$. Given $s = (s_1, \dots, s_n)$ and a subset $\mathcal{I} \subseteq [n]$, we often write $s_{\mathcal{I}}$ to denote the shares $(s_i)_{i \in \mathcal{I}}$.

The subset of parties allowed to reconstruct the secrets by pulling their shares together form the so-called access structure.

Definition 1 (Access structure). We say \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 sets $\mathcal{I} \in \mathcal{A}$ authorized or qualified, and unauthorized or unqualified otherwise.

Intuitively, a secure secret sharing scheme must be such that all qualified subsets of players can efficiently reconstruct the secret, whereas all unqualified subset 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 $\Sigma = (\text{Init}, \text{Share}, \text{Rec})$ is a secret sharing scheme realizing access structure \mathcal{A} in the CRS model, 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 in the CRS model with the following properties.

- (i) **Correctness:** For all $\lambda \in \mathbb{N}$, all $\omega \in \text{Init}(1^\lambda)$, all messages $m \in \mathcal{M}$, and for all subsets $\mathcal{I} \in \mathcal{A}$, we have that $\text{Rec}(\omega, (\text{Share}(\omega, m))_{\mathcal{I}}) = m$, with overwhelming probability over the randomness of the sharing algorithm.

(ii) **Privacy:** For all PPT adversaries $A = (A_1, A_2)$, we have

$$\{\mathbf{Privacy}_{\Sigma, A}(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Privacy}_{\Sigma, A}(\lambda, 1)\}_{\lambda \in \mathbb{N}},$$

where the experiment $\mathbf{Privacy}_{\Sigma, A}(\lambda, b)$ is defined by

$$\mathbf{Privacy}_{\Sigma, A}(\lambda, b) := \left\{ \begin{array}{l} \omega \leftarrow \text{Init}(1^\lambda); (m_0, m_1, \mathcal{U} \notin \mathcal{A}, \alpha_1) \leftarrow A_1(\omega) \\ s \leftarrow \text{Share}(\omega, m_b); b' \leftarrow A_2(\alpha_1, s_{\mathcal{U}}) \end{array} \right\}.$$

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

Moreover, we say that Σ is a secret sharing scheme realizing access structure \mathcal{A} in the plain model, if for all $\lambda \in \mathbb{N}$ algorithm Init simply returns $\omega = 1^\lambda$.

Remark 1. In the plain model, the above definition of privacy is equivalent to saying that for all pairs of messages $m_0, m_1 \in \mathcal{M}$, and for all unqualified subsets $\mathcal{U} \notin \mathcal{A}$, it holds that

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

2.2 Non-Interactive Commitments

A non-interactive commitment scheme $\Pi = (\text{Gen}, \text{Com})$ is a pair of polynomial-time algorithms specified as follows: (i) The randomized algorithm Gen takes as input 1^λ and outputs a public key $pk \in \mathcal{K}$; (ii) The randomized algorithm Com takes as input the public key pk and a message $m \in \mathcal{M}$, and outputs a commitment $c = \text{Com}(pk, m; r) \in \mathcal{C}$ using random coins $r \in \mathcal{R}$. The pair (m, r) is called the opening. In the plain model, we omit the algorithm Gen and simply set $pk = 1^\lambda$.

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 definitions follow.

Definition 3 (Binding). We say that a non-interactive commitment scheme $\Pi = (\text{Gen}, \text{Com})$ is *computationally binding* if the following probability is negligible for all PPT adversaries A :

$$\mathbb{P} \left[m_0 \neq m_1 \wedge \text{Com}(pk, m_0; r_0) = \text{Com}(pk, m_1; r_1) : \begin{array}{l} pk \leftarrow \text{Gen}(1^\lambda) \\ (m_0, r_0, m_1, r_1) \leftarrow A(pk) \end{array} \right].$$

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

Definition 4 (Hiding). We say that a non-interactive commitment scheme $\Pi = (\text{Gen}, \text{Com})$ is *computationally hiding* if the following holds for all PPT adversaries A :

$$\left\{ \begin{array}{l} pk \leftarrow \text{Gen}(1^\lambda); (m_0, m_1, \alpha_1) \leftarrow A_1(pk) \\ c \leftarrow \text{Com}(pk, m_0); b' \leftarrow A_2(\alpha_1, c) \end{array} \right\} \\ \approx_c \left\{ \begin{array}{l} pk \leftarrow \text{Gen}(1^\lambda); (m_0, m_1, \alpha_1) \leftarrow A_1(pk) \\ c \leftarrow \text{Com}(pk, m_1); b' \leftarrow A_2(\alpha_1, c) \end{array} \right\}.$$

In case the above ensembles are statistically close (resp. identically distributed), we speak of *statistical* (resp. *perfect*) hiding.

Note that in the plain model the above definition of hiding is equivalent to saying that for all pairs of messages $m_0, m_1 \in \mathcal{M}$ the following holds:

$$\left\{ c : c \leftarrow \text{Com}(1^\lambda, m_0) \right\}_{\lambda \in \mathbb{N}} \approx_c \left\{ c : c \leftarrow \text{Com}(1^\lambda, m_1) \right\}_{\lambda \in \mathbb{N}}.$$

2.3 Non-Interactive Zero Knowledge

Let R be a relation, corresponding to an NP language \mathcal{L} . A non-interactive zero-knowledge (NIZK) proof system for R is a tuple of efficient algorithms $\Pi = (\text{CRSGen}, \text{Prove}, \text{Ver})$ specified as follows. (i) The randomized algorithm CRSGen takes as input the security parameter and outputs a common reference string ω ; (ii) The randomized algorithm $\text{Prove}(\omega, \phi, (x, w))$, given $(x, w) \in R$ and a label $\phi \in \{0, 1\}^*$, outputs a proof π ; (iii) The deterministic algorithm $\text{Ver}(\omega, \phi, (x, \pi))$, given an instance x , a proof π , and a label $\phi \in \{0, 1\}^*$, outputs either 0 (for “reject”) or 1 (for “accept”). We say that a NIZK for relation R is *correct* if for every $\lambda \in \mathbb{N}$, all ω as output by $\text{Init}(1^\lambda)$, any label $\phi \in \{0, 1\}^*$, and any $(x, w) \in R$, we have that $\text{Ver}(\omega, \phi, (x, \text{Prove}(\omega, \phi, (x, w)))) = 1$.

We define two properties of a NIZK proof system. The first property says that honest proofs do not reveal anything beyond the fact that $x \in \mathcal{L}$.

Definition 5 (Adaptive multi-theorem zero-knowledge). A NIZK with labels Π for a relation R satisfies adaptive multi-theorem zero-knowledge if there exists a PPT simulator $S := (S_0, S_1)$ such that the following holds:

- (i) S_0 outputs ω , a simulation trapdoor ζ and an extraction trapdoor ξ .
- (ii) For all PPT distinguishers D , we have that

$$\left| \mathbb{P} \left[D^{\text{Prove}(\omega, \cdot, (\cdot, \cdot))}(\omega) = 1 : \omega \leftarrow S_0(1^\lambda) \right] - \mathbb{P} \left[D^{\mathcal{O}_{\text{sim}}(\zeta, \cdot, \cdot, \cdot)}(\omega) = 1 : (\omega, \zeta) \leftarrow S_0(1^\lambda) \right] \right|$$

is negligible in λ , where the oracle $\mathcal{O}_{\text{sim}}(\zeta, \cdot, \cdot, \cdot)$ takes as input a tuple (ϕ, x, w) and returns $S_1(\zeta, \phi, x)$ iff $R(x, w) = 1$ (and otherwise it returns \perp).

Groth [Gro06] introduced the concept of simulation-extractable NIZK, which informally states that knowledge soundness should hold even if the adversary can see simulated proofs for possibly false statements of its choice. For our purpose, it will suffice to consider the weaker notion of true simulation extractability, as defined by Dodis *et al.* [DHLW10].

Definition 6 (True simulation extractability). Let Π be a NIZK proof systems for a relation R , that satisfies adaptive multi-theorem zero-knowledge w.r.t. a simulator $S := (S_0, S_1)$. We say that Π is *true simulation extractable* if there exists a PPT algorithm K such that every PPT adversary A has a negligible probability of winning in the following game:

- The challenger runs $(\omega, \zeta, \xi) \leftarrow S_0(1^\lambda)$, and gives ω to A .
- Adversary A can ask polynomially many queries of the form (ϕ, x, w) , upon which the challenger returns $S_1(\zeta, \phi, x)$ if $(x, w) \in R$ and \perp otherwise.
- Adversary A outputs a tuple (ϕ^*, x^*, π^*) .
- The challenger runs $w \leftarrow S_1(\xi, \phi^*, (x^*, \pi^*))$.

We say that A wins iff: (a) (ϕ^*, x^*) was not queried in the second step; (b) $\text{Ver}(\omega, \phi^*, (x^*, \pi^*)) = 1$; (c) $(x^*, w) \notin R$.

2.4 Randomness Extractors

An (average-case, strong) seeded extractor is a polynomial-time deterministic algorithm $\text{Ext} : \{0, 1\}^d \times \{0, 1\}^\rho \rightarrow \{0, 1\}^\alpha$ taking as input a source x and a seed r , and outputting an α -bit string.

Definition 7 (Seeded extractors). Let $\kappa \in \mathbb{N}$ and $\epsilon \in [0, 1]$. We say that Ext is (κ, ϵ) -secure, if for any random variable $\mathbf{\Lambda}$ such that $\tilde{\mathbb{H}}_\infty(\mathbf{X}|\mathbf{\Lambda}) \geq \kappa$ the statistical distance between the distribution $(\mathbf{R}, \text{Ext}(\mathbf{X}, \mathbf{R}), \mathbf{\Lambda})$ and $(\mathbf{R}, \mathbf{A}, \mathbf{\Lambda})$ is at most ϵ , where $\mathbf{R} \equiv \mathbf{U}_\rho$ and $\mathbf{A} \equiv \mathbf{U}_\alpha$.

An (average-case) two-source extractor is a deterministic polynomial-time algorithm $2\text{Ext} : \{0, 1\}^d \times \{0, 1\}^d \rightarrow \{0, 1\}^\alpha$ taking as input two sources x and y and outputting an α -bit string.

Definition 8 (Two-source extractors). Let $\kappa_1, \kappa_2 \in \mathbb{N}$ and $\epsilon \in [0, 1]$. We say that 2Ext is $(\kappa_1, \kappa_2, \epsilon)$ -secure if for any random variable $\mathbf{\Lambda}$ such that $\tilde{\mathbb{H}}_\infty(\mathbf{X}|\mathbf{\Lambda}) \geq \kappa_1$, $\tilde{\mathbb{H}}_\infty(\mathbf{Y}|\mathbf{\Lambda}) \geq \kappa_2$, and $(\mathbf{X}|\mathbf{\Lambda}), (\mathbf{Y}|\mathbf{\Lambda})$ are independent, the statistical distance between the distribution $(\mathbf{X}, 2\text{Ext}(\mathbf{X}, \mathbf{Y}), \mathbf{\Lambda})$ and $(\mathbf{X}, \mathbf{A}, \mathbf{\Lambda})$ is at most ϵ , where $\mathbf{A} \equiv \mathbf{U}_\alpha$.

3 Continuous Tampering under Selective Partitioning

In this section we define a new notion of non-malleability against joint memory tampering and leakage for secret sharing. Our definition generalizes the one in [FV19] which was tailored to threshold access structures and to independent leakage/tampering from the shares.

Very roughly, in our model the attacker is allowed to partition the set of share holders into β (non-overlapping) blocks with size at most k , covering the entire set $[n]$. This is formalized through the notion of a k -partition.

Definition 9 (k -partition). Let $n, k, \beta \in \mathbb{N}$. We call $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_\beta)$ a k -partition of $[n]$ when:

- (i) $\bigcup_{i=1}^{\beta} \mathcal{B}_i = [n]$;
- (ii) $\forall i_1, i_2 \in [\beta]$, with $i_1 \neq i_2$, we have $\mathcal{B}_{i_1} \cap \mathcal{B}_{i_2} = \emptyset$.
- (iii) $\forall i \in [\beta] : |\mathcal{B}_i| \leq k$.

3.1 The Definition

To define non-malleability, we consider an attacker \mathbf{A} playing the following game. At the beginning of the experiment, \mathbf{A} chooses two messages m_0, m_1 possibly depending on the CRS ω of the underlying secret sharing scheme, and a k -partition $(\mathcal{B}_1, \dots, \mathcal{B}_\beta)$ of the set $[n]$. Hence, the adversary interacts with a target secret sharing $s = (s_1, \dots, s_n)$ of either m_0 or m_1 , via the following queries:

- **Leakage queries.** For each $j \in [\beta]$, the attacker can leak jointly from the shares $s_{\mathcal{B}_j}$. This can be done repeatedly and in an adaptive fashion, the only limitation being that the overall amount of leakage on each block is at most $\ell \in \mathbb{N}$ bits.
- **Tampering queries.** For each $j \in [\beta]$, the attacker can tamper jointly the shares $s_{\mathcal{B}_j}$. 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. an arbitrary reconstruction set $\mathcal{T} \in \mathcal{A}$ that is also chosen adversarially. This can be done for at most $p \in \mathbb{N}$ times, and in an adaptive fashion.

The above naturally yields a notion of joint bounded-leakage and tampering admissible adversary, as defined below. Note that, in order to rule out trivial attacks, we must require that the partition \mathcal{B} chosen by the attacker be such that no block of the partition is an authorized set for the underlying access structure.

Definition 10 (Joint bounded-leakage and tampering admissible adversaries). Let $n, k, \ell, p \in \mathbb{N}$, and fix an arbitrary message space \mathcal{M} , sharing domain $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ and access structure

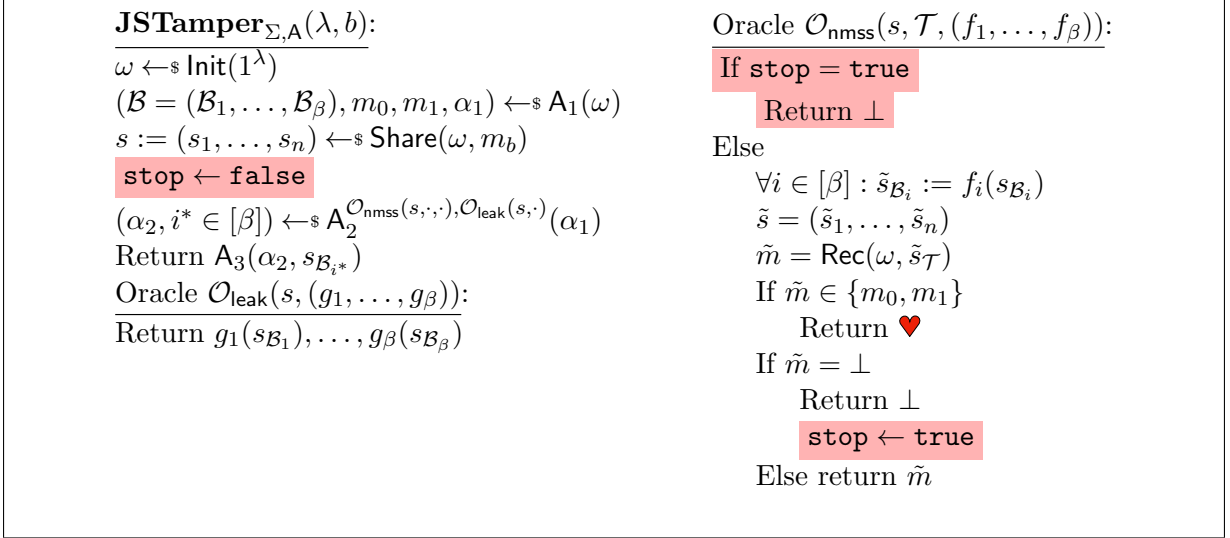


Figure 1: Experiment defining leakage-resilient (continuously) non-malleable secret sharing under adaptive concurrent reconstruction. The instructions boxed in red are considered only for continuous non-malleability, in which case the oracle $\mathcal{O}_{\text{nmss}}$ is implicitly parameterized by the flag **stop**.

\mathcal{A} for n parties. We say that a (possibly unbounded) adversary $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ is k -joint ℓ -bounded-leakage and p -tampering admissible ((k, ℓ, p) -BLTA for short) if it satisfies the following conditions:

- (i) \mathbf{A}_1 outputs two messages $m_0, m_1 \in \mathcal{M}$ and a k -partition $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_\beta)$ of $[n]$ such that $\forall j \in [\beta]$ we have $\mathcal{B}_j \notin \mathcal{A}$.
- (ii) \mathbf{A}_2 outputs a sequence of poly-many leakage queries, chosen adaptively, $(g_1^{(q)}, \dots, g_\beta^{(q)})_{q \in \text{poly}(\lambda)}$ such that $\forall j \in [\beta]$ it holds that $\sum_q |g_j^{(q)}(\cdot)| \leq \ell$, where $g_j^{(q)} : \times_{i \in \mathcal{B}_j} \mathcal{S}_i \rightarrow \{0, 1\}^*$.
- (iii) \mathbf{A}_2 outputs a sequence of p tampering queries, chosen adaptively, $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_\beta^{(q)}))_{q \in [p]}$ such that $\mathcal{T}^{(q)} \in \mathcal{A}$, and $\forall j \in [\beta]$ it holds that $f_j^{(q)} : \times_{i \in \mathcal{B}_j} \mathcal{S}_i \rightarrow \times_{i \in \mathcal{B}_j} \mathcal{S}_i$.

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 . In the definition below, the attacker is further allowed to obtain in full the shares belonging to one of the partitions, at the end of the experiment. This is reminiscent of augmented (leakage-resilient) non-malleability, as considered in [FMNV14, AAG⁺16, GPR16, Cfv19].

Definition 11 (Leakage-resilient non-malleability under selective partitioning). Let $n, k, \ell, p \in \mathbb{N}$ be parameters, and \mathcal{A} be an access structure for n parties. We say that $\Sigma = (\text{Init}, \text{Share}, \text{Rec})$ is an *augmented* ℓ -bounded leakage-resilient p -time non-malleable secret sharing scheme realizing \mathcal{A} against selective k -joint leakage and tampering in the CRS model (resp., in the plain model)—augmented (k, ℓ, p) -BLR-CNMSS for short—if it is an n -party secret sharing scheme realizing \mathcal{A} in the CRS model (resp., in the plain model) as per Def. 2, and additionally for all (k, ℓ, p) -BLTA adversaries $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ we have:

$$\{\mathbf{JSTamper}_{\Sigma, \mathcal{A}}(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{JSTamper}_{\Sigma, \mathcal{A}}(\lambda, 1)\}_{\lambda \in \mathbb{N}},$$

where, for $b \in \{0, 1\}$, experiment $\mathbf{JSTamper}_{\Sigma, \mathcal{A}}(\lambda, b)$ is depicted in Fig. 1.

In case the above definition holds for all $p(\lambda) \in \text{poly}(\lambda)$, but w.r.t. all PPT adversaries A (i.e., \approx_s is replaced with \approx_c in the above equation), we call Σ (augmented, bounded leakage-resilient) continuously non-malleable. As shown by [FV19], already for the simpler case of independent tampering, it is impossible to achieve this notion without assuming self-destruct (i.e., the oracle $\mathcal{O}_{\text{nmss}}$ must stop answering tampering queries after the first such query yielding an invalid reconstructed message).

It is also well-known that computational security is inherent for obtaining continuously non-malleable secret sharing realizing threshold access structures [BS19]. Unless stated otherwise, when we refer to non-malleable secret sharing in this paper we implicitly assume security holds in the computational setting (both for privacy and non-malleability).

3.2 On Augmented Non-Malleability

When dropping the adversary A_3 from the above definition, we obtain the standard (non-augmented) notion of (leakage-resilient, continuous) non-malleability. The theorem below, however, says that augmented security is essentially for free whenever non-malleability is considered together with leakage resilience. Intuitively, this is because in the reduction we can simply simulate all leakage queries, and then ask a final leakage query which reveals the output guess of an hypothetical distinguisher attacking augmented non-malleability.⁶ A similar proof strategy was used in [KMS18, Lemma 7].

Theorem 3. *Let Σ be a $(k, \ell + 1, p)$ -BLR-CNMS realizing access structure \mathcal{A} for n parties in the CRS model (resp. plain model). Then, Σ is an augmented (k, ℓ, p) -BLR-CNMS realizing \mathcal{A} in the CRS model (resp. plain model).*

Proof. Without loss of generality, we prove the statement in the CRS model. Let $A^+ = (A_1^+, A_2^+, A_3^+)$ be a (β, ℓ, p) -BLTA attacker violating Def. 11 for Σ ; we construct an adversary $A = (A_1, A_2)$ breaking the non-augmented variant of Def. 11 for Σ . Attacker A works as follows:

- Upon receiving ω from the challenger, $A_1(\omega)$ outputs the same tuple (\mathcal{B}, m_0, m_1) as returned by $A_1^+(\omega)$.
- Upon input a leakage query (g_1, \dots, g_β) from A_2^+ , forward the same query to the target leakage oracle and return the answer to A_2^+ .
- Upon input a tampering query $(\mathcal{T}, (f_1, \dots, f_\beta))$ from A_2^+ , forward the same query to the target tampering oracle and return the answer to A_2^+ .
- Let (α_2, i^*) be the final output of A_2^+ . Define the leakage function $\hat{g}_{i^*}^{\alpha_2, A_3^+}$ which hard-wires α_2 and a description of A_3^+ , takes as input the shares $s_{\mathcal{B}_{i^*}}$, and returns the decision bit $b' \leftarrow A_3^+(\alpha_2, s_{\mathcal{B}_{i^*}})$.
- Forward $(\varepsilon, \dots, \hat{g}_{i^*}^{\alpha_2, A_3^+}, \varepsilon, \dots, \varepsilon)$ to the target leakage oracle, obtaining a bit b' which is then sent to the challenger.

The statement follows by observing that A 's simulation to A^+ 's leakage/tampering queries is perfect, and moreover A_2 leaks a total of at most $\ell + 1$ bits from each partition of the shares. \square

3.3 Related Notions

We finally argue that known definitions from the literature can be cast by either restricting, or slightly tweaking, Def. 11.

⁶While we state the theorem for the case of bounded leakage, an identical statement holds in the noisy-leakage setting.

Independent leakage and tampering. The definition below restricts the adversary to leak/tamper from/with each of the shares individually; this is sometimes known as local or independent leakage/tampering. The condition on leakage admissibility, though, is more general, in that the attacker can leak an arbitrary amount of information as long as the total leakage reduces the uncertainty on each share (conditioned on the other shares) by at most ℓ bits.

Definition 12 (Independent noisy-leakage and tampering admissible adversaries). Let $n, \ell, p \in \mathbb{N}$, and fix an arbitrary message space \mathcal{M} , sharing domain $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ and access structure \mathcal{A} for n parties. We say that a (possibly unbounded) adversary $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ is independent ℓ -noisy-leakage and tampering admissible ((n, ℓ, p) -NLTA for short) if it satisfies the following conditions:

- (i) \mathbf{A}_1 outputs two messages $m_0, m_1 \in \mathcal{M}$ and the partition $\mathcal{B} = (\{1\}, \dots, \{n\})$.
- (ii) \mathbf{A}_2 outputs a sequence of poly-many leakage queries (chosen adaptively) $(g_1^{(q)}, \dots, g_n^{(q)})_{q \in \text{poly}(\lambda)}$ such that $\forall i \in [n]$ we have $g_i^{(q)} : \mathcal{S}_i \rightarrow \{0, 1\}^*$, and $\forall m \in \mathcal{M}$ it holds that:

$$\tilde{\mathbb{H}}_\infty \left(\mathbf{S}_i | (\mathbf{S}_j)_{j \neq i}, g_i^{(1)}(\mathbf{S}_i), \dots, g_i^{(p)}(\mathbf{S}_i) \right) \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_i | (\mathbf{S}_j)_{j \neq i}) - \ell,$$

where $(\mathbf{S}_1, \dots, \mathbf{S}_n)$ is the random variable corresponding to $\text{Share}(\text{Init}(1^\lambda), m)$.

- (iii) \mathbf{A}_2 outputs a sequence of tampering queries (chosen adaptively) $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_n^{(q)}))_{q \in [p]}$ such that $\mathcal{T}^{(q)} \in \mathcal{A}$, and $\forall i \in [n]$ it holds that $f_i^{(q)} : \mathcal{S}_i \rightarrow \mathcal{S}_i$.

