

Mind Your Coins: Fully Leakage-Resilient Signatures with Graceful Degradation

Antonio Faonio^{*1}, Jesper Buus Nielsen^{†2}, and Daniele Venturi¹

²*Aarhus University*

¹*Sapienza University of Rome*

November 5, 2014

Abstract

We construct new leakage-resilient signature schemes. Our schemes remain unforgeable against an adversary leaking arbitrary (yet bounded) information on the entire state of the signer (sometimes known as *fully* leakage resilience).

The main feature of our constructions, is that they offer a graceful degradation of security in situations where standard existential unforgeability is impossible. This property was recently put forward by Nielsen *et al.* (PKC 2014) to deal with settings in which the secret key is much larger than the size of a signature. One remarkable such case is the so-called Bounded Retrieval Model (BRM), where one intentionally inflates the size of the secret key while keeping constant the signature size and the computational complexity of the scheme.

Our main constructions have leakage rate $1 - o(1)$, and are proven secure in the standard model. We additionally give a construction in the BRM, relying on a random oracle. All of our schemes are described in terms of generic building blocks, but also admit efficient instantiations under fairly standard number-theoretic assumptions. Finally, we explain how to extend some of our schemes to the setting of noisy leakage, where the only restriction on the leakage functions is that the output does not decrease the min-entropy of the secret key by too much.

^{*}Part of the work done while at Aarhus University.

[†]Partially supported by Danish Council for Independent Research via DFF Starting Grant 10-081612. Partially supported by the European Research Commission Starting Grant 279447.

Contents

| | | | | | |
|----------|--|-----------|----------|---|-----------|
| 1 | Introduction | 2 | 4.3 | A Variant | 22 |
| 1.1 | Our Contribution | 3 | 4.4 | Proof Sketch of Theorem 4.2 . . . | 23 |
| 1.2 | Technical Overview | 4 | | | |
| 1.3 | Related Work | 5 | 5 | A Scheme without Erasures | 25 |
| 1.4 | Roadmap | 6 | 5.1 | <i>Secret Sharing</i> Hybrid Commitment | 25 |
| 2 | Preliminaries | 7 | 5.2 | The Signature Scheme | 27 |
| 2.1 | Notation | 7 | 5.3 | Proof Sketch | 28 |
| 2.2 | Random Variables and Min-Entropy | 7 | 6 | A Scheme in the Bounded Retrieval Model | 30 |
| 2.3 | Commitment Schemes | 7 | 6.1 | Ingredients | 31 |
| 2.4 | NIWI and NIZK Arguments . . . | 9 | 6.2 | Basic Construction | 32 |
| 3 | Fully-Leakage One-More Unforgeability | 10 | 6.3 | The Full-Fledged Scheme | 34 |
| 3.1 | Definition without Erasures . . . | 10 | 7 | Instantiations | 36 |
| 3.2 | Definition with Erasures | 11 | 8 | Extension to Noisy Leakage | 38 |
| 4 | A Scheme with Erasures | 12 | 8.1 | Proof Sketch of Theorem 8.1 . . . | 40 |
| 4.1 | The Basic Construction | 12 | 9 | Conclusions and Open Problems | 41 |
| 4.2 | Proof of Theorem 4.1 | 14 | | | |

1 Introduction

Cryptography relies on secret information and random sources to accomplish its tasks. In order for a given cryptographic primitive to be secure, it is typically required that its secrets and randomness are well-protected, and cannot be influenced by an attacker. In practice, however, it is not always possible to fulfil this requirement, and partial information about the secret state of a cryptosystem can leak to an external adversary, e.g., via so-called side-channel attacks exploiting physical characteristics of a crypto-device, such as power consumption [44], electromagnetic radiation [55], and running times [43].

Recently a lot of effort has been put into constructing cryptographic primitives that come along with some form of leakage resilience, meaning that the scheme should remain secure even in case the adversary obtains some type of leakage on the secrets used within the system. A common way to model leakage attacks, is to empower the adversary with access to a leakage oracle, taking as input (adaptively chosen) functions f_i and returning $f_i(st)$ where st is the current secret state of the cryptosystem under attack. Clearly some restriction on the functions f_i has to be put, as otherwise there is no hope for security. By now, a plethora of leakage models (corresponding to different ways how to restrict the functions f_i) have been proposed. We review the ones that are more relevant to our work below (and refer the reader to Section 1.3 for a bigger picture).