When restricting Def. 11 to all PPT $(n, \ell, \text{poly}(\lambda))$ -NLTA adversaries, we obtain the notion of (augmented) ℓ -noisy leakage-resilient continuously non-malleable secret sharing against individual leakage and tampering (with adaptive concurrent reconstructions) [FV19]. Finally, if we consider $n = 2$ and the threshold access structure with reconstruction parameter $\varrho = 2$ (i.e., both shares are required in order to reconstruct the message), we immediately obtain noisy leakage-resilient continuously non-malleable codes in the split-state model [FMNV14, OPVV18]. In what follows, we write $\mathbf{Tamper}(\lambda, b)$ to denote the random variable in the security experiment of Def. 11 with an (n, ℓ, p) -NLTA adversary.

Leakage-resilient secret sharing. Further, when no tampering is allowed (i.e., $p = 0$), we obtain the notion of leakage-resilient secret sharing [DDV10, KMS18, SV19, ADN⁺19, NS19] as a special case. In particular, we write $\mathbf{JSLeak}(\lambda, b)$ to denote the random variable in the security experiment of Def. 11 with a $(k, \ell, 0)$ -BLTA adversary, and $\mathbf{Leak}(\lambda, b)$ to denote the random variable in the security experiment of Def. 11 with an $(n, \ell, 0)$ -NLTA adversary.

Recall that, by Theorem 3, the augmented variant is without loss of generality as long as leakage resilience holds for $\ell \geq 2$.

4 Construction in the CRS Model

4.1 Description of the Scheme

We show how to obtain leakage-resilient continuously non-malleable secret sharing for arbitrary access structures in the CRS model, with security against selective joint leakage and tampering. Our construction combines a commitment scheme (Gen, Com) (cf. §2.2), a non-interactive proof system $(\text{CRSGen}, \text{Prove}, \text{Ver})$ for proving knowledge of a committed value (cf. §2.3), and an auxiliary n -party secret sharing scheme $\Sigma = (\text{Share}, \text{Rec})$, as depicted in Fig. 2.

The main idea behind the scheme is as follows. The CRS includes the CRS ω for the proof system and the public key pk for the commitment scheme. Given a message $m \in \mathcal{M}$, the sharing

Let $\Sigma = (\text{Share}, \text{Rec})$ be an auxiliary secret sharing scheme realizing access structure \mathcal{A} , with message space \mathcal{M} and share space $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$. Let (Gen, Com) be a commitment scheme with domain $\{0, 1\}^*$, and $(\text{CRSGen}, \text{Prove}, \text{Ver})$ be a non-interactive argument system for the language $\mathcal{L}_{\text{com}}^{pk} = \{c \in \{0, 1\}^\gamma : \exists i \in [n], s \in \mathcal{S}_i, r \in \mathcal{R} \text{ s.t. } \text{Com}(pk, i || s; r) = c\}$ that supports labels in $\{0, 1\}^\gamma$. Define the following secret sharing scheme $\Sigma^* = (\text{Init}^*, \text{Share}^*, \text{Rec}^*)$ in the CRS model.

Initialization algorithm Init^* : Sample $\omega \leftarrow_{\$} \text{CRSGen}(1^\lambda)$ and $pk \leftarrow_{\$} \text{Gen}(1^\lambda)$, and return $\omega^* = (\omega, pk)$.

Sharing algorithm Share^* : Upon input $\omega^* = (\omega, pk)$ and a value $m \in \mathcal{M}$, compute $(s_1, \dots, s_n) \leftarrow_{\$} \text{Share}(m)$. For each $i \in [n]$, generate $r_i \leftarrow_{\$} \mathcal{R}$ and define $c_i = \text{Com}(pk, i || s_i; r_i)$. For each $i, j \in [n]$ such that $i \neq j$, define $\pi_i^j \leftarrow_{\$} \text{Prove}(\omega, c_j, (c_i, i || s_i, r_i))$. Return the shares $s^* = (s_1^*, \dots, s_n^*)$, where for each $i \in [n]$ we set $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.

Reconstruction algorithm Rec^* : Upon input $\omega^* = (\omega, pk)$ and shares $(s_i^*)_{i \in \mathcal{I}}$ parse $s_i^* = (s_i, r_i, (c_j^i)_{j \neq i}, (\pi_j^i)_{j \neq i})$ for each $i \in \mathcal{I}$. Hence, proceed as follows:

- (a) If $\exists i_1, i_2 \in \mathcal{I}$ and $j \in [n]$ such that $c_j^{i_1} \neq c_j^{i_2}$, output \perp ; else let the input shares be $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$ for each $i \in \mathcal{I}$.
- (b) If $\exists i \in \mathcal{I}$ such that $\text{Com}(pk, i || s_i; r_i) \neq c_i$, output \perp .
- (c) If $\exists i, j \in \mathcal{I}$ such that $i \neq j$ and $\text{Ver}(\omega, c_j, (c_i, \pi_i^j)) = 0$, output \perp .
- (d) Else, output $\text{Rec}((s_i)_{i \in \mathcal{I}})$.

Figure 2: Leakage-resilient continuously non-malleable secret sharing for arbitrary access structures against selective joint leakage and tampering, in the CRS model.

procedure first shares m using Share , obtaining shares (s_1, \dots, s_n) . Then, it commits to the i -th share s_i along with the position i using randomness r_i , and finally generates $n - 1$ proofs $(\pi_j^i)_{j \neq i}$ for the statement c_i using each time the value $c_j = \text{Com}(pk, j || s_j; r_j)$ as label. The final share of player i consists of s_i , along with the randomness r_i used to obtain c_i and all the values $(c_j)_{j \neq i}$ and $(\pi_j^i)_{j \neq i}$. The reconstruction procedure, given a set of shares $s_{\mathcal{I}}^*$, first checks that for each $i \in \mathcal{I}$ the commits $(c_j)_{j \neq i}$ contained in each share are all equal, and moreover each c_i is indeed obtained by committing $i || s_i$ with the randomness r_i ; further, it checks that all the proofs verify correctly w.r.t. the corresponding statement and label. If any of the above checks fails, the algorithm returns \perp and otherwise it outputs the same as $\text{Rec}(s_{\mathcal{I}})$.

Intuitively, our scheme can be seen as a generalization of the original construction of continuously non-malleable codes in the split-state model from [FMNV14]. In particular, when $n = 2$, the two constructions are identical except for two differences: (i) We commit to each share, whereas [FMNV14] uses a collision-resistant hash function; (ii) We include the position of each share in the commitment. Roughly speaking, the first modification is necessary in order to prove privacy (as hash functions do not necessarily hide their inputs). The second modification is needed in order to avoid that an attacker can permute the shares within one of the partitions, which was not possible in the setting of independent tampering. We establish the following result.

Theorem 4. *Let $n, k \in \mathbb{N}$, and \mathcal{A} be any access structure for n parties. Assume that:*

- (i) Σ is an n -party augmented ℓ -bounded leakage-resilient secret sharing scheme realizing access structure \mathcal{A} against selective k -joint leakage in the plain model;
- (ii) (Gen, Com) is a statistically hiding and computationally binding commitment scheme with commitment length $\gamma = O(\lambda)$;

(iii) $(\text{CRSGen}, \text{Prove}, \text{Ver})$ is a true-simulation extractable non-interactive zero-knowledge argument system for the language $\mathcal{L}_{\text{com}}^{\text{pk}} = \{c \in \{0, 1\}^\gamma : \exists i \in [n], s \in \mathcal{S}_i, r \in \mathcal{R} \text{ s.t. } \text{Com}(\text{pk}, i || s; r) = c\}$.

Then, the secret sharing scheme Σ^* described in Fig. 2 is an n -party augmented ℓ^* -bounded leakage-resilient continuously non-malleable secret sharing scheme realizing access structure \mathcal{A} against selective k -joint leakage and tampering in the CRS model, as long as $\ell = 2\ell^* + n\gamma + O(\lambda \log \lambda)$.

4.2 Proof Overview

Before coming to the proof, we discuss the main intuition behind privacy and continuous non-malleability.

Privacy. In order to show privacy, we need to prove that no PPT attacker \mathbf{A} can distinguish between $\mathbf{Privacy}_{\Sigma^*, \mathbf{A}}(1^\lambda, 0)$ and $\mathbf{Privacy}_{\Sigma^*, \mathbf{A}}(1^\lambda, 1)$ (cf. Def. 2). Recall that in this experiment, after seeing the CRS, the attacker can select an unauthorized subset \mathcal{U} along with messages (m_0, m_1) , after which it is given the shares $s_{\mathcal{U}}^*$ (either corresponding to message m_0 or to m_1).

Consider a game hop in which we generate the CRS together with a simulation trapdoor, and we replace the proofs π_i^j in the target secret sharing s^* with simulated proofs. A reduction to adaptive multi-theorem zero knowledge of the NIZK (cf. Def. 5 in §2.3) shows that no attacker can distinguish this modified experiment from the original game. Next, we replace the values c_i with commitments to dummy values. A simple hybrid argument, relying on the statistical hiding property of the commitment (cf. Def. 4 in §2.2), shows that this experiment is statistically close to the previous one.

Finally, we reduce to the privacy of Σ . In particular, note that the reduction can locally sample the CRS and compute the values c_i (with exactly the same distribution as in the final experiment).

Non-malleability. The goal is to show that no PPT attacker \mathbf{A} can distinguish the experiments $\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, 0)$ and $\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, 1)$ (cf. Def. 11). Recall that \mathbf{A} , after seeing the CRS, can selectively partition the set $[n]$ into β blocks $\mathcal{B}_1, \dots, \mathcal{B}_\beta$, and then jointly leak from and tamper with the shares within each block. The proof proceeds via a hybrid argument, as outlined below.

Hyb $_{\Sigma^*, \mathbf{A}}^1(\lambda, b)$: In the first hybrid, we modify the distribution of the target secret sharing $s^* = (s_1, \dots, s_n)$. In particular, we first let the NIZK simulator \mathbf{S}_0 program the CRS ω yielding simulation trapdoor ζ and extraction trapdoor ξ , and then we compute each of the proofs π_i^j by running the NIZK simulator \mathbf{S}_1 upon input ζ , statement c_i and label c_j .

Hyb $_{\Sigma^*, \mathbf{A}}^2(\lambda, b)$: In the second hybrid, we modify the way tampering queries are answered. In particular, let $(\mathcal{T}, (f_1, \dots, f_\beta))$ be a generic tampering query, and $\tilde{s}^* = (\tilde{s}_1^*, \dots, \tilde{s}_n^*)$ be the mauled secret sharing after tampering jointly with the shares; here, each \tilde{s}_i^* can be parsed⁷ as $\tilde{s}_i^* = (\tilde{s}_i, (\tilde{c}_j)_{j \neq i}, (\tilde{\pi}_j^i)_{j \neq i})$.

The tampering oracle proceeds as follows: (a) In case the set of mauled commitments $(\tilde{c})_{i \in [n]}$ is equal to the set of initial commitments $(c_i)_{i \in [n]}$ considered in their order, the answer is \heartsuit ; (b) In case there exist at least two distinct indices i_1, i_2 such that $\tilde{c}_{i_1} = c_{i_2}$, abort; (c) Otherwise, there exists at least one commitment \tilde{c}_{i^*} that is different from all

⁷As in the original experiment, the tampering oracle first checks that the commitments in each of the shares are all equal, that the NIZK proofs are accepting, and that each \tilde{c}_j is indeed a valid commitment of $j || \tilde{s}_j$; in case any of these checks fails, the oracle self-destructs.

the initial commitments $(c_i)_{i \in [n]}$. In the latter case, we can use the extraction trapdoor and the knowledge extractor of the NIZK to extract⁸ from each set $\tilde{s}_{\mathcal{B}_i}^*$ of tampered shares the missing shares $\tilde{s}_{\mathcal{T} \setminus \mathcal{B}_i}^*$, which yields a candidate answer for the attacker's tampering query: In case the answers are all consistent for each of the β blocks in the partition, the oracle returns this value, and otherwise it outputs \perp and self-destructs.

As a first step, we argue that $\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, b)$ and $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)$ are computationally close. This follows readily from adaptive multi-theorem zero knowledge of the NIZK (cf. Def. 5 in §2.3), as the only difference between the two experiments is the fact that in the latter the proofs π_j^i are simulated. As a second step, we prove that $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)$ and $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, b)$ are also computationally indistinguishable. More in details, we show how to bound the probability that the output of the tampering oracle in the two experiments differs in the above described cases (a), (b) and (c). In case (a), we rely on computational binding (cf. Def. 3 in §2.2) of $(\mathbf{Gen}, \mathbf{Com})$ to argue that if the mauled shares are equal to the original, the underlying sharing is also unchanged. In case (b), we rely again on computational binding of $(\mathbf{Gen}, \mathbf{Com})$ using the fact that each value c_i is obtained by committing to both the share s_i and the position $i \in [n]$, so that permuting the commitments within a given block of the partition \mathcal{B} yields an invalid output. In case (c), we rely on simulation extractability (cf. Def. 6 in §2.3) to argue that whenever the outcome of a tampering query in the first hybrid is a valid message $\tilde{m} \in \mathcal{M}$, this value must be the same that can be extracted from the proofs associated to the commitments that have been modified; here is where we exploit the fact that the NIZK supports labels, since the extractor works as long as either the statement or the label mauled by the adversary are fresh.

Next, we show that no PPT attacker \mathbf{A} can distinguish between $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, 0)$ and $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, 1)$ with better than negligible probability. To this end, we build a reduction $\hat{\mathbf{A}}$ to leakage resilience of the underlying secret sharing Σ . In order to keep the exposition simple, let us first assume that \mathbf{A} is not allowed to ask leakage queries. Very roughly, the reduction works as follows.

- **(Simulate the CRS.)** At the beginning, $\hat{\mathbf{A}}$ receives the challenge CRS ω , samples the public key pk , and runs \mathbf{A} upon (ω, pk) with fresh randomness r . Upon receiving (\mathcal{B}, m_0, m_1) , then $\hat{\mathbf{A}}$ forwards the same tuple to the challenger.
- **(Learn the self-destruct index.)** Note that in the last hybrid, the tampering oracle computes the answers to \mathbf{A} 's tampering queries by computing a candidate answer for each set \mathcal{B}_i of the partition, and then this value is returned if the candidate values are all consistent (and otherwise a self-destruct is triggered). Since $\hat{\mathbf{A}}$ outputs the same partition \mathcal{B} chosen by \mathbf{A} , it can compute the different candidates by running \mathbf{A} with hard-wired randomness r inside the leakage oracle,⁹ and then use a pairwise independent hash function to determine using a binary search the first index where the candidates differ. By pairwise independence, this yields the index of the query p^* in which \mathbf{A} provokes a self-destruct with overwhelming probability, and by leaking at most $O(\lambda \log \lambda)$ bits from each block.
- **(Play the game.)** Once the self-destruct index p^* is known, the idea is that now $\hat{\mathbf{A}}$ can restart \mathbf{A} outside the leakage oracle, with the same randomness r , and answer to the first $p^* - 1$ tampering queries given the shares $\tilde{s}_{\mathcal{B}_i}$ within one of the partitions, after which the

⁸More in details, the extraction procedure extracts the missing shares from the proofs whose statement or label corresponds to one of the mauled commitments; since we might extract one of the shares from more than one proof, the procedure further checks that those the extracted shares are all consistent. Cf. Fig. 4 on page 21 for a precise description.

⁹To be more precise, this also requires to leak the initial commitments to each share, to simulate the NIZK proofs as done in the last hybrid, and to hard-wire those values in each leakage query. However, the latter can be done easily by the reduction.

answer to all remaining tampering queries is set to be \perp , so that \hat{A} can output the same guess of A and keep its advantage. Here, we rely on the fact that Σ is *augmented* leakage resilient, i.e. leakage resilience still holds even if \hat{A} is given the shares belonging to one of the partitions after the leakage is done. By Thm. 3 (cf. §3.2), this assumption is without loss of generality.

Finally, we show how to remove the simplifying assumption that A cannot leak from the shares. The difficulty when considering leakage is that we cannot run anymore the entire experiment with A inside the leakage oracle, as the answer to A 's leakage queries depends on the shares outside a given partition. However, note that in this case we can stop the execution whenever A asks a leakage query and inform the reduction to leak from the other shares whatever information is needed to continue the execution of each copy of A inside the leakage oracle.

This allows to obtain the answers to all leakage queries of A up to a self-destruct occurs. In order to obtain the answers to the remaining queries, we must re-run A inside the leakage oracle and adjust the simulation consistently with the self-destruct index being p^* . In the worst case, this requires $2\ell^*$ bits of leakage from the shares, yielding the final bound of $2\ell^* + n\gamma + O(\lambda \log \lambda)$. At the end, the reduction knows the answer to all leakage queries of A with hard-wired randomness r , and can thus win the game with the challenger in the same way as explained above.

4.3 Security Analysis

Correctness w.r.t. access structure \mathcal{A} is immediate. We thus need to prove that Σ^* satisfies privacy and continuous non-malleability.

Privacy. We need to show that

$$\{\mathbf{Privacy}_{\Sigma^*, \mathcal{A}}(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Privacy}_{\Sigma^*, \mathcal{A}}(\lambda, 1)\}_{\lambda \in \mathbb{N}}.$$

Below, we describe the experiment $\mathbf{Privacy}_{\Sigma^*, \mathcal{A}}(\lambda, b)$ after expanding the specification of \mathbf{Init}^* and \mathbf{Share}^* .

- Generate $\omega^* \leftarrow \mathbf{Init}^*(1^\lambda)$, where \mathbf{Init}^* is computed as follows:
 - upon receiving 1^λ , compute $\omega \leftarrow \mathbf{CRSGen}(1^\lambda)$, $pk \leftarrow \mathbf{Gen}(1^\lambda)$;
 - output $\omega^* = (\omega, pk)$.
- Run $(\alpha_1, m_0, m_1, \mathcal{U} \notin \mathcal{A}) \leftarrow \mathbf{A}_1(\omega^*)$.
- Compute $(s_1^*, \dots, s_n^*) \leftarrow \mathbf{Share}^*(\omega^*, m_b)$, where \mathbf{Share}^* is defined as follows:
 - upon receiving (ω^*, m) , parse $\omega^* = (\omega, pk)$ and compute $(s_1, \dots, s_n) \leftarrow \mathbf{Share}(m)$;
 - for all $i \in [n]$, generate $r_i \leftarrow \mathcal{R}$ and compute $c_i = \mathbf{Com}(pk, i || s_i; r_i)$ and, for all $j \in [n]$ such that $j \neq i$, compute $\pi_i^j = \mathbf{Prove}(\omega, c_j, (c_i, i || s_i, r_i))$;
 - for all $i \in [n]$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$;
 - output (s_1^*, \dots, s_n^*) .
- Run $b' \leftarrow \mathbf{A}_2(\alpha_1, (s_u^*)_{u \in \mathcal{U}})$.

Let $S = (S_0, S_1)$ be the simulator for the adaptive multi-theorem zero-knowledge property of Π , and consider the hybrid experiment $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, b)$ in which the algorithms \mathbf{Init}^* and \mathbf{Share}^* are replaced by \mathbf{Init}' and \mathbf{Share}' , defined below.

- Algorithm \mathbf{Init}' :
 - upon receiving 1^λ , compute $(\omega, \zeta) \leftarrow S_0(1^\lambda)$, $pk \leftarrow \mathbf{Gen}(1^\lambda)$;

- output $\omega^* = (\omega, pk)$.
- Algorithm **Share'** (with hard-wired ζ):
 - upon receiving (ω^*, m) , parse $\omega^* = (\omega, pk)$ and compute $(s_1, \dots, s_n) \leftarrow \text{Share}(m)$;
 - for all $i \in [n]$, generate $r_i \leftarrow \mathcal{R}$, compute $c_i = \text{Com}(pk, i || s_i; r_i)$, and for all $j \in [n]$ with $j \neq i$ let $\pi_i^j \leftarrow \mathcal{S}_1(\zeta, c_j, c_i)$;
 - for all $i \in [n]$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$;
 - output (s_1^*, \dots, s_n^*) .

Finally, consider the experiment $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}^t(\lambda, b)$, defined below.

- Generate $(\omega^*, \zeta) \leftarrow \text{Init}'(1^\lambda)$.
- Run $(\alpha_1, m_0, m_1, \mathcal{U} \notin \mathcal{A}) \leftarrow \mathcal{A}_1(\omega^*)$.
- Compute $(s_1, \dots, s_n) \leftarrow \text{Share}(m_b)$.
- Generate $\hat{m} \leftarrow \mathcal{M}$ and compute $(\hat{s}_1, \dots, \hat{s}_n) \leftarrow \text{Share}(\hat{m})$.
- For all $i \in [n]$, generate $r_i \leftarrow \mathcal{R}$ and compute c_i as follows. If $i \in \mathcal{U}$ or $i > t$, compute $c_i = \text{Com}(pk, i || s_i, r_i)$; else, compute $c_i = \text{Com}(pk, i || \hat{s}_i, r_i)$.
- For all $i, j \in [n]$ such that $i \neq j$, compute $\pi_i^j \leftarrow \mathcal{S}_1(\zeta, c_j, c_i)$.
- For all $i \in \mathcal{U}$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.
- Run $b' \leftarrow \mathcal{A}_2(\alpha_1, (s_u^*)_{u \in \mathcal{U}})$.

Note that, for $t = 0$, the condition $i > t$ is true for all $i \in [n]$, thus all the commitments are computed using the shares from the message m_b and the experiments $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, b)$ and $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}^0(\lambda, b)$ are equivalent.

We now prove that the original experiment is computationally close to the first hybrid experiment and, for fixed $b \in \{0, 1\}$, all the hybrid experiments defined above are statistically close. Finally we prove that, for $t = n$, the experiments with $b = 0$ and $b = 1$ are computationally close.

Lemma 3. $\forall b \in \{0, 1\}: \{\mathbf{Privacy}_{\Sigma^*, \mathcal{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}}$.

Proof. The proof is down to the adaptive multi-theorem zero-knowledge property of the underlying non-interactive zero-knowledge argument system $\Pi = (\text{CRSGen}, \text{Prove}, \text{Ver})$. By contradiction, assume that there exists a PPT attacker $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ that, for any fixed $b \in \{0, 1\}$, can distinguish between $\{\mathbf{Privacy}_{\Sigma^*, \mathcal{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}}$ and $\{\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}}$ with noticeable advantage. Consider the following PPT distinguisher \mathbf{D} attacking the zero-knowledge property of Π .

- Upon receiving a CRS ω generated either by $\text{CRSGen}(1^\lambda)$ or by the simulator $\mathcal{S}_0(1^\lambda)$ (depending on the experiment that \mathbf{D} is currently playing), generate $pk \leftarrow \mathcal{G}\text{en}(1^\lambda)$ and write $\omega^* = (\omega, pk)$.
- Run $(\alpha_1, m_0, m_1, \mathcal{U}) \leftarrow \mathcal{A}_1(\omega^*)$ and compute $s = (s_1, \dots, s_n) \leftarrow \text{Share}(m_b)$.
- For each $i \in [n]$, generate $r_i \leftarrow \mathcal{R}$ and compute $c_i = \text{Com}(pk, i || s_i; r_i)$.
- For each $i, j \in [n]$ such that $i \neq j$, query the challenger with $(c_j, (c_i, i || s_i, r_i))$, obtaining either $\pi_i^j \leftarrow \text{Prove}(\omega, c_j, (c_i, i || s_i, r_i))$ or $\pi_i^j \leftarrow \mathcal{S}_1(\zeta, c_j, c_i)$ (depending on the experiment that \mathbf{D} is currently playing).
- For each $i \in [n]$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.
- Write $s^* = (s_1^*, \dots, s_n^*)$.
- Run $b' \leftarrow \mathcal{A}_2(\alpha_1, (s_i^*)_{i \in \mathcal{U}})$.
- Output b' .

For the analysis, note that the simulation done by \mathbf{D} is perfect. Thus, \mathbf{D} distinguishes with noticeable advantage. The lemma follows. \square

Lemma 4. $\forall t \in [n]: \{\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}^{t-1}(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}^t(\lambda, b)\}_{\lambda \in \mathbb{N}}$.

Proof. The proof is down to the statistical hiding property of the underlying commitment scheme. Fix $t \in [n]$. By contradiction, assume that there exists a computationally unbounded adversary \mathbf{A} that can distinguish between the two experiments with noticeable advantage. Consider the following unbounded adversary $\hat{\mathbf{A}}$ attacking the statistical hiding property of (Gen, Com) .

- Upon receiving $pk \leftarrow_s \text{Gen}(1^\lambda)$, generate $(\omega, \zeta) \leftarrow_s \mathbf{S}_0(1^\lambda)$ and write $\omega^* = (\omega, pk)$.
- Run $(\alpha_1, m_0, m_1, \mathcal{U} \notin \mathcal{A}) \leftarrow_s \mathbf{A}_1(\omega^*)$.
- Generate $\hat{m} \leftarrow_s \mathcal{M}$ and compute $(\hat{s}_1, \dots, \hat{s}_n) \leftarrow_s \text{Share}(\hat{m})$.
- Send the pair $(t || s_t, t || \hat{s}_t)$ to the challenger, receiving a commitment c_t .
- For all $i \in [n]$ such that $i \neq t$, generate $r_i \leftarrow_s \mathcal{R}$ and compute c_i as follows. If $i \in \mathcal{U}$ or $i > t$, compute $c_i = \text{Com}(pk, i || s_i; r_i)$; else, compute $c_i = \text{Com}(pk, i || \hat{s}_i; r_i)$.
- For all $i, j \in [n]$ such that $i \neq j$, compute $\pi_i^j \leftarrow_s \mathbf{S}_1(\zeta, c_j, c_i)$.
- For all $i \in \mathcal{U}$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.
- Run $b' \leftarrow_s \mathbf{A}_2(\alpha_1, (s_u^*)_{u \in \mathcal{U}})$.
- Output b' .

Note that the only difference between the two experiments is how the commitment c_t is computed. In particular, in case $t \in \mathcal{U}$ the two experiments are identical and thus we can assume wlog. that $t \notin \mathcal{U}$. Furthermore, the simulation performed by $\hat{\mathbf{A}}$ is perfect, so that if \mathbf{A} tells apart the two experiments, $\hat{\mathbf{A}}$ distinguishes if the commitment c_t comes from $t || s_t$ or from $t || \hat{s}_t$, thus breaking the statistical hiding property of (Gen, Com) . The lemma follows. \square

Lemma 5. $\{\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}^n(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}^n(\lambda, 1)\}_{\lambda \in \mathbb{N}}$.

Proof. The proof is down to the privacy of the underlying secret sharing scheme Σ . By contradiction, assume that there exists a PPT adversary \mathbf{A} that can distinguish between the experiment with $b = 0$ and $b = 1$ with noticeable advantage. Consider the following PPT adversary $\hat{\mathbf{A}}$ attacking the privacy of Σ .

- Run $(\omega, \zeta) \leftarrow_s \mathbf{S}_0(1^\lambda)$ and $pk \leftarrow_s \text{Gen}(1^\lambda)$ and write $\omega^* = (\omega, pk)$.
- Run $(\alpha_1, m_0, m_1, \mathcal{U} \notin \mathcal{A}) \leftarrow_s \mathbf{A}_1(\omega^*)$.
- Output (m_0, m_1, \mathcal{U}) to the challenger, receiving the shares $(s_i)_{i \in \mathcal{U}}$.
- Generate $\hat{m} \leftarrow_s \mathcal{M}$ and compute $(\hat{s}_1, \dots, \hat{s}_n) \leftarrow_s \text{Share}(\hat{m})$.
- For all $i \in [n]$, generate $r_i \leftarrow_s \mathcal{R}$ and compute c_i as follows. If $i \in \mathcal{U}$, compute $c_i = \text{Com}(pk, i || s_i; r_i)$; otherwise, compute $c_i = \text{Com}(pk, i || \hat{s}_i; r_i)$.
- For all $i, j \in [n]$ such that $i \neq j$, compute $\pi_i^j \leftarrow_s \mathbf{S}_1(\zeta, c_j, c_i)$.
- For all $i \in \mathcal{U}$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.
- Run $b' \leftarrow_s \mathbf{A}_2(\alpha_1, (s_u^*)_{u \in \mathcal{U}})$.
- Output b' .

For the analysis, note that the reduction is perfect. Therefore, if \mathbf{A} distinguishes between the two experiments with noticeable advantage, $\hat{\mathbf{A}}$ distinguishes the two messages using the unauthorized set of shares with noticeable advantage, thus breaking privacy for Σ . This concludes the proof. \square

Continuous non-malleability. Recall the experiment $\mathbf{JSTamper}_{\Sigma^*, \mathcal{A}}(\lambda, b)$ defining continuous non-malleability for Σ^* . In particular, expanding the definitions of Init^* and Share^* :

<p>Oracle $\mathcal{O}'_{\text{nmss}}(s^*, \mathcal{T}, (f_1, \dots, f_\beta))$:</p> <p>If stop = true return \perp</p> <p>$\forall i \in [\beta] : \tilde{s}_{\mathcal{B}_i}^* := f_i(s_{\mathcal{B}_i}^*)$</p> <p>$\tilde{s}^* = (\tilde{s}_1^*, \dots, \tilde{s}_n^*)$</p> <p>If $\text{Check}_{\mathcal{T}}(\omega^*, \tilde{s}^*) = 0$</p> <p style="padding-left: 20px;">Return \perp and stop \leftarrow true</p> <p>$\forall t \in \mathcal{T}$:</p> <p style="padding-left: 20px;">$s_t^* = (s_t, r_t, (c_j)_{j \neq t}, (\pi_j^t)_{j \neq t})$</p> <p style="padding-left: 20px;">$\tilde{s}_t^* = (\tilde{s}_t, \tilde{r}_t, (\tilde{c}_j)_{j \neq t}, (\tilde{\pi}_j^t)_{j \neq t})$</p> <p>If $\forall j \in [n], \tilde{c}_j = c_j$, return \heartsuit</p> <p>If $\exists j_1, j_2 \in [n] : \tilde{c}_{j_1} = c_{j_2}$, abort</p> <p>$\tilde{\mathcal{I}} = \{i \in [n] : \tilde{c}_i \notin (c_j)_{j \in [n]}\}$</p> <p>$\forall i \in [\beta] : \tilde{m}_i \leftarrow \text{Extract}_{\mathcal{T}, \tilde{\mathcal{I}}}(\xi, \tilde{s}_{\mathcal{B}_i}^*)$</p> <p>If $\exists i, j \in [\beta] : \tilde{m}_i \neq \tilde{m}_j$ or $\tilde{m}_i = \perp$:</p> <p style="padding-left: 20px;">Return \perp and stop \leftarrow true</p> <p>$\tilde{m} = \tilde{m}_i$</p> <p>If $\tilde{m} \in \{m_0, m_1\}$</p> <p style="padding-left: 20px;">Return \heartsuit</p> <p>Else return \tilde{m}</p>	<p>Algorithm $\text{Check}_{\mathcal{T}}(\omega^*, (\tilde{s}_i^*)_{i \in \mathcal{I}})$:</p> <p>$\forall i \in \mathcal{T} \cap \mathcal{I} : \tilde{s}_i^* = (\tilde{s}_i, \tilde{r}_i, (\tilde{c}_j^i)_{j \neq i}, (\tilde{\pi}_j^i)_{j \neq i})$</p> <p>If $\exists i_1, i_2 \in \mathcal{T}, j \in [n] : \tilde{c}_j^{i_1} \neq \tilde{c}_j^{i_2}$</p> <p style="padding-left: 20px;">Return 0</p> <p>$\forall i \in \mathcal{T} \cap \mathcal{I} : \tilde{s}_i^* = (\tilde{s}_i, (\tilde{c}_j)_{j \neq i}, (\tilde{\pi}_j^i)_{j \neq i})$</p> <p>If $\exists i \in \mathcal{T} \cap \mathcal{I} : \tilde{c}_i \neq \text{Com}(pk, i \tilde{s}_i; \tilde{r}_i)$</p> <p style="padding-left: 20px;">Return 0</p> <p>If $\exists j \in \mathcal{T}, i \in [n] : \text{Ver}(\omega, \tilde{c}_j, (\tilde{c}_i, \tilde{\pi}_i^j)) = 0$</p> <p style="padding-left: 20px;">Return 0</p> <p>Return 1</p> <p>Algorithm $\text{Extract}_{\mathcal{T}, \tilde{\mathcal{I}}}(\xi, (\tilde{s}_i^*)_{i \in \mathcal{I}})$:</p> <p>$\forall i \in \mathcal{T} \cap \mathcal{I} : \tilde{s}_i^* = (\tilde{s}_i, \tilde{r}_i, (\tilde{c}_j)_{j \neq i}, (\tilde{\pi}_j^i)_{j \neq i})$</p> <p>$\forall t \in \mathcal{T}$:</p> <p style="padding-left: 20px;">If $t \in \tilde{\mathcal{I}}$:</p> <p style="padding-left: 40px;">$\forall j \in [n], j \neq t$:</p> <p style="padding-left: 60px;">$(\text{pos}_t^j \tilde{s}_t^j, r_t^j) \leftarrow \text{K}(\xi, \tilde{c}_j, (\tilde{c}_t, \tilde{\pi}_t^j))$</p> <p style="padding-left: 20px;">Else:</p> <p style="padding-left: 40px;">$\forall j \in \tilde{\mathcal{I}}$:</p> <p style="padding-left: 60px;">$(\text{pos}_t^j \tilde{s}_t^j, r_t^j) \leftarrow \text{K}(\xi, \tilde{c}_j, (c_t, \tilde{\pi}_t^j))$</p> <p>If $\exists j, j_1, j_2 : \tilde{s}_t^{j_1} \neq \tilde{s}_t^{j_2}$ or $r_t^{j_1} \neq r_t^{j_2}$</p> <p style="padding-left: 40px;">or $\text{pos}_t^j \neq t$</p> <p style="padding-left: 20px;">Return \perp</p> <p style="padding-left: 20px;">Write $\tilde{s}_t = \tilde{s}_t^j$ for any j</p> <p>Return $\text{Rec}((\tilde{s}_t)_{t \in \mathcal{T}})$</p>
---	---

Figure 3: Construction of the oracle $\mathcal{O}'_{\text{nmss}}$ used in the experiment $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, b)$.

JSTamper $_{\Sigma^*, \mathbf{A}}(\lambda, b)$:

$\omega \leftarrow \text{CRSGen}(1^\lambda), pk \leftarrow \text{Gen}(1^\lambda), \omega^* := (\omega, pk)$

$(\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_\beta), m_0, m_1, \alpha_1) \leftarrow \mathbf{A}_1(\omega^*)$

$(s_1, \dots, s_n) \leftarrow \text{Share}(m_b)$

$\forall i \in [n] : r_i \leftarrow \mathcal{R}, c_i = \text{Com}(pk, i || s_i; r_i)$

$\forall i, j \in [n], i \neq j : \pi_i^j \leftarrow \text{Prove}(\omega, c_j, (c_i, i || s_i, r_i))$

$\forall i \in [n] : s_i^* := (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$

stop \leftarrow **false**

$(\alpha_2, i^* \in [\beta]) \leftarrow \mathbf{A}_2^{\mathcal{O}_{\text{nmss}}(s^*, \cdot, \cdot), \mathcal{O}_{\text{leak}}(s^*, \cdot)}(\alpha_1)$

Return $\mathbf{A}_3(\alpha_2, s_{\mathcal{B}_{i^*}})$

Consider the following hybrid experiments:

- $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)$: consider the simulator $\mathbf{S} = (\mathbf{S}_0, \mathbf{S}_1)$ for the NIZK and replace the instructions $\omega \leftarrow \text{CRSGen}(1^\lambda)$ and $\pi_i^j \leftarrow \text{Prove}(\omega, c_j, (c_i, i || s_i))$ with $(\omega, \zeta, \xi) \leftarrow \mathbf{S}_0(1^\lambda)$ and $\pi_i^j \leftarrow \mathbf{S}_1(\zeta, c_j, c_i)$ respectively.
- $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, b)$: replace the oracle $\mathcal{O}_{\text{nmss}}$ with $\mathcal{O}'_{\text{nmss}}$ described in Fig. 3.

We first prove that the above experiments are computationally close, i.e. for all $b \in \{0, 1\}$,

$$\{\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, b)\}_{\lambda \in \mathbb{N}}.$$

Then, we prove that

$$\{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, 1)\}_{\lambda \in \mathbb{N}}$$

by reduction to the underlying LRSS scheme Σ , thus proving continuous non-malleability:

$$\{\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, 1)\}_{\lambda \in \mathbb{N}}.$$

Lemma 6. $\forall b \in \{0, 1\}: \{\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)\}_{\lambda \in \mathbb{N}}.$

Proof. The proof is down to the adaptive multi-theorem zero-knowledge property of the underlying non-interactive zero-knowledge argument system $\Pi = (\text{CRSGen}, \text{Prove}, \text{Ver})$. By contradiction, assume that there exists a PPT distinguisher $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3)$ that, for any fixed $b \in \{0, 1\}$, can distinguish between $\{\mathbf{JSTamper}_{\Sigma^*, \mathbf{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}}$ and $\{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)\}_{\lambda \in \mathbb{N}}$ with noticeable advantage. Consider the following PPT distinguisher \mathbf{D} attacking the zero-knowledge property of Π .

- Upon receiving a CRS ω that is generated either from $\text{CRSGen}(1^\lambda)$ or from the simulator $\mathbf{S}_0(1^\lambda)$ (depending on the experiment that \mathbf{D} is currently playing), generate $pk \leftarrow \text{Gen}(1^\lambda)$ and write $\omega^* = (\omega, pk)$.
- Run $(\mathcal{B}, m_0, m_1, \alpha_1) \leftarrow \mathbf{A}_1(\omega^*)$ and compute $s = (s_1, \dots, s_n) \leftarrow \text{Share}(m_b)$.
- For each $i \in [n]$, generate $r_i \leftarrow \mathcal{R}$ and compute $c_i = \text{Com}(pk, i || s_i; r_i)$.
- For each $i, j \in [n]$ such that $i \neq j$, query the challenger with $(c_j, (c_i, i || s_i, r_i))$, obtaining either $\pi_i^j \leftarrow \text{Prove}(\omega, c_j, (c_i, i || s_i, r_i))$ or $\pi_i^j \leftarrow \mathbf{S}_1(\zeta, c_j, c_i)$ (depending on the experiment that \mathbf{D} is currently playing).
- For each $i \in [n]$, write $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.
- Write $s^* = (s_1^*, \dots, s_n^*)$ and set $\text{stop} \leftarrow \text{false}$.
- Run $(\alpha_2, i^*) \leftarrow \mathbf{A}_2^{\mathcal{O}_{\text{nmss}}(s^*, \cdot), \mathcal{O}_{\text{leak}}(s^*, \cdot)}(\alpha_1)$.
- Return the same as $\mathbf{A}_3(\alpha_2, s_{\mathcal{B}_{i^*}})$.

For the analysis, note that the simulation done by \mathbf{D} is perfect. In particular, \mathbf{D} can impersonate the oracles $\mathcal{O}_{\text{nmss}}$ and $\mathcal{O}_{\text{leak}}$ and can answer all the queries. Thus, \mathbf{D} distinguishes with noticeable probability. The lemma follows. \square

Lemma 7. $\forall b \in \{0, 1\}: \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^1(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, b)\}_{\lambda \in \mathbb{N}}.$

Proof. The only difference between the two experiments is how the oracle answers the tampering queries. In both experiments the following checks are performed:

- if there exist indices $i_1, i_2 \in \mathcal{T}, j \in [n]$ such that $c_j^{i_1} \neq c_j^{i_2}$, self-destruct;
- if there exist an index $i \in \mathcal{T}$ such that $\text{Com}(pk, i || s_i; r_i) \neq c_i$, self-destruct;
- if there exist indices $i, j \in \mathcal{T}$ such that $\text{Ver}(\omega, c_j, (c_i, \pi_i^j)) = 0$, self-destruct.

In particular, the check phase is perfectly simulated. Then, the oracle $\mathcal{O}'_{\text{nmss}}$ finds the indices in which a mauled commitment appears.

- If $\forall j \in [n], \tilde{c}_j = c_j$, then $\text{Com}(pk, j || \tilde{s}_j; \tilde{r}_j) = \text{Com}(pk, j || s_j; r_j)$; call **Bad**₁ the event that $\tilde{s}_j \neq s_j$. If **Bad**₁ does not happen, the adversary didn't modify any of the shares, and thus the reconstructed message also remains the same $m_b \in \{m_0, m_1\}$.
- If there exist two distinct $j_1, j_2 \in [n]$ such that $\tilde{c}_{j_2} = c_{j_1}$, abort the experiment; call this event **Bad**₂.
- If none of the above holds, then there exists at least one fresh commitment \tilde{c}_j , namely a commitment $\tilde{c}_j \notin (c_i)_{i \in [n]}$; call $\tilde{\mathcal{I}}$ the set of indices for which the commitment is fresh.

Next, the oracle $\mathcal{O}'_{\text{nmss}}$ reconstructs a message for each block \mathcal{B}_i by extracting the needed shares from the proofs. In particular, for all $t \in \mathcal{T}$, the oracle extracts the pair $(\text{pos}_t^j || \tilde{s}_t^j, r_t^j)$ from $(\tilde{c}_t, \tilde{\pi}_t^j)$ using \tilde{c}_j as label, and this happens for all $j \in [n]$ if \tilde{c}_t is fresh and for all $j \in \tilde{\mathcal{I}}$ otherwise. At this point of the execution, all proofs verify (otherwise, the oracle would have already provoked a self-destruct during the check phase) and all tuples $(\tilde{c}_j, (\tilde{c}_t, \tilde{\pi}_t^j))$ with the above condition on j are fresh, so that the extractor K can extract the pair $(\text{pos}_t^j || \tilde{s}_t^j, \tilde{r}_t^j)$. Call **Bad**₃ the event that $\text{Com}(pk, \text{pos}_t^j || \tilde{s}_t^j; \tilde{r}_t^j) \neq \tilde{c}_t = \text{Com}(pk, t || \tilde{s}_t; \tilde{r}_t)$ and **Bad**₄ the event that $\text{pos}_t^j || \tilde{s}_t^j \neq t || \tilde{s}_t$ for any j . If neither **Bad**₃ nor **Bad**₄ happen, the algorithm **Extract** managed to extract all the shares $(\tilde{s}_t)_{t \in \mathcal{T}}$ and can reconstruct the message. In fact, if **Bad**₃ does not happen, the extracted shares are consistent with the mauled commits and, if **Bad**₄ does not happen, for all j_1, j_2 , $\text{pos}_t^{j_1} || \tilde{s}_t^{j_1} = \text{pos}_t^{j_2} || \tilde{s}_t^{j_2} = t || \tilde{s}_t$. Finally, the oracle computes $\tilde{m}_i = \text{Rec}((\tilde{s}_t)_{t \in \mathcal{T}})$ and, since $\text{Rec}((\tilde{s}_t)_{t \in \mathcal{T}})$ does not depend on the set \mathcal{B}_i chosen to extract the shares, we can simply write $\tilde{m} = \tilde{m}_i$.

The last part of the reconstruction is identical to the one performed in the original oracle: if $\tilde{m} = \perp$, self-destruct; if $\tilde{m} \in \{m_0, m_1\}$, output \heartsuit ; otherwise, output \tilde{m} . Call **Bad** = **Bad**₁ \cup **Bad**₂ \cup **Bad**₃ \cup **Bad**₄. If **Bad** does not happen, the two oracles $\mathcal{O}_{\text{nmss}}$ and $\mathcal{O}'_{\text{nmss}}$ are equivalent. It remains to show that the event **Bad** happens with negligible probability.

- The event **Bad**₁ happens when $\text{Com}(pk, j || \tilde{s}_j; \tilde{r}_j) = \text{Com}(pk, j || s_j; r_j)$ but $\tilde{s}_j \neq s_j$. A straightforward reduction to the computational binding property of (Gen, Com) shows that this event happens with negligible probability. By contradiction, suppose that there exists an adversary $\hat{\mathsf{A}}$ that triggers **Bad**₁ with noticeable probability. Then, we can construct an algorithm $\hat{\mathsf{A}}$ that, upon receiving a setup string $pk \leftarrow_{\$} \text{Gen}(1^\lambda)$, emulates the experiment $\text{Hyb}_{\Sigma^*, \mathsf{A}}^1(\lambda, b)$ and, upon receiving a tampering query by A that triggers the event **Bad**₁, halts the execution of A and outputs the values $(j || \tilde{s}_j, \tilde{r}_j)$ and $(j || s_j, r_j)$ that result in the same commitment c_j , thus breaking the computational binding property of (Gen, Com) .
- The event **Bad**₂ happens when $\text{Com}(pk, j_1 || \tilde{s}_{j_1}; \tilde{r}_{j_1}) = \text{Com}(pk, j_2 || s_{j_2}; r_{j_2})$ but $j_1 || \tilde{s}_{j_1} \neq j_2 || s_{j_2}$. Again, a straightforward reduction to the computational binding property of (Gen, Com) shows that this event happens with negligible probability. The proof is identical to the previous one and thus omitted.
- The event **Bad**₃ happens when the pair of commitments $(\tilde{c}_j, \tilde{c}_i)$ is fresh and $\text{Ver}(\omega, \tilde{c}_j, (\tilde{c}_i, \pi_i^j)) = 1$, but the relation over $\mathcal{L}_{\text{Com}}^{pk}$ does not hold for $(\tilde{c}_i, (i || \tilde{s}_i, \tilde{r}_i))$, i.e. $\text{Com}(pk, i || \tilde{s}_i; \tilde{r}_i) \neq \tilde{c}_i$. A reduction to the true simulation extractability of $(\text{CRSGen}, \text{Prove}, \text{Ver})$ shows that this event happens with negligible probability. By contradiction, suppose that there exists an adversary A that triggers **Bad**₃ with noticeable probability. Then, we can construct the following algorithm $\hat{\mathsf{A}}$.
 - Upon receiving ω such that $(\omega, \zeta, \xi) \leftarrow_{\$} \mathsf{S}_0(1^\lambda)$, generate $pk \leftarrow_{\$} \text{Gen}(1^\lambda)$ and run $(\mathcal{B}, m_0, m_1, \alpha_1) \leftarrow_{\$} \mathsf{A}_1((\omega, pk))$.
 - Compute $(s_1, \dots, s_n) \leftarrow_{\$} \text{Share}(m_b)$ and, for all $i \in [n]$, generate $r_i \leftarrow_{\$} \mathcal{R}$ and compute $c_i = \text{Com}(pk, i || s_i; r_i)$.
 - For all $i, j \in [n], i \neq j$, query the challenger with $(c_j, (c_i, i || s_i, r_i))$, receiving $\pi_i^j \leftarrow_{\$} \mathsf{S}_1(\zeta, c_j, c_i)$.
 - For all $i \in [n]$, $s_i^* = (s_i, r_i, (c_j)_{j \neq i}, (\pi_j^i)_{j \neq i})$.
 - Run $\mathsf{A}_2(\alpha_1)$, answering to all its leakage queries as in the original experiment. Upon receiving a tampering query:
 - * perform the same steps of $\mathcal{O}_{\text{nmss}}$;
 - * upon mauling s_t^* , for all $t \in \mathcal{T}, j \in [n], j \neq t$, if $(\tilde{c}_t, \tilde{c}_j)$ is fresh, the commitments are correct and the proof $\tilde{\pi}_j^t$ verifies, save for later the tuple $(\tilde{c}_t, \tilde{c}_j, \tilde{\pi}_j^t)$.
 - Randomly choose and output one of the tuples $(\tilde{c}_t, \tilde{c}_j, \tilde{\pi}_j^t)$.

For the analysis, we next prove that the reduction breaks true simulation extractability with noticeable probability. Note first that all the proofs simulated by S_0 are relative to true statements, thus we can apply true simulation extractability. Call p the total number of tampering queries asked by A and call q the query from which \hat{A} chooses the output. If \mathbf{Bad}_3 happens, the probability of \hat{A} guessing the right query is at least $\frac{1}{p}$; call p^* one of such queries. Furthermore, if \mathbf{Bad}_3 happens in query $p^* = q$, \hat{A} can choose among up to $n^2 - n$ tuples; call (j^*, t^*) the indices referring to one of the tuples in query p^* triggering the event \mathbf{Bad}_3 and (j, t) the indices referring to the tuple chosen by \hat{A} . The probability of \hat{A} breaking true simulation extractability is

$$\begin{aligned} \mathbb{P}[\hat{A} \text{ wins}] &\geq \mathbb{P}[\mathbf{Bad}_3 \wedge q = p^* \wedge (j, t) = (j^*, t^*)] \\ &\geq \mathbb{P}[\mathbf{Bad}_3] \mathbb{P}[q = p^* | \mathbf{Bad}_3] \mathbb{P}[(j, t) = (j^*, t^*) | \mathbf{Bad}_3 \wedge q = p^*]. \end{aligned}$$

Since $\mathbb{P}[(j, t) = (j^*, t^*) | \mathbf{Bad}_3 \wedge q = p^*] = \frac{1}{n^2 - n}$ and n is polynomial in λ , the third factor is $\frac{1}{\text{poly}(\lambda)}$. Similarly, $\mathbb{P}[q = p^* | \mathbf{Bad}_3] = \frac{1}{p(\lambda)} = \frac{1}{\text{poly}(\lambda)}$, being $p(\lambda)$ polynomial in λ . Finally, if \mathbf{Bad}_3 happens with noticeable probability, all the three factors are $\frac{1}{\text{poly}(\lambda)}$ and their product is $\frac{1}{\text{poly}(\lambda)}$. This concludes the reduction.