- **Bounded leakage:** One natural restriction is to just assume that the total length of the information leaked via the functions f_i is smaller than some a priori determined leakage bound ℓ . Usually the leakage bound ℓ is also related to the secret key size, so that a relatively large fraction of the secret key can be leaked. Leakage-resilient schemes in this model include storage schemes, public-key and identity-based encryption, signature schemes, and more (see, e.g., [49, 42, 2, 46, 1, 19, 10, 21, 9, 39, 8, 47, 24, 34, 30, 50, 51, 16]).
- **Noisy leakage:** A drawback of the bounded leakage model is that physical leakage rarely obeys to the restriction that the length of the leaked information is a priori bounded (e.g.,

a power trace could be much longer than the secret key). A milder restriction (which weakens the above) is to just assume that the total amount of leaked information does not reduce the entropy of the secret key by too much. Leakage-resilient primitives in this model include one-way relations, public-key encryption, and signature schemes [49, 21, 34].

The focus of this paper is on constructing leakage-resilient signatures in the bounded leakage and noisy leakage model, where one demands that a signature scheme remains unforgeable even against an adversary leaking arbitrary (yet restricted as above) information on the signing key and the overall randomness used within the life-time of the system (this flavour is sometimes known as *fully* leakage resilience). We give a more in-depth description of our contributions in Section 1.1; an overview of our techniques can be found in Section 1.2.

1.1 Our Contribution

Note that in order for a signature scheme to remain existentially unforgeable in the bounded leakage model, it must necessarily be the case that the length of a signature is larger than the length of the secret key (as otherwise an adversary could simply leak a forgery). A first consequence of this is that signatures are very long, as the goal is to enlarge the secret key to tolerate more and more leakage, which is impractical. A second consequence is that we cannot make any meaningful security statement (w.r.t. security in the bounded leakage model) for schemes where the size of the secret key is much larger than the size of a single signature. One remarkable such case is the setting of the Bounded Retrieval Model [15, 26, 27] (BRM), where one intentionally inflates the size of the secret key while keeping constant the size of a signature and the verification key, as well as the computational complexity of the scheme (w.r.t. signature computation/verification).

A similar concern applies to the noisy leakage model, for those schemes where signatures (statistically) reveal little information about the secret key. In such cases leaking a forgery is, again, a valid leakage query, as a signature does not decrease the uncertainty of the secret key by too much. Still, we would like to not consider a scheme completely insecure if the adversary cannot do better than that.

A first step towards addressing the above issues, has recently been taken by Nielsen *et al.* [51] (for the bounded leakage model) who introduced a “graceful degradation” property, essentially requiring that an adversary should not be able to produce more forgeries than what he could have leaked via leakage queries. More precisely, in order to break unforgeability, an adversary has to produce n forgeries where $n \approx \lambda/(\gamma \cdot s) + 1$ for signatures of size s , a total of λ bits of leakage, and a “slack parameter” $\gamma \in (0, 1]$ measuring how close to optimal security a scheme is. The main advantage is that one can design schemes where the size of the secret key is independent of the signature size, leading to shorter signatures. Moreover, this flavour of leakage resilience still allows for interesting applications (e.g., to leaky identification [51]).

New definitions. We start by generalizing the above graceful degradation property to the setting of fully-leakage resilience (both in the bounded and noisy leakage model). Our main notion, dubbed fully-leakage one-more unforgeability, is essentially the same as the one of [51], with the twist that leakage functions can be applied to the entire state of the signer (both for noisy and length-bounded leakage).

We also establish a “middle-ground” notion, which models a setting where secure erasures of the state are available. In particular, we imagine a situation in which the random coins sampled by the signer are completely erased after each invocation of the signing algorithm. Note that in this setting the leakage can essentially depend only on the secret key and the random coins used to compute a single signature. While requiring perfect erasure is a strong assumption (see, e.g., [12]), we believe our notion might still make sense for some applications, as it in particular allows to design simpler and more efficient schemes.