- The event \mathbf{Bad}_4 happens when the extractor K successfully extracts a pair $(\text{pos}_t^j || \tilde{s}_t^j, \tilde{r}_t^j)$ that verifies the relation with \tilde{c}_t (i.e. $\text{Com}(pk, \text{pos}_t^j || \tilde{s}_t^j; \tilde{r}_t^j) = \tilde{c}_t$), but $\text{pos}_t^j || \tilde{s}_t^j \neq t || \tilde{s}_t$. Since $\tilde{c}_t = \text{Com}(pk, t || \tilde{s}_t; \tilde{r}_t)$, a straightforward reduction to the computational binding property of (Gen, Com) shows that this event happens with negligible probability. The proof is similar to the ones about \mathbf{Bad}_1 and \mathbf{Bad}_2 and thus omitted.

Finally, since \mathbf{Bad} is the union of the above events, its probability is less than or equal to the sum of the probabilities of the \mathbf{Bad}_i , thus negligible. The lemma follows. \square

Lemma 8. $\{\text{Hyb}_{\Sigma^*, A}^2(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Hyb}_{\Sigma^*, A}^2(\lambda, 1)\}_{\lambda \in \mathbb{N}}$.

Proof. The proof is down to the ℓ -bounded leakage-resilience property of Σ . Here we consider an auxiliary family of weakly universal hash functions¹⁰ Ψ . By contradiction, assume that there exists a PPT adversary A that can tell apart if the scheme has been applied on m_0 or on m_1 . Consider the following PPT adversary \hat{A} attacking ℓ -bounded leakage-resilience of Σ .

1. Run $pk \leftarrow \text{Gen}(1^\lambda)$ and $(\omega, \zeta, \xi) \leftarrow S_0(1^\lambda)$ and write $\omega^* = (\omega, pk)$.
2. Generate $r^A = (r_1^A, r_2^A, r_3^A) \leftarrow \mathcal{R}^A$ and run $(\mathcal{B}, m_0, m_1, \alpha_1) \leftarrow A_1(\omega^*; r_1^A)$.
3. Output (\mathcal{B}, m_0, m_1) , obtaining access to the leakage oracle.
4. For all $i \in [n]$, generate $r_i \leftarrow \mathcal{R}$.
5. Query $(\hat{g}_i^{\text{com}}(pk, r, \cdot))_{i \in [\beta]}$ to the leakage oracle, obtaining (c_1, \dots, c_n) .
6. Simulate the proofs $(\pi_i^j)_{j \neq i} \leftarrow S_1(\zeta, c_j, c_i)$.
7. Initialize the values $\text{rcp} = ((r_i)_{i \in [n]}, (c_i)_{i \in [n]}, (\pi_i^j)_{j \neq i})$, $\text{adv} = (A_2, \alpha_1, r_2^A)$ and $\text{state} = (\omega^*, \xi, \text{adv}, \text{rcp})$ and the empty string $\Lambda \leftarrow \varepsilon$.
8. Run the following loop:
 - In step q , query $(\hat{g}_i^{\text{sd}}(\text{state}, \perp, \Lambda, 0, \cdot))_{i \in [\beta]}$ to the leakage oracle, obtaining, for all $i \in [\beta]$, either $(\text{leak}, \Lambda_q^i)$ or $(\text{done}, q_{\text{tamper}}^i)$. Furthermore, for each $i \in [\beta]$, after receiving $(\text{done}, q_{\text{tamper}}^i)$, we can replace \hat{g}_i^{sd} with a function returning ε in order to save leakage bits.

¹⁰A family of hash functions $\Psi = \{\psi : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda\}$ is *weakly universal* if for all $x, y \in \{0, 1\}^*$ such that $x \neq y$, $\mathbb{P}[\psi(x) = \psi(y) | \psi \leftarrow \Psi] \leq \frac{1}{2^\lambda}$.

<p>Function $\hat{g}_i^{\text{com}}(pk, r, (s_j)_{j \in \mathcal{B}_i})$:</p> <p>Parse $r = (r_1, \dots, r_n)$</p> <p>$\forall j \in \mathcal{B}_i : c_j = \text{Com}(pk, j s_j; r_j)$</p> <p>Return $(c_j)_{j \in \mathcal{B}_i}$</p> <p>Function $\hat{g}_i^{\text{sd}}(\text{state}, \psi, \Lambda, q^*, (s_j)_{j \in \mathcal{B}_i})$:</p> <p>Parse $\text{state} = (\omega^*, \xi, \text{adv}, \text{rcp})$</p> <p>Parse $\omega^* = (\omega, pk)$ and $\text{adv} = (\mathbf{A}, \alpha, r^{\mathbf{A}})$</p> <p>Parse $\text{rcp} = ((r_j)_{j \in [n]}, (c_j)_{j \in [n]}, (\pi_j^i)_{j \neq i})$</p> <p>Parse $\Lambda = (\Lambda_1, \Lambda_2, \dots)$</p> <p>$\forall j \in \mathcal{B}_i : s_j^* = (s_j, r_j, (c_\beta)_{\beta \neq j}, (\pi_\beta^j)_{\beta \neq j})$</p> <p>$q_{\text{leak}} \leftarrow 0, q_{\text{tamper}} \leftarrow 0$</p> <p>stop \leftarrow false</p> <p>$x_i \leftarrow \varepsilon$</p> <p>Run $\mathbf{A}(\alpha, r)$ as follows:</p> <p style="padding-left: 20px;">Upon leak query (g_1, \dots, g_k):</p> <p style="padding-left: 40px;">$q_{\text{leak}} \leftarrow q_{\text{leak}} + 1$</p> <p style="padding-left: 40px;">If Λ contains $\Lambda_{q_{\text{leak}}}$, answer $\Lambda_{q_{\text{leak}}}$</p> <p style="padding-left: 40px;">Else, quit and return $(\text{leak}, g_i(s_{\mathcal{B}_i}^*))$</p> <p style="padding-left: 20px;">Upon tamper query $(\mathcal{T}, (f_1, \dots, f_k))$:</p> <p style="padding-left: 40px;">$q_{\text{tamper}} \leftarrow q_{\text{tamper}} + 1$</p> <p style="padding-left: 40px;">$(\tilde{m}, \tilde{c}) \leftarrow \hat{f}_i(\text{state}, \mathcal{T}, f_i, (s_j)_{j \in \mathcal{B}_i})$</p> <p style="padding-left: 40px;">$x_i \leftarrow x_i (\tilde{m}, \tilde{c})$</p> <p style="padding-left: 40px;">If $q_{\text{tamper}} = q^*$</p> <p style="padding-left: 60px;">$y_i \leftarrow \psi(x_i)$</p> <p style="padding-left: 40px;">Quit and return (hash, y_i)</p> <p style="padding-left: 40px;">Else, answer \tilde{m}.</p> <p>If $q^* = 0$ return $(\text{done}, q_{\text{tamper}})$</p> <p>Else return (done)</p>	<p>Function $\hat{f}_i(\text{state}, \mathcal{T}, f_i, (s_j)_{j \in \mathcal{B}_i})$:</p> <p>If stop = true return \perp</p> <p>Parse state, rcp as in \hat{g}_i^{sd}</p> <p>$\forall j \in \mathcal{B}_i : s_j^* := (s_j, r_j, (c_\beta)_{\beta \neq j}, (\pi_\beta^j)_{\beta \neq j})$</p> <p>$\tilde{s}_{\mathcal{B}_i}^* := f_i(s_{\mathcal{B}_i}^*)$</p> <p>If $\text{Check}_{\mathcal{T}}(\omega^*, \tilde{s}_{\mathcal{B}_i}^*) = 0$</p> <p style="padding-left: 20px;">Return \perp and stop \leftarrow true</p> <p>$\forall t \in \mathcal{T} : \tilde{s}_t^* = (\tilde{s}_t, \tilde{r}_t, (\tilde{c}_j)_{j \neq t}, (\tilde{\pi}_j^t)_{j \neq t})$</p> <p>$\tilde{c} := (\tilde{c}_j)_{j \in [n]}$</p> <p>If $\forall j \in [n], \tilde{c}_j = c_j$, return (\heartsuit, \tilde{c})</p> <p>If $\exists j_1, j_2 \in [n] : \tilde{c}_{j_1} = c_{j_2}$, abort</p> <p>$\tilde{\mathcal{I}} = \{i \in [n] : \tilde{c}_i \notin (c_j)_{j \in [n]}\}$</p> <p>$\tilde{m} \leftarrow \text{Extract}_{\mathcal{T}, \tilde{\mathcal{I}}}(\xi, \tilde{s}_{\mathcal{B}_i}^*)$</p> <p>If $\tilde{m} = \perp$ return \perp and stop \leftarrow true</p> <p>If $\tilde{m} \in \{m_0, m_1\}$, return (\heartsuit, \tilde{c})</p> <p>Else return (\tilde{m}, \tilde{c})</p> <p>Algorithm $\hat{g}_i^{\text{leak}}(\text{state}, \Lambda, p^*, (s_j)_{j \in \mathcal{B}_i})$:</p> <p>Parse $\text{state}, \text{adv}, \text{rcp}, \Lambda$ as in \hat{g}_i^{sd}</p> <p>$\forall j \in \mathcal{B}_i : s_j^* := (s_j, r_j, (c_\beta)_{\beta \neq j}, (\pi_\beta^j)_{\beta \neq j})$</p> <p>$q_{\text{leak}} \leftarrow 0, q_{\text{tamper}} \leftarrow 0, \text{stop} \leftarrow \text{false}$</p> <p>Run $(\alpha_2, i_{\text{aug}}) \leftarrow \mathbf{A}(\alpha, r)$ as follows:</p> <p style="padding-left: 20px;">Upon leak query (g_1, \dots, g_k):</p> <p style="padding-left: 40px;">$q_{\text{leak}} \leftarrow q_{\text{leak}} + 1$</p> <p style="padding-left: 40px;">If Λ contains $\Lambda_{q_{\text{leak}}}$, answer $\Lambda_{q_{\text{leak}}}$</p> <p style="padding-left: 40px;">Else, quit and return $(\text{leak}, g_i(s_{\mathcal{B}_i}^*))$</p> <p style="padding-left: 20px;">Upon tamper query $(\mathcal{T}, (f_1, \dots, f_k))$:</p> <p style="padding-left: 40px;">$q_{\text{tamper}} \leftarrow q_{\text{tamper}} + 1$</p> <p style="padding-left: 40px;">If $q_{\text{tamper}} \geq p^*$, answer \perp</p> <p style="padding-left: 40px;">Else:</p> <p style="padding-left: 60px;">$(\tilde{m}, \tilde{c}) \leftarrow \hat{f}_i(\text{state}, \mathcal{T}, f_i, (s_j)_{j \in \mathcal{B}_i})$</p> <p style="padding-left: 60px;">Answer \tilde{m}</p> <p>Return $(\text{done}, i_{\text{aug}})$</p>
--	---

Figure 4: Construction of the functions used by the reduction to emulate tampering with leakage.

- If the result of the query contains at least an element of the form $(\text{leak}, \Lambda_q^i)$, replace any missing Λ_q^i with ε , write $\Lambda_q = (\Lambda_q^1, \dots, \Lambda_q^\beta)$ and set $\Lambda \leftarrow \Lambda || \Lambda_q$;
 - else, break the loop, obtaining the string Λ and the (temporary) number of tamper queries $q_{\text{tamper}} = \min_i \{q_{\text{tamper}}^i\}$.
9. Set $\text{range} = (q_0, q_1) = (0, q_{\text{tamper}})$.
10. Run the following loop:
- If $q_0 = q_1$, break the loop obtaining $p^* \leftarrow q_0 = q_1$; else, set $q_m = \lfloor \frac{q_0 + q_1}{2} \rfloor$.
 - Sample $\psi \leftarrow \text{s} \Psi$.
 - Query $(\hat{g}_i^{\text{sd}}(\text{state}, \psi, \Lambda, q_m, \cdot))_{i \in [\beta]}$ to the leakage oracle, obtaining, for all $i \in [\beta]$,

- either (hash, y_i) or (done).
- If there exist some $i, i_1, i_2 \in [\beta]$ such that either $y_{i_1} \neq y_{i_2}$ or the i -th result of the query is (done), replace $(q_0, q_1) \leftarrow (q_0, q_m)$;
 - else, replace $(q_0, q_1) \leftarrow (q_m + 1, q_1)$.
11. Set $\Lambda \leftarrow \varepsilon$ (i.e. discard all the leakage queries).
 12. Run the following loop:
 - In step q , query $(\hat{g}_i^{\text{leak}}(\text{state}, \Lambda, p^*, \cdot))_{i \in [\beta]}$ to the leakage oracle, obtaining, for all $i \in [\beta]$, either (leak, Λ_q^i) or (done, i_{aug}).
 - If the result of the query contains (leak, Λ_q^i) $_{i \in [\beta]}$, write $\Lambda_q = (\Lambda_q^1, \dots, \Lambda_q^\beta)$ and set $\Lambda \leftarrow \Lambda \parallel \Lambda_q$;
 - else, break the loop, obtaining the string Λ and the index i_{aug} of the chosen block.
 13. Tell the leakage oracle that there are no more queries and obtain the shares $s_{\mathcal{B}_{i_{\text{aug}}}}$.
 14. Write, for all $j \in \mathcal{B}_{i_{\text{aug}}}$, $s_j^* = (s_j, r_j, (c_\beta)_{\beta \neq j}, (\pi_\beta^j)_{\beta \neq j})$
 15. Run $(\alpha_2, i_{\text{aug}}) \leftarrow \mathbf{A}_2^{\hat{\mathbf{A}}}(\alpha_1; r_2^{\mathbf{A}})$, answering as follows.
 - Upon receiving the q -th leakage query, write $\Lambda = (\Lambda_1, \dots, \Lambda_q, \dots)$ and return Λ_q .
 - Upon receiving the q -th tampering query $(\mathcal{T}, f_1, \dots, f_\beta)$: if $q \geq p^*$, return \perp ; else, compute $\tilde{s}_{\mathcal{B}_{i_{\text{aug}}}}^* = f_{i_{\text{aug}}}(s_{\mathcal{B}_{i_{\text{aug}}}^*})$ and $(\tilde{m}, \tilde{c}) \leftarrow \hat{f}_{i_{\text{aug}}}(\text{state}, \mathcal{T}, f_{i_{\text{aug}}}, s_{\mathcal{B}_{i_{\text{aug}}}})$ and return \tilde{m} .
 16. Run $b' \leftarrow \mathbf{A}_3(\alpha_2, s_{\mathcal{B}_{i_{\text{aug}}}}; r_3^{\mathbf{A}})$.
 17. Return b' .

For the analysis, we must show that $\hat{\mathbf{A}}$ is ℓ -admissible and that the reduction is correct. $\hat{\mathbf{A}}$ makes leakage queries in steps 5, 8, 10 and 12. The leakage amount in step 5 is $\gamma = O(\lambda)$ bits per share. The leakage amount in step 8 is bounded¹¹ by the ℓ^* -admissibility of \mathbf{A} . In step 10, the reduction $\hat{\mathbf{A}}$ uses the binary search algorithm to find the first index p^* in which the mauled commitments or messages from two different blocks differ, thus making up to $\log(q_{\text{tamper}}) = O(\log(\lambda))$ queries; in each query, the information obtained amounts to $\gamma = O(\lambda)$ bits, therefore the overall leakage amount in step 10 is $O(\lambda \log(\lambda))$. The leakage amount in step 12 is bounded again by the ℓ^* -admissibility of \mathbf{A} .

Summing up the above, the overall leakage amount made by $\hat{\mathbf{A}}$ from the block \mathcal{B}_i is

$$|\mathcal{B}_i| \gamma + 2\ell^* + O(\lambda \log(\lambda)) \leq n\gamma + 2\ell^* + O(\lambda \log(\lambda)) = \ell,$$

therefore $\hat{\mathbf{A}}$ is ℓ -admissible.

Now it remains to show that the reduction $\hat{\mathbf{A}}$ behaves, with overwhelming probability, exactly like the challenger in the experiment $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}^2(\lambda, b)$. First of all, call **Abort** the event in which the computation of \hat{f}_i aborts. This event is similar to the event **Bad**₂ in proof to Lemma 7 and happens when \mathbf{A} manages to break the binding property of the underlying commitment scheme (Gen, Com), hence happening with negligible probability.

- In steps 1-6, the reduction $\hat{\mathbf{A}}$ determines the randomness of the algorithm \mathbf{A} and perfectly simulates the challenger. On the other side, generates the needed randomness, obtains the commitments of the shares from the leakage oracle and simulates the proofs.
- In step 7, the reduction gradually acquires the results of the leakage queries and the number of tamper queries made by \mathbf{A} . Since each \hat{g}_i^{sd} has only access to the set \mathcal{B}_i , at some point may exist $i, j \in [\beta]$ such that the result of a tampering query made by \mathbf{A} in \hat{g}_i^{sd}

¹¹We should also consider $O(\log(\lambda))$ bits given by the indices q_{tamper}^i . Since $O(\log(\lambda)) + O(\lambda \log(\lambda)) = O(\lambda \log(\lambda))$, they are irrelevant and thus omitted. The same situation appears in step 12.

differs from the result of the same tampering query in \hat{g}_j^{sd} , thus the next leakage queries may also differ. In particular, \mathbf{A} may query two different leakage functions for \mathcal{B}_i in \hat{g}_i^{sd} and \hat{g}_j^{sd} , resulting in an inconsistent answer. However, this is not a problem since the adversary receiving inconsistent answers is the one simulated in the leakage oracle and such inconsistent answers are discarded in the next steps; in particular, different answers from the same tampering query q should result in a self-destruct only, thus the tampering answers should remain the same before the query q and \perp (thus, again, the same) from q onwards. Finally, note that, since $q^* = 0$ in \hat{g}_i^{sd} , the condition $q_{\text{tamper}} = q^*$ never verifies, thus there's no need to specify a hash function $\psi \in \Psi$.

- In steps 8-9, the reduction uses the binary search to find the first query p^* in which either the mauled commitments or the resulting reconstructions differ. In particular, the reduction samples $\psi \leftarrow_s \Psi$ and the leakage oracle runs up to q_m tampering queries, collecting, for each $i \in [\beta]$, all the relative commitments and reconstructed messages in a string x_i ; then, the leakage oracle halts \mathbf{A} , computes the hash y_i of x_i through ψ and outputs y_i . Now the reduction compares all these hashes and halves the search range accordingly. Call **Bad** the event in which there exist two different $x_i \neq x_j$ that generate the same $y_i = y_j$.
- In steps 10-11, the reduction discards all the Λ_i and queries them again. In particular, now $\hat{\mathbf{A}}$ knows when \mathbf{A} receives two different answers to the same tampering query (unless **Bad** happens) and repeats the leakage acquisition similarly to the one performed in step 7; however, the new function \hat{g}_i^{leak} is programmed to answer normally to tamper queries until the p^* -th query, after which the answer is always \perp . If **Bad** does not happen, the algorithm \mathbf{A} asks the same leakage and tampering queries for all $i \in [\beta]$ and the results remain the same. At the end of the run, \hat{g}_i^{leak} returns the index i_{aug} of the block chosen by \mathbf{A} (that, again, is the same for all \hat{g}_i^{leak} unless **Bad** happens).
- In steps 12-13, the reduction has no more leakage queries to do and can ask for the shares in the block $\mathcal{B}_{i_{\text{aug}}}$.
- In step 14, the reduction $\hat{\mathbf{A}}$ finally runs \mathbf{A} , answering the leakage and tampering queries as in $\hat{g}_{i_{\text{aug}}}^{\text{leak}}$. Again, if **Bad** does not happen, this is a perfect simulation of the challenger: the answers for the leakage queries have been previously computed and the answers for the tampering queries remain the same for every block \mathcal{B}_i in which they are computed (so, in particular, in $\mathcal{B}_{i_{\text{aug}}}$).
- Finally, in steps 15-16, the reduction $\hat{\mathbf{A}}$ obtains the distinguishing bit b' from \mathbf{A} and outputs the same bit.

The simulation made by $\hat{\mathbf{A}}$ is perfect unless the event $\mathbf{Bad} \cup \mathbf{Abort}$ happens; in particular, if **Bad** does not happen, the commitments and the reconstructed messages are the same in all blocks until a self-destruct happens, therefore the check is equivalent to the one performed in the original experiment. The event **Bad** happens with negligible probability by weak universality of Ψ and the event **Abort** happens with negligible probability by computational binding of (Gen, Com). Therefore, if \mathbf{A} exists, the reduction $\hat{\mathbf{A}}$ breaks the leakage resilience of Σ by querying up to ℓ bits of leakage and telling apart if the scheme has been applied on m_0 or on m_1 with noticeable probability. This concludes the proof. \square

4.4 Concrete Instantiation

Finally, we show how to instantiate Thm. 4 from generic assumptions, thus yielding the statement of Thm. 2 as a corollary. It is well known that true-simulation extractable NIZKs can be obtained from (doubly-enhanced) trapdoor permutations [FLS90, SCO⁺01, DHLW10],

whereas statistically hiding non-interactive commitments—with commitment size $O(\lambda)$ and $2^{-\Omega(\lambda)}$ -statistical hiding—can be instantiated from collision-resistant hash functions [HM96].

As for the underlying leakage-resilient secret sharing, we can use the recent construction from [KMS18] which achieves information-theoretic security in the stronger setting where the attacker can adaptively leak from subsets of shares of size at most $O(\log n)$, in a joint manner. The latter clearly implies leakage resilience against selective $O(\log n)$ -joint leakage.

5 Construction in the Plain Model

5.1 Description of the Scheme

We show how to obtain leakage-resilient continuously non-malleable secret sharing for arbitrary access structures in the plain model, with security against independent leakage and tampering attacks. Our construction combines a non-interactive commitment scheme Com with an auxiliary n -party secret sharing scheme $\Sigma = (\text{Share}, \text{Rec})$, as depicted in Fig. 5. The basic idea is to compute a commitment c to the message m being shared, using random coins r ; hence, we secret share the string $m||r$ using the underlying sharing function Share , yielding shares (s_1, \dots, s_n) . Hence, the final share of the i -th player is $s_i^* = (c, s_i)$.

We establish the following result. Note that when $n = 2$, we get as a special case the construction of split-state continuously non-malleable codes in the plain model that was originally proposed in [OPVV18], and later simplified in [FV19] by relying on noisy leakage. Our proof can be seen as a generalization of the proof strategy in [FV19] to the case $n > 2$.

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

- (i) Com is a perfectly binding and computationally hiding non-interactive commitment;
- (ii) Σ is an n -party ℓ -noisy leakage-resilient one-time non-malleable secret sharing scheme realizing access structure \mathcal{A} against independent leakage and tampering in the plain model, 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. 5 is an n -party ℓ^* -noisy leakage-resilient continuously non-malleable secret sharing scheme realizing access structure \mathcal{A} against independent leakage and tampering with computational security in the plain model, as long as $\ell = \ell^* + 1 + \gamma + O(\log \lambda)$ where $\gamma = \log |\mathcal{C}|$ is the size of a commitment.*

5.2 Proof Overview

Before coming to the proof, we discuss the main intuition behind privacy and continuous non-malleability.

Privacy. In order to show privacy, we need to prove that no PPT attacker can distinguish the distribution of $\text{Share}^*(1^\lambda, m_0)_{\mathcal{U}}$ from $\text{Share}^*(1^\lambda, m_1)_{\mathcal{U}}$, for any choice of messages $m_0, m_1 \in \mathcal{M}$, as well as unauthorized subset $\mathcal{U} \notin \mathcal{A}$. Towards this, consider a modified sharing algorithm $\text{Share}'(1^\lambda, m)$ that is identical to $\text{Share}(1^\lambda, m)$, except that the shares (s_1, \dots, s_n) are now obtained by secret sharing (via Share) an independent (uniformly random) string $\hat{m}||\hat{r}$.

The proof of privacy then proceeds in two steps. In the first step, we show that for any message $m \in \mathcal{M}$ and unauthorized subset $\mathcal{U} \notin \mathcal{A}$ the distributions $\text{Share}^*(1^\lambda, m)_{\mathcal{U}}$ and $\text{Share}'(1^\lambda, m)_{\mathcal{U}}$ are computationally close; this follows by privacy of Σ , as we can easily turn a

Let Com 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 $c = \text{Com}(m; r)$ and $(s_1, \dots, s_n) \leftarrow_{\$} \text{Share}(m||r)$. Return the shares $s^* = (s_1^*, \dots, s_n^*)$, where for each $i \in [n]$ we set $s_i^* = (c, s_i)$.

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

Hence, proceed as follows:

- (a) 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 $s_i^* = (s_i, c)$.
- (b) Run $m||r = \text{Rec}((s_i)_{i \in \mathcal{I}})$; if the outcome equals \perp return \perp .
- (c) If $c = \text{Com}(m; r)$ return m , else return \perp .

Figure 5: Leakage-resilient continuously non-malleable secret sharing for arbitrary access structures against independent leakage and tampering in the plain model.

distinguisher D telling apart $\text{Share}^*(1^\lambda, m)_{\mathcal{U}}$ and $\text{Share}'(1^\lambda, m)_{\mathcal{U}}$ into a distinguisher \hat{D} telling apart $\text{Share}(1^\lambda, m||r)_{\mathcal{U}}$ and $\text{Share}(1^\lambda, \hat{m}||r)_{\mathcal{U}}$, for random r .¹²

In the second step, we show that for all messages $m_0, m_1 \in \mathcal{M}$, and for all unauthorized subset $\mathcal{U} \notin \mathcal{A}$, the distributions $\text{Share}'(1^\lambda, m_0)_{\mathcal{U}}$ and $\text{Share}'(1^\lambda, m_1)_{\mathcal{U}}$ are computationally close; this is because in Share' the input to the sharing algorithm Share is decoupled from the input to the commitment Com , which immediately yields a reduction from a distinguisher telling apart $\text{Share}'(1^\lambda, m_0)_{\mathcal{U}}$ and $\text{Share}'(1^\lambda, m_1)_{\mathcal{U}}$ to a distinguisher telling apart $\text{Com}(1^\lambda, m_0)$ and $\text{Com}(1^\lambda, m_1)$. The latter contradicts computational hiding of the commitment.

Non-malleability. The proof of non-malleability follows along the same lines of the proof in [OPVV18, FV19]. The goal is to show that no PPT attacker A can distinguish the experiments $\mathbf{Tamper}_{\Sigma^*, A}(\lambda, 0)$ and $\mathbf{Tamper}_{\Sigma^*, A}(\lambda, 1)$ (cf. Def. 11). Towards this, we consider a hybrid experiment $\mathbf{Hyb}_{\Sigma^*, A}(\lambda, b)$ which is obtained from $\mathbf{Tamper}_{\Sigma^*, A}(\lambda, b)$ by making two changes: (i) The target secret sharing s^* is computed using the modified algorithm Share' considered already in the proof of privacy; (ii) Whenever the message \hat{m} reconstructed inside the tampering oracle $\mathcal{O}_{\text{nmss}}$ equals the dummy message \hat{m} , we return \heartsuit as long as the mauled commitment $\tilde{c}_1 = \dots = \tilde{c}_n$ equals to the original commitment c , and otherwise we return \perp leading to self-destruct.¹³

The first step is to show that, for any b , the distributions $\mathbf{Tamper}_{\Sigma^*, A}(\lambda, b)$ and $\mathbf{Hyb}_{\Sigma^*, A}(\lambda, b)$ are statistically close. The proof is by induction on the number of tampering queries p asked by the attacker A to the target $\mathcal{O}_{\text{nmss}}$ oracle.

Induction basis. Assume we have an unbounded attacker A telling apart the experiments $\mathbf{Tamper}_{\Sigma^*, A}(\lambda, b)$ and $\mathbf{Hyb}_{\Sigma^*, A}(\lambda, b)$ using a single tampering query $(\mathcal{T}, (f_1, \dots, f_n))$ to the target $\mathcal{O}_{\text{nmss}}(s^*, \cdot, \cdot)$ oracle. We build an unbounded attacker \hat{A} whose goal is to distinguish $\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 0)$ and $\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 1)$. Roughly, the reduction works as follows. When A

¹²In particular, upon receiving $(\hat{s}_u)_{u \in \mathcal{U}}$, distinguisher \hat{D} simply appends $c = \text{Com}(m; r)$ to each share s_u , and finally runs D upon input the resulting set of shares.

¹³The second change is needed as otherwise an attacker could easily distinguish between the original experiment and the hybrid by tampering with the identity function.

outputs $m_0, m_1 \in \mathcal{M}$, we pick uniformly \hat{m}, r, \hat{r} and forward $\hat{m}_0 = m_b || r$ and $\hat{m}_1 = \hat{m} || \hat{r}$ to the challenger. Denote by s the target secret sharing. Note that the above choice of the messages \hat{m}_0, \hat{m}_1 yields a perfectly distributed copy of s^* inside the target leakage and tampering oracles $\mathcal{O}_{\text{leak}}(s, \cdot)$ and $\mathcal{O}_{\text{nmss}}(s, \cdot, \cdot)$, by hard-wiring to each query the value $c = \text{Com}(m_b; r)$. This allows to trivially simulate all the leakage queries asked by A .

In order to simulate A 's tampering query $(\mathcal{T}, (f_1, \dots, f_n))$, we proceed as follows. We first leak the mauled commitments $(\tilde{c}_t)_{t \in \mathcal{T}}$ via the target leakage oracle. Hence, we use the target tampering oracle to obtain the mauled reconstructed value $\tilde{m} || \tilde{r}$ corresponding to $\tilde{s}_{\mathcal{T}}^*$. Finally, we can use knowledge of $\tilde{m} || \tilde{r}$ and $(\tilde{c}_t)_{t \in \mathcal{T}}$ to simulate the output of A 's tampering query correctly, with all but a negligible¹⁴ probability.

Inductive step. Assume now that the real experiment and the hybrid are indistinguishable for any unbounded attacker asking at most p tampering queries, but there exists an unbounded adversary A that can tell them apart using $p+1$ tampering queries. We construct an unbounded attacker \hat{A} attacking leakage-resilient one-time non-malleability of Σ as in the proof for the induction basis. The main idea is to simulate the answer to A 's first p tampering queries via limited leakage; assuming this can be done, we can then use A 's last tampering query in order to define \hat{A} 's unique tampering query exactly as in the previous case.

In particular, for each tampering query $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_n^{(q)}))$ asked by A , the strategy of the reduction is to leak the mauled commitments $(\tilde{c}_t)_{t \in \mathcal{T}^{(q)}}$. Hence, if the commitments are all equal to a single value \tilde{c} , attacker \hat{A} finds by brute force the corresponding¹⁵ opening \tilde{m} and uses this value to answer A 's query, and otherwise it sets $\tilde{m} = \perp$ with consequent self-destruct. To show that this strategy is sound, we need to overcome two challenges:

1. **(Adjusting the final guess.)** Note that by answering the first p tampering queries by inverting the mauled commitment \tilde{c} (when this is possible and such value is uniquely determined), the reduction neglects the possibility that the answer to one of tampering queries could be \perp (with consequent self-destruct) due to the fact that the inner mauled secret sharing \tilde{s} encodes an invalid message.

This might cause an inconsistency in the simulation. However, as originally shown in [OPVV18], the latter problem can be overcome (in the information-theoretic setting) by having the reduction asking an additional leakage query indicating whether the simulation of the first p tampering queries was correct or not. Since by the induction hypothesis the event that A does not provoke a self-destruct within the first p tampering queries must be noticeable, the extra leakage allows to maintain simulation and in particular preserve the advantage of A in the reduction.

2. **(Bounding the entropy loss.)** In order to conclude the proof, we still have to argue that the reduction did not leak too much information. In fact, note that \hat{A} for each tampering query leaks from the i -th share the mauled commitment \tilde{c}_i . Since the number of tampering queries p is an arbitrary polynomial, we cannot simply bound the entropy loss on each share by the length of a commitment, as this would result in too much leakage. Instead, we borrow a trick from [FNSV18] which allows to bound the entropy loss on each share by the size of a single commitment plus $O(\log \lambda)$. Intuitively, this holds because as long as the commitments \tilde{c}_i leaked from each share are all equal, we can interpret the leakage on the s_i as a function of the other shares $(s_j)_{j \neq i}$. On the other hand, if the

¹⁴When $\tilde{m} = \hat{m}$, the reduction returns \heartsuit as long as $\tilde{c} = (\tilde{c}_t)_{t \in \mathcal{T}} = c$; this is a perfect simulation as long as the real experiment never outputs \hat{m} , which however happens with overwhelming probability if the message length is super-logarithmic in the security parameter.

¹⁵Note that this value is unique, by perfect binding.

mauled commitments are not all equal, the leakage query does reveal information, but that accounts to the size of a single commitment and moreover happens at most once, since the reduction self-destructs afterwards.

Concluding the proof. Finally, we show that $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, 0)$ and $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, 1)$ are computationally close. Here, we rely on the computational hiding property of the commitment.

In particular, given a PPT attacker \mathbf{A} telling apart the experiments $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, 0)$ and $\mathbf{Hyb}_{\Sigma^*, \mathcal{A}}(\lambda, 1)$ we construct a PPT distinguisher $\hat{\mathbf{D}}$ distinguishing $\mathbf{Com}(1^\lambda, m_0)$ from $\mathbf{Com}(1^\lambda, m_1)$. This is straightforward because in the hybrid experiment the input to the inner secret sharing algorithm \mathbf{Share} is detached from the input to the commitment. Thus, the reduction can compute $s = (s_1, \dots, s_n)$ as a secret sharing of a random and independent string $\hat{m} \parallel \hat{r}$, and use the challenge commitment \hat{c} to answer all of \mathbf{A} 's leakage and tampering queries locally.

5.3 Security Analysis

Correctness w.r.t. access structure \mathcal{A} is immediate. We thus need to prove that Σ^* satisfies both privacy and continuous non-malleability.

Privacy. We need to prove that for all pairs of messages $m_0, m_1 \in \mathcal{M}$, and for all unqualified subsets $\mathcal{U} \notin \mathcal{A}$, we have that

$$\{(\mathbf{Share}(1^\lambda, m_0))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}} \approx_c \{(\mathbf{Share}(1^\lambda, m_1))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}}.$$

Consider a modified sharing algorithm \mathbf{Share}' where we let $c = \mathbf{Com}(m, r)$, for $r \leftarrow \mathcal{R}$, as in the original scheme, but (s_1, \dots, s_n) are now defined by secret sharing an independent random string $\hat{m} \parallel \hat{r} \leftarrow \mathcal{M} \times \mathcal{R}$; the final shares are set to $s'_i = (c, s_i)$ for all $i \in [n]$.

Lemma 9. $\forall m \in \mathcal{M}, \forall \mathcal{U} \notin \mathcal{A} : \{(\mathbf{Share}^*(1^\lambda, m))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}} \approx_c \{(\mathbf{Share}'(1^\lambda, m))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}}.$

Proof. The proof is down to the privacy property of the underlying secret sharing scheme Σ . By contradiction, assume that there exists a PPT distinguisher \mathbf{D} , a message $m \in \mathcal{M}$, and an unauthorized subset $\mathcal{U} \notin \mathcal{A}$, such that

$$\left| \mathbb{P} \left[\mathbf{D}((\mathbf{Share}^*(1^\lambda, m))_{\mathcal{U}}) = 1 \right] - \mathbb{P} \left[\mathbf{D}((\mathbf{Share}'(1^\lambda, m))_{\mathcal{U}}) = 1 \right] \right| \geq \frac{1}{\text{poly}(\lambda)}.$$

Consider the following PPT distinguisher $\hat{\mathbf{D}}$ attacking privacy of Σ .

- Output messages $\hat{m}_0 = m \parallel r$ and $\hat{m}_1 = \hat{m} \parallel \hat{r}$, where $\hat{m} \leftarrow \mathcal{M}$ and $r, \hat{r} \leftarrow \mathcal{R}$.
- Upon receiving a set of shares $(\hat{s}_u)_{u \in \mathcal{U}}$, compute $c = \mathbf{Com}(m; r)$ and set $\hat{s}_u^* = (c, \hat{s}_u)$ for each $u \in \mathcal{U}$.
- Return the same as $\mathbf{D}((\hat{s}_u^*)_{u \in \mathcal{U}})$.

For the analysis, note that in case the values $(\hat{s}_u)_{u \in \mathcal{U}}$ are distributed like a secret sharing of \hat{m}_0 , the distribution of $(\hat{s}_1^*, \dots, \hat{s}_n^*)$ is identical to that of $\mathbf{Share}^*(1^\lambda, m)$, whereas in case the values $(\hat{s}_u)_{u \in \mathcal{U}}$ are distributed like a secret sharing of \hat{m}_1 , the distribution of $(\hat{s}_1^*, \dots, \hat{s}_n^*)$ is identical to that of $\mathbf{Share}'(1^\lambda, m)$. The lemma follows. \square

Lemma 10. $\forall m_0, m_1 \in \mathcal{M}, \forall \mathcal{U} \notin \mathcal{A} : \{(\mathbf{Share}'(1^\lambda, m_0))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}} \approx_c \{(\mathbf{Share}'(1^\lambda, m_1))_{\mathcal{U}}\}_{\lambda \in \mathbb{N}}.$

<p>Tamper$_{\Sigma^*, \mathcal{A}}(\lambda, b)$ Hyb$_{\Sigma^*, \mathcal{A}}(\lambda, b)$:</p> <p>$(m_0, m_1, \alpha_1) \leftarrow_{\\$} A_1(\omega)$ $r \leftarrow_{\\$} \mathcal{R}; c = \text{Com}(m_b; r)$ $(s_1, \dots, s_n) \leftarrow_{\\$} \text{Share}(m_b r)$ $\hat{m} \hat{r} \leftarrow_{\\$} \mathcal{M} \times \mathcal{R}; (s_1, \dots, s_n) \leftarrow_{\\$} \text{Share}(\hat{m} \hat{r})$ $(s_1^*, \dots, s_n^*) = ((c, s_1), \dots, (c, s_n))$ stop \leftarrow false Return $\alpha_2 \leftarrow_{\\$} A_2^{\mathcal{O}_{\text{nmss}}(s, \cdot, \cdot), \mathcal{O}_{\text{leak}}(s, \cdot)}(\alpha_1)$ Oracle $\mathcal{O}_{\text{leak}}(s, (g_1, \dots, g_n))$: Return $g_1(s_1), \dots, g_n(s_n)$</p>	<p>Oracle $\mathcal{O}_{\text{nmss}}(s, \mathcal{T}, (f_1, \dots, f_n))$: If stop = true Return \perp Else $\forall i \in [n] : \tilde{s}_i^* := f_i(s_{i^*}) = (\tilde{c}_i, \tilde{s}_i)$ $\tilde{m} \tilde{r} = \text{Rec}(\tilde{s}_{\mathcal{T}})$ If $\exists t_1, t_2 \in \mathcal{T}$ s.t. $\tilde{c}_{t_1} \neq \tilde{c}_{t_2}$ Return \perp and stop \leftarrow true If $\tilde{m} = \perp$ or $\tilde{c}_1 \neq \text{Com}(\tilde{m}; \tilde{r})$ Return \perp and stop \leftarrow true If $\tilde{m} \in \{m_0, m_1\}$ Return \heartsuit If $\tilde{m} = \hat{m}$ If $\tilde{c}_1 = c$, return \heartsuit Else Return \perp stop \leftarrow true Else, return \tilde{m}</p>
---	---

Figure 6: Hybrid experiments in the proof of Theorem 5.

Proof. The proof is down to the computational hiding property of the non-interactive commitment scheme Com . By contradiction, assume that there exists a PPT distinguisher D , a pair of messages $m_0, m_1 \in \mathcal{M}$, and an unauthorized subset $\mathcal{U} \notin \mathcal{A}$, such that

$$\left| \mathbb{P} \left[D((\text{Share}'(1^\lambda, m_0))_{\mathcal{U}}) = 1 \right] - \mathbb{P} \left[D((\text{Share}'(1^\lambda, m_1))_{\mathcal{U}}) = 1 \right] \right| \geq \frac{1}{\text{poly}(\lambda)}.$$

Consider the following PPT distinguisher \hat{D} , attacking the hiding property of Com .

- Upon receiving a target commitment \hat{c} that is either a commitment to m_0 or a commitment to m_1 , sample $\hat{m} \leftarrow_{\$} \mathcal{M}$ and $\hat{r} \leftarrow_{\$} \mathcal{R}$, and let $(s_1, \dots, s_n) \leftarrow_{\$} \text{Share}(\hat{m} || \hat{r})$.
- Define $\hat{s}_i^* = (\hat{c}, s_i)$ for each $i \in [n]$.
- Return the same as $D((\hat{s}_u^*)_{u \in \mathcal{U}})$.

For the analysis, note that in case the value \hat{c} is distributed like a commitment to m_0 , the distribution of $(\hat{s}_1^*, \dots, \hat{s}_n^*)$ is identical to that of $\text{Share}'(1^\lambda, m_0)$, whereas in case the value \hat{c} is distributed like a commitment to m_1 , the distribution of $(\hat{s}_1^*, \dots, \hat{s}_n^*)$ is identical to that of $\text{Share}'(1^\lambda, m_1)$. The lemma follows. \square

Fix arbitrary messages $m_0, m_1 \in \mathcal{M}$ and $\mathcal{U} \notin \mathcal{A}$. Combining the two lemmas above, we have obtained:

$$\begin{aligned} \left\{ (\text{Share}^*(1^\lambda, m_0))_{\mathcal{U}} \right\}_{\lambda \in \mathbb{N}} &\approx_c \left\{ (\text{Share}'(1^\lambda, m_0))_{\mathcal{U}} \right\}_{\lambda \in \mathbb{N}} \\ &\approx_c \left\{ (\text{Share}'(1^\lambda, m_1))_{\mathcal{U}} \right\}_{\lambda \in \mathbb{N}} \approx_c \left\{ (\text{Share}^*(1^\lambda, m_1))_{\mathcal{U}} \right\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

Continuous non-malleability. Let $\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, b)$ be the original experiment defining continuous non-malleability of Σ^* . Consider a modified experiment $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, b)$, where we replace (s_1, \dots, s_n) with a secret sharing of a random and independent value $\hat{m} \parallel \hat{r} \leftarrow \mathcal{M} \times \mathcal{R}$; when answering tampering queries, the experiment returns \heartsuit in case the mauled message \tilde{m} happens to be equal to \hat{m} . See Fig. 6 for a more formal description.

We first prove that the above experiments are computationally close by induction over the number of tampering queries $p \in \text{poly}(\lambda)$ asked by adversary \mathbf{A} . Towards this, let us denote by $\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, p, b)$ (resp. $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, p, b)$) the original (resp. hybrid) experiment where adversary \mathbf{A} is limited to ask exactly p queries to oracle $\mathcal{O}_{\text{nmss}}$. The lemma below constitutes the basis of the induction.

Lemma 11. $\forall b \in \{0, 1\} : \{\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, 1, b)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 1, b)\}_{\lambda \in \mathbb{N}}$.

Proof. The proof is down to the statistical leakage-resilient one-time non-malleability of Σ . Fix $b = 0$ (the proof for the other case being identical). Assume that there exists an unbounded adversary $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ which can distinguish $\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, 1, 0)$ and $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 1, 0)$ with non-negligible probability. By an averaging argument, this means that there must exist values $r \in \mathcal{R}$ and $\hat{m} \parallel \hat{r} \in \mathcal{M} \times \mathcal{R}$ such that \mathbf{A} distinguishes the two experiments when we fix these particular values of r and $\hat{m} \parallel \hat{r}$. Let $s = (s_1, \dots, s_n)$ be the target encoding in the tampering experiment relative to $(\text{Share}, \text{Rec})$. Consider the following adversary $\hat{\mathbf{A}}$ attacking Σ , given as input the above fixed values r, \hat{m}, \hat{r} .

1. Run $\mathbf{A}_1(1^\lambda)$, obtaining a pair of messages (m_0, m_1) and auxiliary state α_1 ; send $\hat{m}_0 = m_0 \parallel r$ and $\hat{m}_1 = \hat{m} \parallel \hat{r}$ to the challenger, and compute $c := \text{Com}(m_0; r)$. Run $\mathbf{A}_2(\alpha_1)$.
2. For each leakage query (g_1, \dots, g_n) asked by \mathbf{A}_2 , define the leakage function \hat{g}_i that hard-wires (a description of) g_i, c , and returns the same as $g_i(c, s_i)$. Forward $(\hat{g}_1, \dots, \hat{g}_n)$ to the target leakage oracle.
3. Upon input the only tampering query $(\mathcal{T}, (f_1, \dots, f_n))$ from \mathbf{A}_2 , proceed as follows.
 - (a) For each $t \in \mathcal{T}$, define the leakage function \hat{h}_t that hard-wires (a description of) f_t , and returns the commitment \tilde{c}_t such that $f_t(c, s_t) = (\tilde{c}_t, \tilde{s}_t)$.
 - (b) Forward $(\hat{h}_t)_{t \in \mathcal{T}}$ to the target leakage oracle,¹⁶ obtaining values $(\tilde{c}_t)_{t \in \mathcal{T}}$.
 - (c) Define the tampering function \hat{f}_i that hard-wires c and (a description of) f_i , and, upon input s_i , returns the value \tilde{s}_i specified by $f_i(c, s_i) = (\tilde{c}_i, \tilde{s}_i)$.
 - (d) Forward $(\mathcal{T}, (\hat{f}_1, \dots, \hat{f}_n))$ to the target tampering oracle, obtaining $\tilde{m} \parallel \tilde{r} \in \mathcal{M} \times \mathcal{R} \cup \{\perp, \heartsuit\}$. Hence:
 - If there exist $t_1, t_2 \in \mathcal{T}$ such that $\tilde{c}_{t_1} \neq \tilde{c}_{t_2}$, return \perp to \mathbf{A} ; else, let \tilde{c} be the unique commitment returned by the leakage oracle.
 - If $\tilde{m} \parallel \tilde{r} = \perp$ or is not a valid opening of \tilde{c} , return \perp to \mathbf{A} .
 - If $\tilde{m} \parallel \tilde{r} = \heartsuit$, in case $\tilde{c} = c$ return \heartsuit to \mathbf{A} (and otherwise \perp).
 - If $\tilde{m} \in \{m_0, m_1\}$, return \heartsuit to \mathbf{A} . Else, if $\tilde{m} = \hat{m}$ return \heartsuit to \mathbf{A} in case $\tilde{c} = c$, and otherwise return \perp to \mathbf{A} .
 - Else, return \tilde{m} to \mathbf{A} .
4. Output the same guess as that of \mathbf{A}_2 .

For the analysis, we next prove that the simulation performed by the above reduction is perfect with overwhelming probability. First, depending on the target (s_1, \dots, s_n) being either a secret sharing of \hat{m}_0 or of \hat{m}_1 , for every $i \in [n]$ the input to the tampering function f_i (resp. leakage function g_i) emulated inside \hat{f}_i (resp. \hat{g}_i) is identically distributed to the i -th share of the target

¹⁶More precisely, the leakage oracle takes as input n functions (one for each of the shares), but we can simply assume that \hat{h}_i equals the empty string ε for each $i \notin \mathcal{T}$.

secret sharing in either experiment $\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, 0, 1)$ or $\mathbf{Hyb}(\lambda, 0, 1)$, with our fixed choice of r, \hat{m}, \hat{r} . Second, the answer to \mathbf{A}_2 's tampering query is simulated correctly, with all but a negligible probability. Indeed:

- If $\text{Rec}(\tilde{s}_{\mathcal{T}})$ yields \perp , both the real and the hybrid experiment would return \perp , which is perfectly emulated by the reduction.
- If $\text{Rec}(\tilde{s}_{\mathcal{T}})$ yields \heartsuit , it means that the inner secret sharing reconstructs to either $\hat{m}_0 = m_0 || r$ or to $\hat{m}_1 = \hat{m} || \hat{r}$. Without loss of generality, assume further that the commitments in the tampered shares are all equal to a single value \tilde{c} .¹⁷ There are 4 possible cases: either both experiments output \hat{m}_0 , or both experiments output \hat{m}_1 , or one experiment outputs \hat{m}_0 while the other outputs \hat{m}_1 . However, since the view in the real experiment is independent of the value \hat{m} , except with negligible probability $2^{-\omega(\log(\lambda))}$, we can condition on the event that the real experiment does not output this value. Thus, there are only two cases to consider:
 - (i) Both the real and the hybrid experiment return $\hat{m}_0 = m_0 || r$.
 - (ii) The real experiment returns $\hat{m}_0 = m_0 || r$, whereas the hybrid returns $\hat{m}_1 = \hat{m} || \hat{r}$.

In both cases, the output of the two experiments is equal to \heartsuit in case $\tilde{c} = c$, and otherwise both experiments return \perp . This is exactly what the reduction does. Hence, the simulation is perfect except with negligible probability.