| Scheme | Model | | | G. D. | Assumption | Efficiency | | |
|-----------------------|---------|-------|-------|------------------------|------------|--------------------------|---|---|
| | bounded | noisy | fully | | | leak | $\frac{ wk }{ sk }$ | $\frac{ \sigma }{ sk }$ |
| DHLW10[22] | ✓ | ✗ | ✗ | — | DLIN, SXDH | $1 - o(1)$ | 1 | 9 |
| BSW11[8] | ✓ | ✗ | ✓ | — | DLIN, SXDH | $1 - o(1)$ | $\omega(\log^2 \kappa) \log^2 q_s$ | $\omega(\log^2 \kappa) \log^2 q_s$ |
| MTVY11[47] | ✓ | ✗ | ✓ | — | SXDH | $1 - o(1)$ | $\omega(\log \kappa)$ | $O(1)$ |
| GJS11[34]* | ✓ | ✓ | ✓ | — | SXDH | $1 - o(1)$ | $\Omega\left(\frac{\kappa+\psi}{\psi}\right)$ | $\Omega\left(\frac{ C \kappa}{\psi}\right)$ |
| NVZ14[51] | ✓ | ✗ | ✗ | $\gamma = O(1)$ | DLIN | $\frac{1}{2} - o(1)$ | $\frac{1}{2}$ | $\Omega(1/\psi)$ |
| SS_1^\dagger (§ 4) | ✓ | ✗ | ✓ | $\gamma = O(1/\kappa)$ | DLIN | $1 - \epsilon$ | $(\kappa + 12 + \epsilon^{-1})/\psi$ | $(24\epsilon^{-1} + 90)/\psi$ |
| SS_2^\dagger (§ 4) | ✓ | ✗ | ✓ | $\gamma = O(1/\kappa)$ | DLIN | $1 - \epsilon$ | $(\kappa + 30 + \epsilon^{-1})/\psi$ | $(15\epsilon^{-1} + 133)/\psi$ |
| SS_3 (§ 5) | ✓ | ✓ | ✓ | $\gamma = O(1/q_s)$ | DLIN | $1 - \epsilon$ | $(\kappa + 12 + \epsilon^{-1})/\psi$ | $(15 \log \kappa \epsilon^{-1} + 13)/\psi$ |
| ADW09[2] [‡] | ✓ | ✗ | ✗ | ✗ | BDH | $\frac{1}{2} - \epsilon$ | $\Omega(\epsilon^{-1}/\psi)$ | $\Omega(\epsilon^{-1}/\psi)$ |
| SS_5^\ddagger (§ 6) | ✓ | ✓ | ✓ | $\gamma = O(1)$ | BDH | $1 - \epsilon$ | $\Omega(\log \kappa \epsilon^{-1}/\psi)$ | $\Omega(\log \kappa \epsilon^{-1}/\psi)$ |

Table 1: Comparison of known efficient leakage-resilient signature schemes in the bounded and noisy leakage model. Top: Standard model schemes. Bottom: Schemes in the BRM. When there are multiple instantiations we considered the most efficient one. The parameter ψ is the number of group elements in the secret key; κ is the security parameter, γ is the slack parameter, and q_s the number of signature queries done by the adversary. The \dagger symbol means the scheme relies on memory erasures. The \ddagger symbol means the scheme is in the random oracle model; [2] achieved the weaker notion of entropic unforgeability. G.D. stands for graceful degradation. For the scheme * we wrote $|C|$ for the size of the boolean circuit representing the NP-relation used within the scheme.

New generic constructions. Next, we present new constructions of fully leakage-resilient signature schemes based on generic cryptographic building blocks, improving over previous work in several directions. All of our schemes tolerate leakage on the entire state of the signer, up to a $1 - o(1)$ fraction of the secret key length (in the bounded leakage model). They also offer graceful degradation, allowing to have short signatures of size independent of the size of the secret key.

The first three schemes are secure in the standard model (the first two assuming perfect erasures, the third one not). Our fourth scheme is secure in the BRM, but requires a random oracle. Some of our schemes additionally maintain security in the noisy leakage model. We refer the reader to Table 1.1 for a more detailed overview and comparison with previous works.

Concrete instantiations. We finally explain how to instantiate all building blocks required for our generic schemes. All our instantiations are efficient, and rely on fairly standard assumptions. We refer the reader to Section 7 for the details.

1.2 Technical Overview

We proceed with a high level description of the rationale behind our signatures constructions. For simplicity, we focus here on the bounded leakage model and refer the reader to Section 8 for the extension to the setting of noisy leakage.

Schemes template. All of our schemes follow a common template, which is inspired by the construction of [51]. The main ingredients are: (1) A hybrid¹ linearly homomorphic non-interactive commitment scheme; (2) A statistical non-interactive witness indistinguishable (NIWI) argument system. The secret key consists of two random degree d polynomials, $\delta(X)$ and $r(X)$, over a finite field \mathbb{F} . The verification key includes $t + 1$ verification keys for the com-

¹Intuitively this is a commitment scheme which can be either unconditionally binding or equivocable, depending on the distribution of the public parameters; moreover these two distributions of the parameters are computationally indistinguishable.

mitment scheme, which we denote by $(\vartheta, \{\vartheta_i\}_{i=1}^t)$, and commitments com_i to the coefficients δ_i of $\delta(X)$ computed using the coefficients r_i of $r(X)$ as randomness.

To sign a message m one computes $\tilde{m} = \delta(m)$, and commits to it in two different ways: (1) Using verification key ϑ and randomness $\tilde{r} = r(m)$ (call the result $c\tilde{m}$); (2) Using another verification key $\bar{\vartheta}$ (depending on $\{\vartheta_i\}_{i=1}^t$ and on fresh randomness) and uniform randomness \bar{r} (call the result $c\bar{m}$).² A signature then consists of the commitment $c\tilde{m}$ together with a NIWI argument π that (\tilde{m}, \tilde{r}) and (\tilde{m}, \bar{r}) are valid openings of, respectively, $c\tilde{m}$ and $c\bar{m}$. In order to verify a signature, the verifier can compute the value $c\bar{m}$ non-interactively (using the fact that the commitment scheme is linearly homomorphic), and hence verify the argument π .

Proof outline. To prove security, we define a series of computationally indistinguishable hybrids experiments, where, in each experiment, we either modify the way keys are sampled, or the way signature/leakage queries are handled. The goal is to reach a final hybrid in which signature queries *on average* do not reveal information about the secret polynomial $\delta(X)$. In order to achieve this, we set-up the keys in such a way that the verification key ϑ is always equivocable, whereas the verification key $\bar{\vartheta}$ is binding with some probability p (and equivocable with probability $1 - p$); let us call **Bad** the event that $\bar{\vartheta}$ is binding. In case **Bad** does not happen, we can simply commit to 0 and later equivocate the commitment to sample consistent randomness to simulate a signature/leakage query. In case **Bad** happens, we will be forced to ask an additional leakage query and reveal the value $\tilde{m} = \delta(m)$, which allows to answer a signature query correctly.

In the second part of the proof we analyse the uncertainty of $\delta(X)$ given the view of the adversary in the last hybrid, and we show two facts: (1) The conditional min-entropy of $\delta(X)$ is high with high probability; (2) There exists a predictor able to recover $\delta(X)$ with probability depending on the advantage of the adversary in the unforgeability game. Finally one observes that fact (2) contradicts fact (1) and, whenever the advantage is noticeable, the probability of fact (2) is high, which cannot be true, since also the probability of fact (1) is, and thus the scheme has to be secure.

This is the most critical part of the proof, as in particular, we cannot show directly that the conditional *average* min-entropy of the random variable $\delta(X)$ is high.³ To get around this, we use a careful averaging argument and set the slack parameter γ appropriately.

Amplifying the leakage rate. Any scheme following the above template can tolerate leakage of at most a $1/2 - o(1)$ fraction of the secret key. Our actual constructions achieve optimal leakage rate $1 - o(1)$, by replacing the polynomial $\delta(X)$ with a matrix of polynomials Δ and exploiting a commitment scheme with a message space of “higher dimension”.

1.3 Related Work

On a high level, our scheme follows the template given in [42]; our techniques are mostly related to the ones in [8, 47]. Below, we first briefly describe the standard model scheme of [42], then we provide a more detailed comparison of our schemes and [8, 47]. Finally, we discuss other leakage models and constructions therein.

The signature scheme of Katz and Vaikuntanathan [42] relies on a second-preimage resistant (SPR) function, a CPA-secure public-key encryption scheme and an unbounded simulation-sound NIZK proof system. The secret key is a random pre-image x of a value y under the SPR function (y is part of the public key); a signature is an encryption c of x together with a proof that the ciphertext c is indeed an encryption of a pre-image of y (the message m is used as label in the proof). The scheme does not achieve fully leakage resilience.

²The actual way $\bar{\vartheta}$ is computed varies from scheme to scheme.

³In fact it can be arbitrarily small, due to the presence of the event **Bad**.