- If $\text{Rec}(\tilde{s}_{\mathcal{T}})$ yields some value $\tilde{m} || \tilde{r} \notin \{\heartsuit, \perp\}$, it means in particular that $\tilde{m} || \tilde{r} \notin \{\hat{m}_0, \hat{m}_1\}$. In such a case both experiments return \perp in case the modified commitment \tilde{c} does not match the opening \tilde{m}, \tilde{r} . Otherwise, it means that the modified shares produced by \mathbf{A} lead to a valid message $\tilde{m} \in \mathcal{M}$. Thus, the output of both experiments would either be \heartsuit or \tilde{m} (depending on \tilde{m} being equal to one of the two messages m_0, m_1 or not).

Finally, note that as long as \mathbf{A} is ℓ^* -admissible, $\hat{\mathbf{A}}$ is $\hat{\ell}$ -admissible for $\hat{\ell} = \ell^* + \gamma$, and since $\hat{\ell} < \ell$, attacker $\hat{\mathbf{A}}$ is ℓ -admissible. Thus, we can conclude that the distinguishing advantage of $\hat{\mathbf{A}}$ is the same as that of \mathbf{A} which concludes the proof of the lemma. \square

The lemma below constitutes the inductive step.

Lemma 12. *Fix any $p \in \text{poly}(\lambda)$. Assume that for all $b \in \{0, 1\}$,*

$$\{\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, p, b)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, p, b)\}_{\lambda \in \mathbb{N}}.$$

Then: $\{\mathbf{Tamper}_{\Sigma^, \mathbf{A}}(\lambda, p+1, b)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, p+1, b)\}_{\lambda \in \mathbb{N}}$, for all $b \in \{0, 1\}$.*

Proof. The proof is down to the statistical leakage-resilient one-time non-malleability of Σ . Fix $b = 0$ (the proof for the other case being identical). Assume that there exists an unbounded adversary $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ which can distinguish $\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, p+1, 0)$ and $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, p+1, 0)$ with non-negligible probability. By an averaging argument, this means that there must exist values $r \in \mathcal{R}$ and $\hat{m} || \hat{r} \in \mathcal{M} \times \mathcal{R}$ such that \mathbf{A} distinguishes the two experiments when we fix these particular values of r and $\hat{m} || \hat{r}$. Let $s = (s_1, \dots, s_n)$ be the target encoding in the tampering experiment relative to $(\text{Share}, \text{Rec})$. Consider the following adversary $\hat{\mathbf{A}}$ attacking Σ , given as input the above fixed values r, \hat{m}, \hat{r} .

1. Run $\mathbf{A}_1(1^\lambda)$, obtaining a pair of messages (m_0, m_1) and auxiliary state α_1 ; send $\hat{m}_0 = m_0 || r$ and $\hat{m}_1 = \hat{m} || \hat{r}$ to the challenger, and compute $c := \text{Com}(m_0; r)$. Run $\mathbf{A}_2(\alpha_1)$.

¹⁷In fact, if this is not the case, both experiments return \perp , which is once again perfectly emulated by the reduction.

2. The reduction keeps a list of all possible secret sharing of the messages \hat{m}_0, \hat{m}_1 ; let $\hat{\mathcal{S}}_1^{(1)}, \dots, \hat{\mathcal{S}}_n^{(1)}$ be the initial sets.
3. For each leakage query (g_1, \dots, g_n) asked by A_2 , define the leakage function \hat{g}_i that hard-wires (a description of) g_i, c , and returns the same as $g_i(c, s_i)$. Forward $(\hat{g}_1, \dots, \hat{g}_n)$ to the target leakage oracle.
4. For each $q \in [p]$, upon input the q -th tampering query $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_n^{(q)}))$, proceed as follows.
 - (a) For each $t \in \mathcal{T}^{(q)}$, define the leakage function $\hat{h}_t^{(q)}$ that hard-wires (a description of) $f_t^{(q)}$, and returns the commitment $\tilde{c}_t^{(q)}$ such that $f_t^{(q)}(c, s_t) = (\tilde{c}_t^{(q)}, \tilde{s}_t)$.
 - (b) Forward $(\hat{h}_t^{(q)})_{t \in \mathcal{T}^{(q)}}$ to the target leakage oracle,¹⁸ obtaining values $(\tilde{c}_t^{(q)})_{t \in \mathcal{T}^{(q)}}$. Hence:
 - If there exist $t_1, t_2 \in \mathcal{T}^{(q)}$ such that $\tilde{c}_{t_1}^{(q)} \neq \tilde{c}_{t_2}^{(q)}$, set $\tilde{m}^{(q)} = \perp$ and self-destruct; else, let $\tilde{c}^{(q)}$ be the unique commitment returned by the leakage oracle.
 - Find by brute force the opening $\tilde{m}^{(q)}$ (i.e., $\tilde{c}^{(q)} = \text{Com}(\tilde{m}^{(q)}; \tilde{r}^{(q)})$ for some $\tilde{r}^{(q)} \in \mathcal{R}$); if no such value is found, set $\tilde{m}^{(q)} = \perp$ and self-destruct.
 - If $\tilde{m}^{(q)} \in \{m_0, m_1\}$, re-define $\tilde{m}^{(q)} := \heartsuit$. Else, if $\tilde{m}^{(q)} = \hat{m}$ re-define $\tilde{m}^{(q)} := \heartsuit$ in case $\tilde{c}^{(q)} = c$, and otherwise $\tilde{m}^{(q)} := \perp$ and self-destruct.
 - Return $\tilde{m}^{(q)}$ to A .
 - For each $i \in [n]$, define $\hat{\mathcal{S}}_i^{(q+1)} \subseteq \hat{\mathcal{S}}_i^{(q)}$ the set of shares for the i -th player that are compatible with the answer to the q -th tampering query being $\tilde{m}^{(q)}$, i.e. $(\hat{s}_1, \dots, \hat{s}_n) \in \hat{\mathcal{S}}_1^{(q+1)} \times \dots \times \hat{\mathcal{S}}_n^{(q+1)}$ if and only if $\tilde{m}^{(q)}$ is the output corresponding to $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_n^{(q)}))$ in the hybrid experiment.
5. Upon input the last tampering query $(\mathcal{T}^{(p+1)}, (f_1^{(p+1)}, \dots, f_n^{(p+1)}))$, proceed as follows.
 - (a) Define the same functions $\hat{h}_t^{(p+1)}$ considered in step 4a, and forward $(\hat{h}_t^{(p+1)})_{t \in \mathcal{T}^{(p+1)}}$ to the target leakage oracle, obtaining values $(\tilde{c}_t^{(p+1)})_{t \in \mathcal{T}^{(p+1)}}$.
 - (b) Define the tampering function \hat{f}_i that hard-wires c and (a description of) $f_i^{(p+1)}$, and, upon input s_i , returns the value $\tilde{s}_i^{(p+1)}$ specified by $f_i^{(p+1)}(c, s_i) = (\tilde{c}_i^{(p+1)}, \tilde{s}_i^{(p+1)})$.
 - (c) Forward $(\mathcal{T}^{(p+1)}, (\hat{f}_1, \dots, \hat{f}_n))$ to the target tampering oracle, obtaining $\tilde{m}^{(p+1)} \parallel \tilde{r}^{(p+1)} \in \mathcal{M} \times \mathcal{R} \cup \{\perp, \heartsuit\}$. Hence:
 - If there exist $t_1, t_2 \in \mathcal{T}^{(p+1)}$ such that $\tilde{c}_{t_1}^{(p+1)} \neq \tilde{c}_{t_2}^{(p+1)}$, return \perp to A ; else, let $\tilde{c}^{(p+1)}$ be the unique commitment returned by the leakage oracle.
 - If $\tilde{m}^{(p+1)} \parallel \tilde{r}^{(p+1)} = \perp$ or is not a valid opening of $\tilde{c}^{(p+1)}$, return \perp to A .
 - If $\tilde{m}^{(p+1)} \parallel \tilde{r}^{(p+1)} = \heartsuit$, in case $\tilde{c}^{(p+1)} = c$ return \heartsuit to A (and otherwise \perp).
 - If $\tilde{m}^{(p+1)} \in \{m_0, m_1\}$, return \heartsuit to A . Else, if $\tilde{m}^{(p+1)} = \hat{m}$ return \heartsuit to A in case $\tilde{c}^{(p+1)} = c$, and otherwise return \perp to A .
 - Else, return $\tilde{m}^{(p+1)}$ to A .
6. Check that the simulation up to the first p queries did not cause any inconsistency, due to the fact that the outcome of the q -th tampering queries should have been \perp because $(\tilde{s}_t)_{t \in \mathcal{T}^{(q)}}$ was not a valid secret sharing.
 - (a) Without loss of generality, assume that the output of A_2 is equal to 0 whenever A believes that the target secret sharing is distributed as in the real experiment.
 - (b) Define the following special leakage function $\hat{h}_{\text{check}} : \mathcal{S}_1 \rightarrow \{0, 1\}$.
 - The function hard-wires a description of A_2 , the values (c, m_0, m_1) , a description of the final tampering query $(\mathcal{T}^{(p+1)}, (f_1^{(p+1)}, \dots, f_n^{(p+1)}))$, the answer to previous

¹⁸More precisely, the leakage oracle takes as input n functions (one for each of the shares), but we can simply assume that $\hat{h}_i^{(q)}$ equals the empty string ε for each $i \notin \mathcal{T}^{(q)}$.

tampering queries $(\tilde{m}^{(1)}, \dots, \tilde{m}^{(p)})$, the answer to all leakage queries on each of the shares $\Lambda := (\Lambda_1, \dots, \Lambda_n)$, and the sets $\hat{\mathcal{S}}_2^{(p+1)}, \dots, \hat{\mathcal{S}}_n^{(p+1)}$.

- Let $\hat{s}^* = ((c, s_1), (c, \hat{s}_2), \dots, (c, \hat{s}_n))$ be the target secret sharing for each possible remaining set of compatible shares $(\hat{s}_2, \dots, \hat{s}_n) \in \hat{\mathcal{S}}_2^{(p+1)} \times \dots \times \hat{\mathcal{S}}_n^{(p+1)}$.
- The output of the function is a bit \tilde{b} such that $\tilde{b} = 1$ if and only if $A_2(\tilde{m}^{(1)}, \dots, \tilde{m}^{(p)}, \tilde{m}^*, \Lambda) = 0$ more often when \hat{s}^* is a valid secret sharing of message m_0 , where \tilde{m}^* is the output of the $\mathcal{O}_{\text{nmss}}$ oracle in the hybrid experiment upon input $(\mathcal{T}^{(p+1)}, (f_1^{(p+1)}, \dots, f_n^{(p+1)}))$ with target secret sharing \hat{s}^* .

(c) Forward $(h_{\text{check}}, \varepsilon, \dots, \varepsilon)$ to the target leakage oracle, obtaining a bit \tilde{b} .

7. Upon receiving a bit b' from A_2 , in case $\tilde{b} = 1$ output b' , and else return a random guess.

Attacker \hat{A} runs in exponential time. We now show that its distinguishing advantage is negligibly close to that of A . Indeed:

$$\begin{aligned} & \left| \mathbb{P} \left[\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 1, 0) = 1 \right] - \mathbb{P} \left[\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 1, 1) = 1 \right] \right| \\ &= \left| \mathbb{P} \left[\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 1, 0) = 1 \wedge \tilde{b} = 1 \right] - \mathbb{P} \left[\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 1, 1) = 1 \wedge \tilde{b} = 1 \right] \right| \end{aligned} \quad (1)$$

$$\geq \frac{1}{\text{poly}(\lambda)} \cdot \left| \mathbb{P} \left[\mathbf{Tamper}_{\Sigma, \hat{A}}(\lambda, 1, 0) = 1 \mid \tilde{b} = 1 \right] - \mathbb{P} \left[\mathbf{Hyb}_{\Sigma, \hat{A}}(\lambda, 1, 1) = 1 \mid \tilde{b} = 1 \right] \right| \quad (2)$$

$$\geq \frac{1}{\text{poly}(\lambda)} \cdot \left(\frac{1}{\text{poly}(\lambda)} - \text{negl}(\lambda) \right), \quad (3)$$

where Eq. (1) follows because when $\tilde{b} = 0$, the reduction \hat{A} returns a random guess, and thus its distinguishing advantage is zero; Eq. (2) holds as the induction hypothesis implies that $\tilde{b} = 1$ with non-negligible probability, otherwise A generates an invalid secret sharing $(\tilde{s}_1^*, \dots, \tilde{s}_n^*)$ within the first p tampering queries with overwhelming probability, which in turn means that A can distinguish using less than $p + 1$ outputs from the decoding. Finally, Eq. (2) holds because an analysis identical to that of Lemma 11 shows that the view of A is perfectly simulated (except with negligible probability) conditioned on $\tilde{b} = 1$, and thus in this case \hat{A} retains essentially the same advantage as that of A .

In order to conclude the proof, it remains to show that \hat{A} is ℓ -admissible, for ℓ as in the statement of the theorem. Note that adversary \hat{A} makes leakage queries at steps 3, 4b, 5a, and 6c, but the leakage queries of step 6c is executed only once, and for a total of at most 1 bit of leakage. Let $q^* \in \mathbb{N}$ be the index of the tampering query, if any, where the commitments retrieved in step 4b happen to be different, and set $q^* = p + 1$ in case that never happens; note that q^* is a random variable, which we denote by \mathbf{q}^* . Clearly, the leakage queries of step 4b are executed exactly $\min\{q^*, p\}$ times. For each $i \in [n]$, denote by $\mathbf{\Lambda}_i$ the random variable corresponding to the leakage performed by the reduction on the i -th share of the target secret sharing $(\mathbf{S}_1, \dots, \mathbf{S}_n)$. We can write:

$$\tilde{\mathbb{H}}_{\infty}(\mathbf{S}_i | (\mathbf{S}_j)_{j \neq i}, \mathbf{\Lambda}_i) \geq \tilde{\mathbb{H}}_{\infty}(\mathbf{S}_i | (\mathbf{S}_j)_{j \neq i}, \hat{h}_i^{(1)}(\mathbf{S}_i), \dots, \hat{h}_i^{(\mathbf{q}^*)}(\mathbf{S}_i), \hat{h}_i^{(p+1)}(\mathbf{S}_i)) - \ell^* - 1 \quad (4)$$

$$= \tilde{\mathbb{H}}_{\infty}(\mathbf{S}_i | (\mathbf{S}_j)_{j \neq i}, \mathbf{q}^*, \hat{h}_i^{(\mathbf{q}^*)}(\mathbf{S}_i), \hat{h}_i^{(p+1)}(\mathbf{S}_i)) - \ell^* - 1 \quad (5)$$

$$\geq \tilde{\mathbb{H}}_{\infty}(\mathbf{S}_i | (\mathbf{S}_j)_{j \neq i}) - \ell^* - 1 - \gamma - O(\log \lambda) \quad (6)$$

In the above derivation, Eq. (4) follows by definition of $\mathbf{\Lambda}_i$, and because the adversary A is ℓ^* -admissible and furthermore the leakage performed in step 6c consists of at most 1 bit; Eq. (5) follows by the fact that, for each $q < q^*$, the commitments leaked in step 5a are all identical and thus can be computed as a function of $(\mathbf{S}_j)_{j \neq i}$ and \mathbf{q}^* ; Eq. (6) follows because either $\mathbf{q}^* = p + 1$,

and then the min-entropy drop due to the leakage $\hat{h}_i^{(\mathbf{q}^*)}(\mathbf{S}_i), \hat{h}_i^{(p+1)}(\mathbf{S}_i)$ is bounded by the size γ of a commitment, or $\mathbf{q}^* < p + 1$, in which case the adversary \hat{A} self-destructs and only the value $\hat{h}_i^{(\mathbf{q}^*)}(\mathbf{S}_i)$ is leaked (causing a drop of at most γ in the min-entropy bound). The lemma follows. \square

Combining Lemma 11 and Lemma 12, we get that for all $b \in \{0, 1\}$:

$$\{\mathbf{Tamper}_{\Sigma^*, \mathbf{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}}.$$

The lemma below concludes the proof.

Lemma 13. $\{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 1)\}_{\lambda \in \mathbb{N}}$.

Proof. Assume that there exists a PPT adversary $\mathbf{A} = (\mathbf{A}_1, \mathbf{A}_2)$ telling apart $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 0)$ and $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 1)$ with non-negligible probability. We build a PPT distinguisher \mathbf{D} such that

$$\left| \mathbb{P} \left[\mathbf{D}(1^\lambda, \hat{c}) = 1 : \hat{c} := \mathbf{Com}(m_0; r); r \leftarrow \mathcal{R} \right] - \mathbb{P} \left[\mathbf{D}(1^\lambda, \hat{c}) = 1 : \hat{c} := \mathbf{Com}(m_1; r); r \leftarrow \mathcal{R} \right] \right| \geq 1/\text{poly}(\lambda),$$

where (m_0, m_1) are the messages output by $\mathbf{A}_1(1^\lambda)$. This contradicts the computational hiding property of the non-interactive commitment scheme, and thus concludes the proof of the lemma.

Distinguisher \mathbf{D} , with black-box access to \mathbf{A}_2 , and with input $\hat{c} \in \mathcal{C}$, is described below:

- Sample $(s_1, \dots, s_n) \leftarrow \mathcal{S} \text{Share}(\hat{m} || \hat{r})$, for $\hat{m} || \hat{r} \leftarrow \mathcal{M} \times \mathcal{R}$.
- Upon input a leakage query (g_1, \dots, g_n) from \mathbf{A}_2 , return $(g_1(\hat{c}, s_1), \dots, g_n(\hat{c}, s_n))$.
- Upon input the q -th tampering query $(\mathcal{T}^{(q)}, (f_1^{(q)}, \dots, f_n^{(q)}))$ from \mathbf{A}_2 , proceed as follows:
 - For each $i \in [n]$, compute

$$\tilde{s}_i^* := f_i^{(q)}(\hat{c}, s_i) = (\tilde{c}_i, \tilde{s}_i)$$

and let $\tilde{m}^{(q)} || \tilde{r}^{(q)} = \mathbf{Rec}(\tilde{s}_{\mathcal{T}^{(q)}})$ where $\tilde{s} = (\tilde{s}_1, \dots, \tilde{s}_n)$.

- If $\exists t_1, t_2 \in \mathcal{T}^{(q)}$ s.t. $\tilde{c}_{t_1} \neq \tilde{c}_{t_2}$, return \perp to \mathbf{A}_2 and self-destruct.
- If $\tilde{m}^{(q)} = \perp$ or $\tilde{c}_1 \neq \mathbf{Com}(\tilde{m}^{(q)}; \tilde{r}^{(q)})$, return \perp to \mathbf{A}_2 and self-destruct.
- If $\tilde{m}^{(q)} \in \{m_0, m_1\}$, return \heartsuit to \mathbf{A}_2 .
- If $\tilde{m}^{(q)} = \hat{m}$, return \heartsuit to \mathbf{A}_2 in case $\tilde{c}_1 = \hat{c}$ and otherwise return \perp to \mathbf{A}_2 and self-destruct.
- Else, return $\tilde{m}^{(q)}$ to \mathbf{A}_2 .
- Return the same guess as \mathbf{A}_2 .

For the analysis, note that the simulation done by \mathbf{D} is perfect. In particular, depending on the value \hat{c} being either a commitment to m_0 or a commitment to m_1 , the view of \mathbf{A}_2 is identical to the view in either experiment $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 0)$ or $\mathbf{Hyb}_{\Sigma^*, \mathbf{A}}(\lambda, 1)$. Thus, \mathbf{D} distinguishes with non-negligible probability. This finishes the proof. \square

6 Statistical One-Time Non-Malleability with Noisy Leakage

Since non-interactive, perfectly binding, commitments can be obtained in the plain model assuming one-to-one one-way functions [GMW87], all that remains in order to derive Thm. 1 as a corollary of Thm. 5 is an unconditional construction of noisy-leakage resilient one-time non-malleable secret sharing for arbitrary access structures against independent leakage and tampering. The only known scheme achieving all these properties unconditionally is the one in [KMS18], but unfortunately that scheme only tolerates bounded leakage, and it is unclear how to generalize the proof to the setting of noisy leakage.¹⁹ Hence, we take a different approach and we instead show how to generalize a recent transformation from [BGW19], which is tailored to the case $n = 2$.

6.1 Asymmetric Noisy-Leakage-Resilient Secret Sharing

Our construction exploits so-called leakage-resilient encryption, as recently introduced by Ball, Guo, and Wichs [BGW19]. To keep the exposition more uniform, we cast their definition in terms of a special 2-out-of-2 leakage-resilient secret sharing satisfying three additional properties: (i) One of the shares is uniformly random, and can be sampled independently from the message; (ii) The shares are almost uncorrelated, namely the distribution of one share in isolation and conditioned on the other share have very similar min-entropy; (iii) The size of the shares are asymmetric, namely one share is substantially larger than the other share. Given a 2-out-of-2 secret sharing scheme $\Sigma = (\text{Share}, \text{Rec})$, abusing notation, for any fixed $s_1 \in \mathcal{S}_1$ and $m \in \mathcal{M}$, we write $s_2 \leftarrow^s \text{Share}(m, s_1)$ for the sharing algorithm that computes share s_2 subject to (s_1, s_2) being a valid sharing of m .

Definition 13 (Asymmetric secret sharing). Let $\Sigma = (\text{Share}, \text{Rec})$ be a 2-out-of-2 secret sharing scheme. We call Σ $(\alpha, \sigma_1, \sigma_2)$ -*asymmetric*, if it satisfies the following properties:

- (i) For any $s_1 \in \mathcal{S}_1$, and any $m \in \mathcal{M}$, it holds that $\text{Rec}(s_1, \text{Share}(m, s_1)) = m$;
- (ii) For any message $m \in \mathcal{M}$, and for all $i \in \{1, 2\}$, it holds that $\tilde{\mathbb{H}}_\infty(\mathbf{S}_i | \mathbf{S}_{3-i}) \geq \log |\mathcal{S}_i| - \alpha$, where $\mathbf{S}_1, \mathbf{S}_2$ are the random variables corresponding to sampling $s_1 \leftarrow^s \mathcal{S}_1$ and $s_2 \leftarrow^s \text{Share}(m, s_1)$;
- (iii) It holds that $\log |\mathcal{S}_1| = \sigma_1$ and $\log |\mathcal{S}_2| = \sigma_2$.

As for security we consider the same security experiment of a leakage-resilient secret sharing, however, we consider a more general class of admissible adversaries:

Definition 14 (Independent noisy-leakage admissibility for asymmetric secret sharing). Let $\Sigma = (\text{Share}, \text{Rec})$ be a 2-out-of-2 secret sharing scheme. We say that an unbounded adversary $A = (A_1, A_2)$ is independent (ℓ_1, ℓ_2) -asymmetric noisy-leakage admissible $((\ell_1, \ell_2)$ -NLA for short) if it satisfies Def. 12 without property (iii), and using the following variant of property (ii):

- (ii) A_2 outputs a sequence of leakage queries (chosen adaptively) $(g^{(a)})_{a \in [p]}$, with $p(\lambda) \in \text{poly}(\lambda)$, such that for all $i \in \{1, 2\}$, and for all $m \in \mathcal{M}$:

$$\tilde{\mathbb{H}}_\infty \left(\mathbf{S}_i | \mathbf{S}_{3-i}, g_i^{(1)}(\mathbf{S}_i), \dots, g_i^{(p)}(\mathbf{S}_i) \right) \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_i | \mathbf{S}_{3-i}) - \ell_i,$$

where \mathbf{S}_1 is uniformly random over \mathcal{S}_1 and \mathbf{S}_2 is the random variable corresponding to $\text{Share}(m, \mathbf{S}_1)$.

¹⁹This is because [KMS18] relies on lower bounds in communication complexity.

Let $\Sigma' = (\text{Share}', \text{Rec}')$ be a secret sharing scheme realizing access structure \mathcal{A} , with message space \mathcal{M} and share space $\mathcal{S}' = \mathcal{S}'_1 \times \cdots \times \mathcal{S}'_n$ where $\mathcal{S}'_i \subseteq \mathcal{M}''$. Let $\Sigma'' = (\text{Share}'', \text{Rec}'')$ be a 2-out-of-2 *asymmetric* secret sharing scheme with message space \mathcal{M}'' and share space $\mathcal{S}'' = \mathcal{S}''_1 \times \mathcal{S}''_2$. Define the following secret sharing scheme $\Sigma = (\text{Share}, \text{Rec})$, with message space \mathcal{M} and share space $\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$, where for each $i \in [n]$ we have $\mathcal{S}_i \subseteq (\mathcal{S}''_1)^{n-1} \times (\mathcal{S}''_2)^{n-1}$.

Sharing algorithm Share: Upon input a value $m \in \mathcal{M}$, compute $(s'_1, \dots, s'_n) \leftarrow \text{Share}'(m)$. For each $i \in [n]$ and $j \in [n] \setminus \{i\}$, sample a random share $s''_{i,j,1} \leftarrow \mathcal{S}''_1$ and compute $s''_{i,j,2} \leftarrow \text{Share}''(s'_i, s''_{i,j,1})$. Return the shares $s = (s_1, \dots, s_n)$, where for each $i \in [n]$ we set $s_i = ((s''_{j,i,1})_{j \neq i}, (s''_{i,j,2})_{j \neq i})$.