Comparison to [8] and [47]. The first fully leakage-resilient signature schemes in the standard model, were constructed independently by [8] and [47]. In a nutshell, the signature scheme of Boyle *et al.* [8] instantiates the above paradigm using a *lossy* tag-based public-key encryption scheme and an admissible hash function H [5]. A signature for a message m is an encryption c under the tag $H(m)$ of a pre-image for an SPR function, together with a statistical NIWI argument that the ciphertext c is indeed an encryption of a pre-image of y . The lossy public-key encryption scheme has the property that for many tags the corresponding ciphertexts are lossy (i.e., they do not leak information about the plaintext), while for a few other tags the corresponding ciphertexts are correctly decryptable. The admissible hash function allows to partition lossy tags from non-lossy tags. With noticeable probability all the signature queries will correspond to lossy tags, therefore signatures do not provide any information about the secret key *even if the randomness of the NIWI is leaked*, while the forgery corresponds to a non-lossy tag, and therefore the simulator can extract a pre-image of y that w.h.p. will be different from the one initially sampled.

The signature scheme of [47] can be seen as an alternative instantiation of the same partitioning strategy. The lossy public-key encryption is removed from the construction, by observing that Groth-Sahai NIZKs are either statistically witness indistinguishable or extractable, depending on the choice of the common reference string (CRS). The admissible function is replaced by the Waters hash function [60], that provides a way to map the message space to a CRS. The Waters function allows to show that with noticeable probability all signature queries correspond to a witness indistinguishable CRS while the forgery corresponds to an extractable CRS.

We observe that the approach of [47] cannot work in case of one-more unforgeability, due to the impossibility result of [40]; this is because finding two distinct messages that are mapped into two extractable CRS would imply that the discrete logarithm problem is easy in the group. Using a similar partitioning argument as the one in [8] might work, however the signature size is proportional to the output of the hash function that grows like $\log^2 \kappa \cdot \log^2 q_s$, where q_s is the number of signature queries.

Other leakage models. Leakage-resilient signatures are constructed also in other leakage models. The “only computation leaks” model [48], assumes that the only part of the state subject to leakage are the ones involved in the computation. Leakage-resilient signatures in this model were constructed in [29]. One important limitation of the “only computation leaks information” axiom, is that it does not captures important attacks in which inactive values in memory are still vulnerable [38]. Note that our assumption of perfect erasures might be interpreted as a special case of only-computation leaks, as a deleted value is never used in a computation again. The model of hard-to-invert leakage [23, 20] (a.k.a. auxiliary input model) generalizes both the bounded and noisy leakage model by allowing the adversary to see computationally hard-to-invert functions of the secret key. The only signature schemes secure in this model are the ones from [28]. In the continual leakage model [22, 10, 41, 34] one can additionally refresh the secret key without changing the verification key of the signature scheme, while the scheme tolerates arbitrary (bounded or noisy) leakage on the entire state between two consecutive invocations of the refreshing algorithms.

A different, but related, line of research recently proposed leakage models that better cope with the perspective of cryptographic engineering (see, e.g., [58, 59, 33]).

1.4 Roadmap

We begin with some notation and basic definitions in Section 2. Our definitions for fully-leakage one-more unforgeability in the bounded leakage model can be found in Section 3. Our generic constructions are given in Section 4 (for the erasure case), Section 5 (for the non-erasure case), and Section 6 (for the BRM). Section 7 discusses how one can instantiate our constructions

under standard number-theoretic assumptions, and comments on efficiency. In Section 8 we explain how to leverage our models and constructions to the noisy leakage model. Finally, we state a few open problems and interesting directions for future research in Section 9.

2 Preliminaries

2.1 Notation

For $a, b \in \mathbb{R}$, we let $[a, b] = \{x \in \mathbb{R} ; a \leq x \leq b\}$; for $a \in \mathbb{N}$ we let $[a] = \{1, 2, \dots, a\}$. If x is a string, we denote its length by $|x|$; if \mathcal{X} is a set, $|\mathcal{X}|$ represents the number of elements in \mathcal{X} . When x is chosen randomly in \mathcal{X} , we write $x \leftarrow \mathcal{X}$. When A is an algorithm, we write $y \leftarrow A(x)$ to denote a run of A on input x and output y ; if A is randomized, then y is a random variable and $A(x; r)$ denotes a run of A on input x and randomness r . An algorithm A is *probabilistic polynomial-time* (PPT) if A is randomized and for any input $x, r \in \{0, 1\}^*$ the computation of $A(x; r)$ terminates in at most $\text{poly}(|x|)$ steps.