Reconstruction algorithm Rec: Upon input shares $(s_i)_{i \in \mathcal{I}}$ with $\mathcal{I} \in \mathcal{A}$, parse $s_i = ((s''_{j,i,1})_{j \neq i}, (s''_{i,j,2})_{j \neq i})$. Hence, proceed as follows:

- Compute $s'_i = \text{Rec}''(s''_{i,\text{nxt}(i),1}, s''_{i,\text{nxt}(i),2})$ for $i \in \mathcal{I}$;
- Return $\text{Rec}'((s'_i)_{i \in \mathcal{I}})$.

Figure 7: Noisy-leakage-resilient one-time statistically non-malleable secret sharing for arbitrary access structures against independent leakage and tampering in the plain model.

Finally, we say that a 2-out-of-2 secret sharing is augmented (ℓ_1, ℓ_2) -noisy-leakage resilient if it is secure as per Def. 11, against the class of all unbounded adversaries that are (ℓ_1, ℓ_2) -NLA. The theorem below says that there is an unconditional construction of such a leakage-resilient secret sharing that is also asymmetric as per Def. 13. The proof appears in §A.

Theorem 6. *For any $\alpha \in \mathbb{N}$, and for any large enough $\ell_1, \ell_2 \in \text{poly}(\lambda, \alpha)$, there exists $\sigma_1, \sigma_2 \in \text{poly}(\lambda, \alpha)$ and an $(\alpha, \sigma_1, \sigma_2)$ -asymmetric secret sharing scheme Σ with message space $\{0, 1\}^\alpha$ that is augmented (ℓ_1, ℓ_2) -noisy leakage resilient.*

6.2 Construction

Before presenting our scheme, we establish some notation. Given a reconstruction set $\mathcal{I} = \{i_1, \dots, i_k\}$, we always assume that $i_j \leq i_{j+1}$ for $j \in [k]$. further, we define the function $\text{nxt}_{\mathcal{I}} : \mathcal{I} \rightarrow \mathcal{I}$ as:

$$\text{nxt}_{\mathcal{I}}(i_j) := \begin{cases} i_{j+1} & j < k \\ i_1 & \text{otherwise} \end{cases}$$

and the function $\text{prv}_{\mathcal{I}}$ to be the inverse of $\text{nxt}_{\mathcal{I}}$. Whenever it is clear from the context we omit the reconstruction set \mathcal{I} and simply write nxt and prv .

Intuitively, our construction (cf. Fig. 7) relies on a one-time non-malleable (but not leakage resilient) secret sharing Σ' , and on an asymmetric leakage-resilient secret sharing Σ'' . The sharing of a message m is obtained by first sharing m under Σ' , obtaining n shares (s'_1, \dots, s'_n) , and then sharing each s_i independently $n - 1$ times under Σ'' , obtaining pairs of shares $(s''_{i,j,1}, s''_{i,j,2})_{j \neq i}$; the final share of party i is then set to be the collection of right shares corresponding to i and all the left shares corresponding to the parties $j \neq i$. We can now state the main theorem of this section.

Theorem 7. *Let $n \in \mathbb{N}$, and let \mathcal{A} be an arbitrary access structure for n parties without singletons. Assume that:*

- Σ' is an n -party one-time non-malleable secret sharing scheme realizing access structure \mathcal{A} against independent tampering in the plain model, with information-theoretic security;

(ii) Σ'' is an $(\alpha, \sigma_1, \sigma_2)$ -asymmetric augmented (ℓ_1, ℓ_2) -noisy leakage-resilient secret sharing scheme.

Then, the secret sharing scheme Σ described in Fig. 7 is an n -party ℓ -noisy leakage-resilient one-time non-malleable secret sharing scheme realizing access structure \mathcal{A} against independent leakage and tampering with statistical security in the plain model, as long as $\ell_1 = \ell + (2n - 3)\alpha$ and $\ell_2 = \ell + (2n - 3)\alpha + \sigma_1$.

6.3 Proof Overview

Privacy of Σ follows in a fairly straightforward manner from privacy of Σ' . In fact, recall that the shares $s''_{i,j,1}$, with $i, j \in [n]$ and $i \neq j$, are sampled uniformly at random and independently of s' . Thus, in the reduction we can sample these values locally and then define the shares $(s_u)_{u \in \mathcal{U}}$ as a function of the shares $(s'_u)_{u \in \mathcal{U}}$. As for the proof of leakage-resilient one-time non-malleability, the idea is to reduce to the one-time non-malleability of Σ' and simulate the leakage by sampling dummy values for the shares $s''_{i,j,1}, s''_{i,j,2}$.

The main challenge is to make sure that the answer to tampering query $f = (f_1, \dots, f_n)$ is consistent with the simulated leakage. To this end, in the reduction we define the tampering function $f' = (f'_1, \dots, f'_n)$, acting on the shares $s' = (s'_1, \dots, s'_n)$, as follows. Each function f'_i , upon input s'_i and given the values $(s''_{i,j,1})_{j \neq i}$, samples $(\hat{s}_{i,j,2})_{j \neq i}$ in such a way that for any j the reconstruction $\text{Rec}''(s_{i,j,1}, \hat{s}_{i,j,2})$ yields a share s'_i that is consistent with the simulated leakage using the dummy values. Noisy-leakage resilience of Σ'' guarantees that the function f'_i samples from a valid distribution (namely, a non-empty one). Note that the function f'_i might not be efficiently computable; however, as we are reducing to statistical non-malleability, this is not a problem.

An additional difficulty is that the functions $(f'_t)_{t \in \mathcal{T}}$ need to communicate in order to produce their outputs. In fact, for any $t \in \mathcal{T}$, the function f'_t returns a tampered share for Σ' that depends on the mauled share $\tilde{s}_{\text{prv}(t),t,1}$ (generated by $f_{\text{prv}(t)}$). To overcome this problem, we let the reduction perform an additional leakage query on the dummy values before tampering. Thanks to this extra leakage, the reduction learns the values $\tilde{s}_{\text{prv}(t),t,1}$ for all $t \in \mathcal{T}$, which can be hard-coded in the description of $(f'_t)_{t \in \mathcal{T}}$. Here is where we rely on the asymmetric property of Σ'' , which allows us to leak σ_1 bits from the second share.

At this point, a reader familiar with [BGW19] might notice that the two proofs proceed very similarly. However, our proof requires extra care when bounding the amount of leakage performed by the reduction. The key ideas are that: (i) Each of the shares under Σ' is shared using $n - 1$ independent invocations of Σ'' ; and (ii) our reconstruction procedure depends only on one of those (chosen as function of the reconstruction set). Property (i) allows to reduce independent leakage on n shares under Σ to independent leakage on 2 shares under Σ'' by sampling locally the missing $n - 2$ shares when reducing to noisy-leakage resilience of Σ'' . Property (ii) allows to bound the amount of information the reduction needs to simulate the tampering query to a single short leakage from each of the shares (i.e., the value $\tilde{s}_{\text{prv}(t),t,1}$ for $t \in \mathcal{T}$).

6.4 Security Analysis

The proof relies on the following technical lemma.

Lemma 14. *Let Σ be a 2-out-of-2 secret sharing scheme with message space \mathcal{M} . If there exists $\tilde{\sigma} \in \mathbb{R}$ such that for all $m \in \mathcal{M}$ we have $\tilde{\mathbb{H}}_\infty(\mathbf{S}_a | \mathbf{S}_{3-a}) \geq \tilde{\sigma}$, where $\mathbf{S}_1, \mathbf{S}_2$ are the random variables corresponding to $\text{Share}(m)$, then for all $k \in \mathbb{N}$, any $a_1, \dots, a_k \in \{1, 2\}^k$ and any*

distribution \mathbf{D} over \mathcal{M}^k , we have:

$$\tilde{\mathbb{H}}_\infty((\mathbf{S}_{a_j,j})_{j \in [k]} | (\mathbf{S}_{3-a_j,j})_{j \in [k]}) \geq \tilde{\mathbb{H}}_\infty((\mathbf{S}_{a_j,j})_{j \in [k-1]} | (\mathbf{S}_{3-a_j,j})_{j \in [k-1]}) + \tilde{\sigma},$$

where $\mathbf{S}_{1,j}, \mathbf{S}_{2,j}$ are the random variables corresponding to $\text{Share}(\mathbf{M}_j)$ and $\mathbf{M} = \mathbf{M}_1, \dots, \mathbf{M}_k \leftarrow^{\$} \mathbf{D}$.

Proof. Let $\bar{\mathbf{S}}_a = (\mathbf{S}_{a_j,j})_{j \in [k-1]}$, $\bar{\mathbf{S}}_{3-a} = (\mathbf{S}_{3-a_j,j})_{j \in [k-1]}$, and similarly let $\bar{\mathbf{S}}'_a = (\mathbf{S}_{a_j,j})_{j \in [k]}$ and $\bar{\mathbf{S}}'_{3-a} = (\mathbf{S}_{3-a_j,j})_{j \in [k]}$. By definition of conditional average min-entropy, we need to show that:

$$\max_{\mathbf{P}} \mathbb{P} [\mathbf{P}(\bar{\mathbf{S}}_{3-a}, \mathbf{S}_{3-a_k,k}) = \bar{\mathbf{S}}_a, \mathbf{S}_{a_k,k}] \leq 2^{-\tilde{\sigma}} \max_{\mathbf{P}} \mathbb{P} [\mathbf{P}(\bar{\mathbf{S}}_{3-a}) = \bar{\mathbf{S}}_a].$$

Given a predictor \mathbf{P} for the left-side of the equation above, let \mathbf{P}' be the predictor that runs as \mathbf{P} but outputs only the first $k-1$ values, and let \mathbf{P}'' be the predictor that runs as \mathbf{P} but outputs only the last value. As we can wlog. assume that \mathbf{P} is deterministic, we have that $\mathbf{P}(\cdot) = \mathbf{P}'(\cdot) \parallel \mathbf{P}''(\cdot)$. We can write:

$$\begin{aligned} & \max_{\mathbf{P}} \mathbb{P} [\mathbf{P}(\bar{\mathbf{S}}'_{3-a}) = \bar{\mathbf{S}}_a, \mathbf{S}_{a_k,k}] = \\ & \max_{\mathbf{P}} \mathbb{E}_m \left[\mathbb{P} [\mathbf{P}(\bar{\mathbf{S}}'_{3-a}) = \bar{\mathbf{S}}_a, \mathbf{S}_{a_k,k} \mid \mathbf{M} = m] \right] = \\ & \max_{\mathbf{P}} \mathbb{E}_m \left[\mathbb{P} [\mathbf{P}'(\bar{\mathbf{S}}'_{3-a}) = \bar{\mathbf{S}}_a \mid \mathbf{M} = m, \mathbf{P}''(\bar{\mathbf{S}}'_{3-a}) = \mathbf{S}_{a_k,k}] \right. \\ & \quad \left. \cdot \mathbb{P} [\mathbf{P}''(\bar{\mathbf{S}}'_{3-a}) = \mathbf{S}_{a_k,k} \mid \mathbf{M} = m] \right] = \\ & \max_{\mathbf{P}', \mathbf{P}''} \mathbb{E}_m \left[\mathbb{P} [\mathbf{P}'(\bar{\mathbf{S}}_{3-a}) = \bar{\mathbf{S}}_a \mid \mathbf{M} = m, \mathbf{P}''(\mathbf{S}_{3-a_k,k}) = \mathbf{S}_{a_k,k}] \right. \\ & \quad \left. \cdot \mathbb{P} [\mathbf{P}''(\mathbf{S}_{3-a_k,k}) = \mathbf{S}_{a_k,k} \mid \mathbf{M} = m] \right] = \\ & \max_{\mathbf{P}'} \mathbb{E}_m \left[\mathbb{P} [\mathbf{P}'(\bar{\mathbf{S}}_{3-a}) = \bar{\mathbf{S}}_a \mid \mathbf{M} = m, \mathbf{P}''(\mathbf{S}_{3-a_k,k}) = \mathbf{S}_{a_k,k}] \right] \cdot 2^{-\tilde{\sigma}}, \end{aligned}$$

where in the third equality we used the fact that, once we fix $\mathbf{M} = m$, the random variables $(\mathbf{S}_{1,j}, \mathbf{S}_{2,j})$ for any j are independently distributed, and in the last equality we used the hypothesis of the lemma. \square

Proof of Thm. 7. We use a hybrid argument. Recall that the experiment $\mathbf{Tamper}_{\Sigma, \mathbf{A}}(\lambda, b)$ computes the target secret sharing by first sampling $s'_1, \dots, s'_n \leftarrow^{\$} \text{Share}'(m_b)$, and $s''_{i,j,1}, s''_{i,j,2}$ for $i, j \in [n]$ and $j \neq i$, and finally defines s_1, \dots, s_n by arranging the shares as described in Fig 7.

First hybrid. Let $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1(\lambda, b)$ be the same as $\mathbf{Tamper}_{\Sigma, \mathbf{A}}(\lambda, b)$, with the following differences.

1. The hybrid records all answers to \mathbf{A} 's leakage queries. In particular, let τ be the transcript containing all such answers; we can parse τ as τ_1, \dots, τ_n , where τ_i is the leakage obtained from the share s_i .
2. When \mathbf{A} outputs its tampering query $(\mathcal{T}, (f_1, \dots, f_n))$:
 - (a) For $t \in \mathcal{T}$, compute the value $\tilde{s}_t = ((\tilde{s}''_{i,j,1})_{j \neq i}, (\tilde{s}''_{j,i,2})_{j \neq i}) = f_t(s_t)$, and set $\tau_t \leftarrow \tau_t \parallel \tilde{s}''_{\text{prv}(t), t, 1}$.
 - (b) For $t \in \mathcal{T}$, sample $(\hat{s}_{t,j,2})_{j \neq t}$ conditioned on:
 - i. $\text{Rec}(s''_{t,j,1}, \hat{s}_{t,j,2}) = s'_t$ for any j , where s'_t is the value computed by the sharing algorithm;

- ii. The values $(\hat{s}_{t,j,2})_{j \neq i}$ being consistent with the leakage transcript τ_t , namely, let $\hat{s}_t := ((s''_{j,i,1})_{j \neq i}, (\hat{s}_{i,j,2})_{j \neq i})$, by applying all the leakage functions to $\hat{s}_t = ((s''_{j,i,1})_{j \neq i}, (\hat{s}_{i,j,2})_{j \neq i})$ we obtain the value τ_t .
- (c) Output $\underline{\text{Rec}}((f_t(\hat{s}_t))_{t \in \mathcal{T}})$ (instead of $\text{Rec}((f_t(s_t))_{t \in \mathcal{T}})$).

Lemma 15. $\forall b \in \{0, 1\}: \{\mathbf{Tamper}_{\Sigma, \mathbf{A}}(\lambda, b)\}_{\lambda \in \mathbb{N}} \equiv \{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1(\lambda, b)\}_{\lambda \in \mathbb{N}}$.

Proof. Notice that the hybrid samples $(\hat{s}_t)_{t \in \mathcal{T}}$ from exactly the same conditional distribution where the shares $(s_t)_{t \in \mathcal{T}}$ are picked. Thus the output of the tampering oracle is distributed identically in the two experiments.

Formally, let \mathbf{Z} be the random variable describing the tuple $(s'_i, (s''_{i,j,1})_{j \neq i}, \tau_i)_{i \in [n]}$, let \mathbf{S} be the random variable $(s''_{t,j,2})_{t \in \mathcal{T}, j \neq t}$, and let $\hat{\mathbf{S}}$ be the random variable $(\hat{s}_{t,j,2})_{t \in \mathcal{T}, j \neq t}$. We can think of \mathbf{Z} being the output of a randomized function of \mathbf{S} : let F be such function, then $\mathbf{Z} = F(\mathbf{S})$. Notice that the two hybrids are equivalent if the distributions $(\mathbf{S}, F(\mathbf{S}))$ and $(\hat{\mathbf{S}}, F(\mathbf{S}))$ are equivalent. To prove this, we observe that for any assignments z and $s = ((s_{t,j})_{j \neq t})_{t \in \mathcal{T}}$ the following holds:

$$\Pr[\mathbf{S} = s | \mathbf{Z} = z] = \prod_t \mathbb{P}[(\mathbf{S}''_{t,j,2})_{j \neq t} = (s_{t,j})_{j \neq t} | \mathbf{Z} = z] \quad (7)$$

$$= \prod_t \mathbb{P}[(\hat{\mathbf{S}}_{t,j,2})_{j \neq t} = (s_{t,j})_{j \neq t} | \mathbf{Z} = z] \quad (8)$$

$$= \Pr[\hat{\mathbf{S}} = s | \mathbf{Z} = z]. \quad (9)$$

Eq. (7) follows because for any $i \in [n]$, and for any assignments of $s' = s'_1, \dots, s'_n$ and $(s''_{i,j,1})_{j \neq i}$, the random variables $((\mathbf{S}''_{i,j,2})_{j \neq i} | \mathbf{S}' = s', (\mathbf{S}''_{i,j,1})_{j \neq i} = (s''_{i,j,1})_{j \neq i})$ are independent, and thus by Lemma 1 the random variables $((\mathbf{S}''_{i,j,2})_{j \neq i} | \mathbf{Z} = z)$ for any $i \in [n]$ are independent too. Eq. (8) uses the fact that, once we fix \mathbf{Z} , for any t the random variables $(\mathbf{S}''_{t,j,2})_{j \neq t}$ and $(\hat{\mathbf{S}}_{t,j,2})_{j \neq t}$ are chosen from the same distribution. Finally, Eq. (9) follows because for any t the random variables $(\hat{\mathbf{S}}_{t,j,2})_{j \neq t}$ are independently sampled. \square

Second hybrid. Let $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, b)$ be the same as $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1(\lambda, b)$, except that:

1. The hybrid additionally samples for any $i \in [n]$ and $j \in [n] \setminus \{i\}$ the share $s_{i,j,2}^* \leftarrow \text{Share}(0^\mu, s''_{i,j,1})$, and sets $s_i^* = ((s''_{j,i,1})_{j \neq i}, (s_{i,j,2}^*)_{j \neq i})$.
2. The answers to A 's leakage queries are answered using the shares s_1^*, \dots, s_n^* .

Lemma 16. $\forall b \in \{0, 1\}: \{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1(\lambda, b)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, b)\}_{\lambda \in \mathbb{N}}$.

Proof. We use a hybrid argument over all indices i, j , where $j \neq i$ and $i, j \in \mathbb{N}$. To simplify notation, let $N = \binom{n}{2}$ and define $\binom{[n]}{2} := \{\mathcal{H}_1, \dots, \mathcal{H}_N\}$. Set $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(0)} := \mathbf{Hyb}_{\Sigma, \mathbf{A}}^2$, and for any $l > 0$ consider the hybrid $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l)}$ to be the same as $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l-1)}$ but where $s_{h_1, h_2, 2}^* \leftarrow \text{Share}(s'_{h_1}, s''_{h_1, h_2, 1})$ for $\mathcal{H}_l = \{h_1, h_2\}$. It is clear that $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(N)}$ is distributed exactly as $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1$. We prove that $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l)}$ and $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l-1)}$ are statistically close via a reduction to leakage resilience of Σ'' . Consider the adversary $\mathbf{B} = (\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3)$ defined below.

Adversary $\mathbf{B}_1(1^\lambda)$:

1. Sample $(s'_1, \dots, s'_n) \leftarrow \text{Share}'(m_b)$;
2. Return state $\alpha_1 = (s'_1, \dots, s'_n)$, and challenge messages $(s'_{h_1}, 0^\mu)$.

Adversary $\mathbf{B}_2^{\mathcal{O}_{\text{leak}}((s''_{h_1, h_2, 1}, s_{h_1, h_2, 2}^*), \cdot)}(\alpha_1)$:

3. Sample $s''_{h'_1, h'_2, 1}$ and $s^*_{h'_1, h'_2, 2}$ for any $\{h'_1, h'_2\} \neq \mathcal{H}_l$, as described in $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l)}$.
 4. Run \mathbf{A} and answer to its leakage queries. In particular, upon input (g_1, \dots, g_n) , for any index $k \notin \mathcal{H}_l$ the output of leakage function g_k can be computed without the help of the leakage oracle. Let z_k be this output. Hence:
 - Let $\bar{g}_1(s''_{h_1, h_2, 1})$ be the function that first defines s_{h_1} by plugging $s''_{h_1, h_2, 1}$ and hard-coding the remaining values, and then outputs $g_{h_1}(s_{h_1})$.
 - Let $\bar{g}_2(s''_{h_1, h_2, 2})$ be the function that first defines s_{h_2} by plugging $s''_{h_1, h_2, 2}$ and hard-coding the remaining values, and then outputs $g_{h_2}(s_{h_2})$.
- Forward the leakage query (\bar{g}_1, \bar{g}_2) to the challenger, and denote by z_{h_1}, z_{h_2} the corresponding answer. Forward the values z_1, \dots, z_n to \mathbf{A} , and update the transcript $\tau \leftarrow \tau \parallel (z_1, \dots, z_n)$.
5. Eventually the adversary \mathbf{A} outputs a tampering query $(\mathcal{T}, (f_1, \dots, f_n))$
 6. Consider the leakage function $g_2^*(s''_{h_1, h_2, 2})$ that first defines s_{h_2} by plugging $s''_{h_1, h_2, 2}$ and hard-coding the remaining values, then lets $\tilde{s}_{h_2} = f_{h_2}(s_{h_2})$, and finally outputs $\tilde{s}''_{\text{prv}(h_2), h_2, 1}$. Forward the leakage query (ε, g_2^*) to the challenger, and let $(\varepsilon, \tilde{s}''_{\text{prv}(h_2), h_2, 1})$ be the corresponding answer. Set $\tau_{h_2} \leftarrow \tau_{h_2} \parallel \tilde{s}''_{\text{prv}(h_2), h_2, 1}$.
 7. Let α_2 be the full state of \mathbf{B} up to this point, including the transcript τ , and the state $\alpha_{\mathbf{A}}$ of attacker \mathbf{A} .

Adversary $\mathbf{B}_3(\alpha_2, s''_{h_1, h_2, 1})$:

8. Compute $\tilde{s}''_{\text{prv}(t), t, 1}$ for $t \in \mathcal{T} \setminus \{h_2\}$, as described in $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1$; in particular, $\tau_t \leftarrow \tau_t \parallel \tilde{s}''_{\text{prv}(t), t, 1}$.
9. Sample $(\hat{s}_{t, j, 2})_{j \neq t}$ for $t \in \mathcal{T}$, conditioned on the knowledge of $s'_t, (s''_{t, j, 1})_{j \neq t}$ and τ , as described in $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1$.
10. Set \hat{s}_t for $t \in \mathcal{T}$, as described in step 2b of $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1$, and compute the tampered message \tilde{m} using these shares.
11. Resume \mathbf{A} using $\alpha_{\mathbf{A}}$, forward to \mathbf{A} the tampered message \tilde{m} , and output whatever \mathbf{A} does.

Claim 1. *The following holds $\forall b \in \{0, 1\}$:*

$$\begin{aligned} \{\mathbf{Leak}_{\Sigma'', \mathbf{B}}(\lambda, 0)\}_{\lambda \in \mathbb{N}} &\equiv \{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l-1)}(\lambda, b)\}_{\lambda \in \mathbb{N}} \\ \{\mathbf{Leak}_{\Sigma'', \mathbf{B}}(\lambda, 1)\}_{\lambda \in \mathbb{N}} &\equiv \{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^{(l)}(\lambda, b)\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

Proof. The proof of the claim holds by inspection of adversary \mathbf{B} . □

Claim 2. *If Σ'' has α -correlated shares and adversary \mathbf{A} is ℓ -NLA, then adversary \mathbf{B} is $(\ell + (2n - 3)\alpha, \ell + (2n - 3)\alpha + \sigma_1)$ -NLA.*

Proof. We analyze the leakage performed by \mathbf{B}_2 on $s''_{h_1, h_2, 1}$ and $s^*_{h_1, h_2, 2}$. Let \mathbf{S}''_1 be the random variable describing the value $s''_{h_1, h_2, 1}$, let \mathbf{S}''_2 be the random variable describing the value $s''_{h_1, h_2, 2}$, let $g^{(q)} := (g_1^{(q)}, \dots, g_n^{(q)})$ be the q -th leakage query output by \mathbf{A} , and $(\bar{g}_1^{(q)}, \bar{g}_2^{(q)})$ be the q -th leakage query output by \mathbf{B} in the reduction. For each $a \in \{1, 2\}$, let $\mathbf{\Lambda}_a := (\bar{g}_a^{(1)}(\mathbf{S}''_a), \dots, \bar{g}_a^{(p)}(\mathbf{S}''_a))$ be the complete transcript containing all answers to leakage queries, and let $\mathbf{\Lambda}_2^* := (\mathbf{\Lambda}_2, g_2^*(\mathbf{S}''_2))$. Finally, for $k \in [n]$, let \mathbf{S}_k be the random variable describing the share s_k^* as sampled by \mathbf{B} , and let $\mathbf{S}_{h_1}^*$ (resp. $\mathbf{S}_{h_2}^*$) be the part of share $s_{h_1}^*$ (resp. $s_{h_2}^*$) which does not include $s''_{h_1, h_2, 2}$ (resp. $s''_{h_1, h_2, 1}$). First, observe that:

$$\tilde{\mathbb{H}}_{\infty}(\mathbf{S}''_2 \mid \mathbf{S}''_1, \mathbf{\Lambda}_1^*) \geq \tilde{\mathbb{H}}_{\infty}(\mathbf{S}''_2 \mid \mathbf{S}''_1, \mathbf{\Lambda}_1) - \sigma_1, \quad (10)$$

which holds by Lemma 2. Moreover, for $a \in \{1, 2\}$, the following holds:

$$\begin{aligned}
& \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'' \mid \mathbf{S}_{3-a}'', \mathbf{\Lambda}_a) \\
& \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'' \mid \mathbf{S}_{3-a}'', \mathbf{S}_{h_{3-a}}^*, (\mathbf{S}_k)_{k \in [n] \setminus \mathcal{H}_1}, \mathbf{\Lambda}_a) \\
& = \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'' \mid (\mathbf{S}_k)_{k \neq h_a}, \mathbf{\Lambda}_a) \\
& \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'', \mathbf{S}_{h_a}^* \mid (\mathbf{S}_k)_{k \neq h_a}, \mathbf{\Lambda}_a) - ((n-1)\sigma_{3-a} + (n-2)\sigma_a) \tag{11} \\
& \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'', \mathbf{S}_{h_a}^* \mid (\mathbf{S}_k)_{k \neq h_a}) - ((n-1)\sigma_{3-a} + (n-2)\sigma_a) - \ell. \tag{12}
\end{aligned}$$

Eq. (11) follows by Lemma 2, whereas Eq. (12) follows by admissibility of \mathbf{A} . For any $k \neq h_a$, we can split \mathbf{S}_k into two pieces: the random variable $\mathbf{S}_k^{(1)}$ corresponding $(s_{k,h_a,1}'', s_{h_a,k,2}'')$, and the random variable $\mathbf{S}_k^{(2)}$ corresponding to $(s_{k,j,1}'')_{j \neq k, j \neq h_a}, (s_{j,k,2}^*)_{j \neq k, j \neq h_a}$. Recall that Σ has α -correlated shares; for convenience, we define $\tilde{\sigma}_i = \sigma_i - \alpha$. Let \mathbf{S}' be the random variable associated to the output of $\text{Share}'(m)$. For any assignments of $s' = (s'_1, \dots, s'_n) \leftarrow^* \text{Share}(m)$, we can write:

$$\begin{aligned}
& \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'', \mathbf{S}_{h_a}^* \mid (\mathbf{S}_k^{(1)}, \mathbf{S}_k^{(2)})_{k \neq h_a}, \mathbf{S}' = s') \\
& = \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'', \mathbf{S}_{h_a}^* \mid (\mathbf{S}_k^{(1)})_{k \neq h_a}, \mathbf{S}' = s') \tag{13} \\
& \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'' \mid \mathbf{S}_{3-a}'', \mathbf{S}' = s') + (n-1)\tilde{\sigma}_{3-a} + (n-2)\tilde{\sigma}_a. \tag{14}
\end{aligned}$$

Eq. (13) follows because, once the value s' is fixed, then the random variables $\mathbf{S}_k^{(2)}$ are independent, whereas Eq. (14) follows by Lemma 14. Finally, by averaging over all possible assignments of s' , we get:

$$\tilde{\mathbb{H}}_\infty(\mathbf{S}_a'', \mathbf{S}_{h_a}^* \mid (\mathbf{S}_k^{(1)}, \mathbf{S}_k^{(2)})_{k \neq h_a}) \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_a'' \mid \mathbf{S}_{3-a}'') + (n-1)\tilde{\sigma}_{3-a} + (n-2)\tilde{\sigma}_a,$$

which together with Eq. (12) and Eq. (10) proves the claim. \square

\square

Lemma 17. $\{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, 0)\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, 1)\}_{\lambda \in \mathbb{N}}$.

Proof. Consider the following adversary \mathbf{B} against the non-malleability of Σ' given black-box access to an attacker \mathbf{A} playing in $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, b)$.

1. At the beginning \mathbf{B} runs $\mathbf{A}_1(1^\lambda)$, and outputs the same values $(m_0, m_1, \alpha_{\mathbf{A}})$ returned by \mathbf{A}_1 .
2. Sample $s_{i,j,1}'' \leftarrow^* \mathbf{S}_1''$ and $s_{i,j,2}^* \leftarrow^* \text{Share}(0^\mu, s_{i,j,1}'')$, for $i, j \in [n]$ and $i \neq j$, and set $s_i^* \leftarrow ((s_{j,i,1}'')_{j \neq i}, (s_{i,j,2}^*)_{j \neq i})$.
3. Resume the adversary \mathbf{A} by running $\mathbf{A}_2(\alpha_{\mathbf{A}})$, and reply to its leakage queries using the shares s_1^*, \dots, s_n^* .
4. Eventually, \mathbf{A}_2 outputs $(\mathcal{T}, (f_1, \dots, f_n))$. Sample the values $\tilde{s}_{\text{prv}(t), t, 1}''$ for $t \in \mathcal{T}$, as described in $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^1$, and consider the tampering function $f'_t(s'_t)$ that samples $(\hat{s}_{t,j,2})_{j \neq t}$ conditioned on: (i) $\text{Rec}(s_{i,j,1}, \hat{s}_{i,j,2}) = s'_i$ for all $j \neq i$; and (ii) the leakage transcript $\tau_t \parallel \tilde{s}_{\text{prv}(t), t, 1}''$ being consistent with $(\hat{s}_{i,j,2})_{j \neq i}$. The function sets $\hat{s}_i := ((s_{j,i,1}'')_{j \neq i}, (\hat{s}_{i,j,2})_{j \neq i})$, computes $\tilde{s}_i \leftarrow f(\hat{s}_i)$, and outputs $\text{Rec}(\tilde{s}_{i, \text{nxt}(i), 1}'', \tilde{s}_{i, \text{nxt}(i), 2}'')$.
5. Forward $(\mathcal{T}, (f'_1, \dots, f'_n))$ to the challenger, and send the answer back to \mathbf{A} .
6. Output whatever \mathbf{A} outputs.

We claim that for any $b \in \{0, 1\}$, we have $\mathbf{Tamper}_{\Sigma', \mathbf{B}}(\lambda, b) \equiv \mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, b)$. In fact, the adversary \mathbf{B} runs exactly as in $\mathbf{Hyb}_{\Sigma, \mathbf{A}}^2(\lambda, b)$. The only (syntactical) difference is that \mathbf{B} samples the values $(\hat{s}_{i,j,2})_{j \neq i}$ and uses the target tampering oracle to answer \mathbf{A} 's tampering query. The lemma then follows by statistical non-malleability of Σ' . \square

The statement of the theorem follows by putting together the above lemmas, and applying the triangular inequality. \square

7 Conclusions and Open Problems

We have shown new constructions of leakage-resilient continuously non-malleable secret sharing schemes, for general access structures. Our first scheme is in the plain model, and guarantees security against independent noisy leakage and tampering with all of the shares. Our second scheme is in the CRS model, and guarantees security against joint bounded leakage and tampering using a fixed partition of the n shares into non-overlapping blocks of size $O(\log n)$.

The two major questions left over by our work are whether continuous non-malleability against joint tampering is achievable in the plain model, or against adaptive (rather than selective) joint tampering with the shares. Interestingly, our proof strategy breaks down in the case of adaptive tampering, and this holds true even assuming that the inner leakage-resilient secret sharing is secure in the presence of adaptive joint leakage. Intuitively, the reason is that in the reduction we must run different copies of the adversary inside the leakage oracle; in particular, we use each block of the shares in order to simulate the answer to all tampering queries asked by each copy of the attacker, and this is clearly possible only if the adversary does not change the partition within each query.

It would also be interesting to achieve continuous non-malleability under joint selective partitioning for better values of the parameter k (namely, the attacker can tamper jointly with blocks of size super-logarithmic in n). Note that this would follow immediately by our result if we plug in our construction a leakage-resilient secret sharing scheme tolerating joint leakage from subsets of shares with size $\omega(\log n)$. Unfortunately, the only known secret sharing scheme achieving joint-leakage resilience is the one in [KMS18], and as the authors explain improving the parameters in their construction would lead to progress on longstanding open problems in complexity theory. We leave it open to establish whether this holds true even in the case of selective partitioning (recall that the scheme of [KMS18] achieves adaptive leakage resilience), or whether the current state of affairs can be improved in the computational setting (with or without trusted setup).

A further open question is to improve the rate of our constructions. Note that by applying the rate compiler of [FV19], we do get *rate-one* continuously non-malleable secret sharing for general access structures, against independent tampering in the plain model. However, this is well-known to be sub-optimal in the computational setting, where the optimal share size would be $O(\mu/n)$, with μ being the size of the message [Kra93]. Note that it is unclear whether the same rate compiler works also for our construction against joint tampering under selective partitioning. This is because the analysis in [FV19] crucially relies on the resilience of the initial rate-zero non-malleable secret sharing against noisy leakage, whereas our construction only achieves security in the bounded-leakage model.

References

- [AAG⁺16] Divesh Aggarwal, Shashank Agrawal, Divya Gupta, Hemanta K. Maji, Omkant Pandey, and Manoj Prabhakaran. Optimal computational split-state non-malleable codes. In *TCC*, pages 393–417, 2016.
- [ADKO15a] Divesh Aggarwal, Yevgeniy Dodis, Tomasz Kazana, and Maciej Obremski. Non-malleable reductions and applications. In *STOC*, pages 459–468, 2015.
- [ADKO15b] Divesh Aggarwal, Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Leakage-resilient non-malleable codes. In *TCC*, pages 398–426, 2015.
- [ADL14] Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. In *STOC*, pages 774–783, 2014.
- [ADN⁺19] Divesh Aggarwal, Ivan Damgård, Jesper Buus Nielsen, Maciej Obremski, Erick Purwanto, João Ribeiro, and Mark Simkin. Stronger leakage-resilient and non-malleable secret sharing schemes for general access structures. In *CRYPTO*, pages 510–539, 2019.
- [BGW19] Marshall Ball, Siyao Guo, and Daniel Wichs. Non-malleable codes for decision trees. In *CRYPTO*, pages 413–434, 2019.
- [BS19] Saikrishna Badrinarayanan and Akshayaram Srinivasan. Revisiting non-malleable secret sharing. In *EUROCRYPT*, pages 593–622, 2019.
- [CFV19] Sandro Coretti, Antonio Faonio, and Daniele Venturi. Rate-optimizing compilers for continuously non-malleable codes. In *ACNS*, pages 3–23, 2019.
- [CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.
- [CG14] Mahdi Cheraghchi and Venkatesan Guruswami. Non-malleable coding against bit-wise and split-state tampering. In *TCC*, pages 440–464, 2014.
- [CL18] Eshan Chattopadhyay and Xin Li. Non-malleable codes, extractors and secret sharing for interleaved tampering and composition of tampering. Cryptology ePrint Archive, Report 2018/1069, 2018. <https://eprint.iacr.org/2018/1069>.
- [DDV10] Francesco Davì, Stefan Dziembowski, and Daniele Venturi. Leakage-resilient storage. In *SCN*, pages 121–137, 2010.
- [DHLW10] Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs. Efficient public-key cryptography in the presence of key leakage. In *ASIACRYPT*, pages 613–631, 2010.
- [DKO13] Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Non-malleable codes from two-source extractors. In *CRYPTO*, pages 239–257, 2013.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.

- [DP07] Stefan Dziembowski and Krzysztof Pietrzak. Intrusion-resilient secret sharing. In *FOCS*, pages 227–237, 2007.
- [DPW10] Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *Innovations in Computer Science*, pages 434–452, 2010.
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *FOCS*, pages 308–317, 1990.
- [FMNV14] Sebastian Faust, Pratyay Mukherjee, Jesper Buus Nielsen, and Daniele Venturi. Continuous non-malleable codes. In *TCC*, pages 465–488, 2014.
- [FNSV18] Antonio Faonio, Jesper Buus Nielsen, Mark Simkin, and Daniele Venturi. Continuously non-malleable codes with split-state refresh. In *ACNS*, pages 1–19, 2018.
- [FV19] Antonio Faonio and Daniele Venturi. Non-malleable secret sharing in the computational setting: Adaptive tampering, noisy-leakage resilience, and improved rate. In *CRYPTO*, pages 448–479, 2019.
- [GK18a] Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing. In *STOC*, pages 685–698, 2018.
- [GK18b] Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing for general access structures. In *CRYPTO*, pages 501–530, 2018.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [GPR16] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *STOC*, pages 1128–1141, 2016.
- [Gro06] Jens Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *ASIACRYPT*, pages 444–459, 2006.
- [HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *CRYPTO*, pages 201–215, 1996.
- [KMS18] Ashutosh Kumar, Raghu Meka, and Amit Sahai. Leakage-resilient secret sharing. Cryptology ePrint Archive, Report 2018/1138, 2018. <https://ia.cr/2018/1138>.
- [Kra93] Hugo Krawczyk. Secret sharing made short. In *CRYPTO*, pages 136–146, 1993.
- [LCG⁺19] Fuchun Lin, Mahdi Cheraghchi, Venkatesan Guruswami, Reihaneh Safavi-Naini, and Huaxiong Wang. Non-malleable secret sharing against affine tampering. *CoRR*, abs/1902.06195, 2019.
- [Li17] Xin Li. Improved non-malleable extractors, non-malleable codes and independent source extractors. In *STOC*, pages 1144–1156, 2017.
- [LL12] Feng-Hao Liu and Anna Lysyanskaya. Tamper and leakage resilience in the split-state model. In *CRYPTO*, pages 517–532, 2012.

- [NS19] Jesper Buus Nielsen and Mark Simkin. Lower bounds for leakage-resilient secret sharing. Cryptology ePrint Archive, Report 2019/181, 2019. <https://eprint.iacr.org/2019/181>.
- [OPVV18] Rafail Ostrovsky, Giuseppe Persiano, Daniele Venturi, and Ivan Visconti. Continuously non-malleable codes in the split-state model from minimal assumptions. In *CRYPTO*, pages 608–639, 2018.
- [SCO⁺01] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *CRYPTO*, pages 566–598, 2001.
- [SV19] Akshayaram Srinivasan and Prashant Nalini Vasudevan. Leakage resilient secret sharing and applications. In *CRYPTO*, pages 480–509, 2019.

A Construction of Asymmetric Secret Sharing

We show that the construction of [BGW19] is secure in the noisy-leakage model. Let Ext be a seeded extractor with d -bit source, ρ -bit seed, and μ -bit output (cf. Def. 7 in §2.4). Let 2Ext be a two-source extractor with σ_1 -bit source and ρ -bit output (cf. Def. 8 in §2.4). Consider the following 2-out-of-2 secret sharing scheme $\Sigma = (\text{Share}, \text{Rec})$, with $\mathcal{S}_1 = \{0, 1\}^{\sigma_1}$, $\mathcal{S}_2 = \{0, 1\}^{d+\sigma_1+\mu}$, and $\mathcal{M} = \{0, 1\}^\mu$.

- Algorithm $\text{Share}(m)$ samples $s_1 \leftarrow_{\$} \{0, 1\}^{\sigma_1}$, $x \leftarrow_{\$} \{0, 1\}^d$, $y \leftarrow_{\$} \{0, 1\}^{\sigma_1}$, computes $r = 2\text{Ext}(s_1, y)$ and $z = \text{Ext}(x, r) \oplus m$, and finally outputs s_1 and $s_2 = (x, y, z)$;
- Algorithm $\text{Rec}(s_1, s_2)$ parses $s_2 = (x, y, z)$, and returns $z \oplus \text{Ext}(x, 2\text{Ext}(s_1, y))$.

Claim 3. *The above secret sharing scheme Σ is $(\alpha, \sigma_1, \sigma_2)$ -asymmetric, for $\alpha = \mu$, and $\sigma_2 = d + \sigma_1 + \mu$.*

Proof. The size of the shares follows by inspection. Moreover, the share s_1 is sampled uniformly at random and used as the first source for the two-source extractor. As for the conditional min-entropy, we have:

$$\begin{aligned} \tilde{\mathbb{H}}_\infty(\mathbf{S}_1|\mathbf{S}_2) &= \tilde{\mathbb{H}}_\infty((\mathbf{X}, \mathbf{Y}, \mathbf{Z})|\mathbf{S}_1) \geq \tilde{\mathbb{H}}_\infty(\mathbf{X}, \mathbf{Y}|\mathbf{S}_1) \geq |\mathbf{X}| + |\mathbf{Y}| = \sigma_2 - \mu \\ \tilde{\mathbb{H}}_\infty(\mathbf{S}_2|\mathbf{S}_1) &= \tilde{\mathbb{H}}_\infty(\mathbf{S}_1|(\mathbf{X}, \mathbf{Y}, \mathbf{Z})) \geq \tilde{\mathbb{H}}_\infty(\mathbf{S}_1|\mathbf{X}, \mathbf{Y}) - |\mathbf{Z}| \geq \sigma_1 - \mu. \end{aligned}$$

□

Following [BGW19, DDV10], we define a slightly different game $\mathbf{Leak}'_{A, \Sigma}(\lambda, b)$ which is the same as $\mathbf{Leak}_{A, \Sigma}$ but where the leakage oracle is instantiated with $s_1, (x, y)$ instead of s_1, s_2 , and at the end of the experiment the adversary A additionally receives z in full.

Claim 4. *Assume that for all (ℓ_1, ℓ_2) -NLA adversaries A' , we have*

$$|\mathbb{P}[\mathbf{Leak}'_{A', \Sigma}(\lambda, 0) = 1] - \mathbb{P}[\mathbf{Leak}'_{A', \Sigma}(\lambda, 1) = 1]| \leq \epsilon,$$

for some $\epsilon' \in [0, 1]$. Then, for all $(\ell_1, \ell_2 - \mu)$ -NLA adversaries A we have:

$$|\mathbb{P}[\mathbf{Leak}_{A, \Sigma}(\lambda, 0) = 1] - \mathbb{P}[\mathbf{Leak}_{A, \Sigma}(\lambda, 1) = 1]| \leq \epsilon \cdot 2^\mu.$$

Proof. By contradiction, let A be an $(\ell_1, \ell_2 - \mu)$ -NLA adversary for the original game with advantage $\epsilon \cdot 2^\mu$. We build an attacker A' in the modified game $\mathbf{Leak}'_{A, \Sigma}(\lambda, b)$. The reduction A' samples z' uniformly at random, and simply runs A by forwarding its leakage queries to its own leakage oracle with hard-wired value z' . Eventually, A' receives the real value z . If $z = z'$, then A' continues running A , and else it aborts outputting a random bit. It is easy to see that the winning probability of A' equals the probability that $z = z'$ times the winning probability of A , which yields the bound in the claim.

Note that conditioned on $z = z'$, the adversary A' is $(\ell_1, \ell_2 - \mu)$ -NLA (and thus (ℓ_1, ℓ_2) -NLA) since z is a deterministic function of the target shares. On the other hand, when $z \neq z'$, we can write:

$$\tilde{\mathbb{H}}_\infty((\mathbf{X}, \mathbf{Y})|\mathbf{S}_1, \mathbf{\Lambda}) \geq \tilde{\mathbb{H}}_\infty((\mathbf{X}, \mathbf{Y}, \mathbf{Z})|\mathbf{S}_1, \mathbf{\Lambda}) - |\mathbf{Z}| \geq \tilde{\mathbb{H}}_\infty((\mathbf{X}, \mathbf{Y})|\mathbf{S}_1) - \ell_2.$$

□

Claim 5. *If 2Ext is a $(\sigma_1 - \ell_1, \sigma_1 - \ell_2, \epsilon_2)$ -secure two-source extractor, and Ext is a $(d - \ell_2, \epsilon_1)$ -secure seeded extractor, then for all (ℓ_1, ℓ_2) -NLA adversaries A :*

$$|\mathbb{P}[\mathbf{Leak}'_{A, \Sigma}(\lambda, 0) = 1] - \mathbb{P}[\mathbf{Leak}'_{A, \Sigma}(\lambda, 1) = 1]| \leq 2(\epsilon_1 + \epsilon_2).$$

Proof. We proceed with a hybrid argument.

- Let $\mathbf{Hyb}_{\mathbf{A},\Sigma}^1(\lambda, b)$ be the same as $\mathbf{Leak}'_{\mathbf{A},\Sigma}(\lambda, b)$, except that we pick $r \leftarrow_{\$} \{0, 1\}^\rho$ instead of computing $r = 2\mathbf{Ext}(s_1, y)$. The statistical distance between $\mathbf{Leak}'_{\mathbf{A},\Sigma}(\lambda, b)$ and $\mathbf{Hyb}_{\mathbf{A},\Sigma}^1(\lambda, b)$ is at most ϵ_2 . To see this, let $\mathbf{\Lambda}$ be the leakage performed by \mathbf{A} . By Lemma 1, the random variables $\mathbf{Y}|\mathbf{\Lambda}$ and $\mathbf{S}_1|\mathbf{\Lambda}$ are independently distributed. Moreover, since \mathbf{A} is (ℓ_1, ℓ_2) -NLA, we have:

$$\begin{aligned} \tilde{\mathbb{H}}_\infty(\mathbf{S}_1|\mathbf{\Lambda}) &\geq \sigma_1 - \ell_1 \\ \tilde{\mathbb{H}}_\infty(\mathbf{Y}|\mathbf{\Lambda}) &\geq \tilde{\mathbb{H}}_\infty(\mathbf{X}, \mathbf{Y}|\mathbf{S}_1, \mathbf{\Lambda}) - |\mathbf{X}| \\ &\geq \tilde{\mathbb{H}}_\infty(\mathbf{X}, \mathbf{Y}|\mathbf{S}_1) - \ell_2 - d = \sigma_1 - \ell_2, \end{aligned}$$

so that we can apply the security of the two-source extractor $2\mathbf{Ext}$.

- Let $\mathbf{Hyb}_{\mathbf{A},\Sigma}^2(\lambda, b)$ be the same as $\mathbf{Hyb}_{\mathbf{A},\Sigma}^1(\lambda, b)$, except that we now pick $z \leftarrow_{\$} \{0, 1\}^\mu$. The statistical distance between $\mathbf{Hyb}_{\mathbf{A},\Sigma}^1(\lambda, b)$ and $\mathbf{Hyb}_{\mathbf{A},\Sigma}^2(\lambda, b)$ is at most ϵ_1 . In fact, neither the seed r nor the output z of \mathbf{Ext} are in the view of the adversary during the leakage phase. Furthermore, a derivation similar to the one above shows that the random variable $\mathbf{X}|\mathbf{\Lambda}$ has conditional average min-entropy at least $d - \ell_2$, so that we can apply the security of the seeded extractor \mathbf{Ext} .

The claim now follows by noting that for all (ℓ_1, ℓ_2) -NLA attackers \mathbf{A} we must have $\mathbf{Hyb}_{\mathbf{A},\Sigma}^2(\lambda, 0) \equiv \mathbf{Hyb}_{\mathbf{A},\Sigma}^2(\lambda, 1)$, as the view of the adversary in the last hybrid is independent of b . \square

Invoking the result of [CG88, DORS08], for any ℓ_1, ℓ_2 , any $\alpha = \mu$, and any ϵ_1, ϵ_2 , we can always find ρ, d, σ_1 polynomial in ℓ_1, ℓ_2, α and $\log(1/\epsilon_1), \log(1/\epsilon_2)$ such that \mathbf{Ext} is a $(d - \ell_2, \epsilon_1)$ -secure seeded extractor and $2\mathbf{Ext}$ is a $(\sigma_1 - \ell_1, \sigma_1 - \ell_2, \epsilon_2)$ -secure two-source extractor. By combining the above claims and setting $\epsilon_1 = \epsilon_2 = 2^{-\alpha - \lambda - 2}$, we obtain that Σ is both $(\alpha, \sigma_1, \sigma_2)$ -asymmetric and (ℓ_1, ℓ_2) -noisy leakage resilient. This concludes the proof of Thm. 6.