Throughout the paper we let κ denote the security parameter. We say that a function $\nu : \mathbb{N} \rightarrow \mathbb{R}$ is negligible in the security parameter κ if $\nu(\kappa) = \kappa^{-\omega(1)}$. A function that is not negligible is called noticeable. For two ensembles $\mathcal{X} = \{X_\kappa\}_{\kappa \in \mathbb{N}}$ and $\mathcal{Y} = \{Y_\kappa\}_{\kappa \in \mathbb{N}}$, we write $\mathcal{X} \equiv \mathcal{Y}$ if they are identically distributed and $\mathcal{X} \approx \mathcal{Y}$ to denote that the two distributions are statistically or computationally close.

Vectors and matrices are typeset in boldface. For a vector $\mathbf{v} = (v_1, \dots, v_n)$ we sometimes write $\mathbf{v}[i]$ for the i -th element of \mathbf{v} .

2.2 Random Variables and Min-Entropy

The min-entropy of a random variable X over a set \mathcal{X} is $\mathbb{H}_\infty(X) := -\log \max_x \mathbb{P}[X = x]$ and represents the best chance of guessing X by an unbounded adversary. Conditional average min-entropy captures how hard it is to guess X on average, given some side information Z (possibly related to X), and it is denoted as $\widetilde{\mathbb{H}}_\infty(X|Z) := -\log \mathbb{E}_z [\max_x \mathbb{P}[X = x|Z = z]]$.

We state a lemma from [42]:

Lemma 2.1. *Let X be a random variable with $H := \mathbb{H}_\infty(X)$, and fix $\Delta \in [0, H]$. Let f be a function whose range has size 2^λ , and set $Y := \{y \in \{0, 1\}^\lambda \mid \mathbb{H}_\infty(X|f(X) = y) \leq H - \Delta\}$. Then*

$$\mathbb{P}[f(X) \in Y] \leq 2^{\lambda - \Delta}.$$

2.3 Commitment Schemes

A (non-interactive) commitment scheme \mathcal{COM} is a tuple of algorithms (**Setup**, **Commit**), defined as follows: (1) Algorithm **Setup** takes as input the security parameter and outputs a verification key ϑ ; (2) Algorithm **Commit** takes as input a message $m \in \mathcal{M}$, randomness $r \in \mathcal{R}$, the verification key ϑ and outputs a value $com \in \mathcal{C}$. To open a commitment com we output (m, r) ; an opening is valid if and only if $com = \text{Commit}(\vartheta, m; r)$.

A commitment scheme has two properties, known as binding and hiding:

Binding Property. Consider the following probability:

$$\mathbb{P}[\text{Commit}(\vartheta, m_0; r_0) = \text{Commit}(\vartheta, m_1; r_1) : \vartheta \leftarrow \text{Setup}(1^\kappa); ((m_0, r_0), (m_1, r_1)) \leftarrow A(\vartheta)].$$

A commitment scheme is *computationally* binding in case the above is negligible for all PPT adversaries A . In case the probability is 0, for all even unbounded A , the commitment scheme is called *perfectly* binding.

Hiding Property. For all messages $m_0, m_1 \in \mathcal{M}$, we have that

$$\{\vartheta, \text{Commit}(\vartheta, m_0)\}_{\kappa \in \mathbb{N}} \approx \{\vartheta, \text{Commit}(\vartheta, m_1)\}_{\kappa \in \mathbb{N}},$$

where the two ensembles are considered as random variables over the choice of the randomness to generate $\vartheta \leftarrow \text{Setup}(1^\kappa)$ and to compute the commitment. A commitment scheme is *statistically* (resp., *computationally*) hiding in case \approx refers to statistical (resp., computational) indistinguishability. In case the two ensembles are identically distributed, the commitment scheme is called *perfectly* hiding.

Whenever \mathcal{M} and \mathcal{R} are a finite field \mathbb{F} , we say that \mathcal{COM} is *linearly homomorphic* in the following sense: Given commitments com and com' and a field element $c \in \mathbb{F}$, one can compute commitments com^* and com'' such that being able to open com and com' to m and m' (respectively) allows to open com^* to $m + m'$ and com'' to $c \cdot m$. We will write the mapping $(com, com') \mapsto com^*$ as $com \cdot com'$ and the mapping $(c, com) \mapsto com''$ as com^c . Similarly, for the opening information we will write the mappings as $com^* = \text{Commit}(\vartheta, m + m'; r + r')$ and $com'' = \text{Commit}(\vartheta, c \cdot m; c \cdot r)$. The above can be generalized to abstract operations over the spaces \mathcal{M} , \mathcal{R} and \mathcal{C} , but for simplicity, and to be consistent with the concrete instantiation given in this paper, we stick to the above formulation here.

Hybrid commitments. A *hybrid* [13] commitment scheme $\mathcal{COM} = (\text{Setup}, \text{Commit}, \text{ESetup}, \text{ECommit}, \text{Equiv})$ is a tuple of algorithms specified as follows: (1) $(\text{Setup}, \text{Commit})$ is a non-interactive commitment scheme with message space \mathcal{M} , randomness space \mathcal{R} , and commitment space \mathcal{C} ; (2) Algorithm ESetup takes as input the security parameter and outputs a pair $(\vartheta, \tau) \leftarrow \text{ESetup}(1^\kappa)$; (3) Algorithm ECommit takes as input a pair (ϑ, τ) and outputs a pair $(com, r') \leftarrow \text{ECommit}(\vartheta, \tau)$; (4) Algorithm Equiv takes as input (τ, m, r') and outputs $r \leftarrow \text{Equiv}(\tau, m, r')$.

Definition 2.1. We say that $\mathcal{COM} = (\text{Setup}, \text{Commit}, \text{ESetup}, \text{ECommit}, \text{Equiv})$ is a hybrid commitment scheme if the following holds.

Perfectly Binding: The tuple $(\text{Setup}, \text{Commit})$ forms a perfectly binding commitment scheme;

Trapdoor Hiding: The tuple $(\text{ESetup}, \text{Commit}, \text{ECommit}, \text{Equiv})$ forms a trapdoor commitment scheme, namely for all $(\vartheta, \tau) \leftarrow \text{ESetup}(1^\kappa)$ and for all $m \in \mathcal{M}$ the following probability distributions are indistinguishable:

$$\left\{ (com, r) \left| \begin{array}{l} r \leftarrow \mathcal{R}, \\ c := \text{Commit}(\vartheta, m; r) \end{array} \right. \right\} \text{ and } \left\{ (com, r) \left| \begin{array}{l} (com, r') \leftarrow \text{ECommit}(\vartheta, \tau), \\ r := \text{Equiv}(\tau, m, r') \end{array} \right. \right\}$$

Hybridness: The following probability distributions are computationally indistinguishable:

$$\{\vartheta \mid (\vartheta, \tau) \leftarrow \text{ESetup}(1^\kappa)\} \text{ and } \{\vartheta \mid \vartheta \leftarrow \text{Setup}(1^\kappa)\}.$$

We call a verification key *equivocable* (resp. *binding*) if it is generated by the algorithm ESetup (resp. Setup). In the same way a commitment is *equivocable* (resp. *binding*) if it is generated using an equivocable (resp. binding) verification key.

For a hybrid linearly homomorphic commitment scheme we will require the following additional property. Let $(\vartheta, \tau) \leftarrow \text{ESetup}(1^\kappa)$, $(com_1, r'_1) \leftarrow \text{ECommit}(\vartheta, \tau)$ and $(com_2, r'_2) \leftarrow \text{ECommit}(\vartheta, \tau)$. Then we can use randomness $r'_1 + r'_2$ to equivocate $com_1 \cdot com_2$, and randomness $c \cdot r'_1$ to equivocate com_1^c (for any c).

Consider the following experiment $\text{Exp}_{\mathcal{COM}, \mathbf{A}}^{\text{comp-hiding}}(\kappa)$ running with a PPT adversary \mathbf{A} and parametrized by the security parameter $\kappa \in \mathbb{N}$:

1. $(m_0, m_1) \leftarrow \mathbf{A}(\vartheta)$ where $\vartheta \leftarrow \text{Setup}(1^\kappa)$;
2. $com \leftarrow \text{Commit}(\vartheta, m_b)$ where $b \leftarrow \{0, 1\}$;

3. Let $b' \leftarrow_{\$} A(\vartheta, com)$, output 1 iff $b = b'$.

Lemma 2.2 (Adaptive Computational Hiding). *Let \mathcal{COM} be a hybrid commitment scheme. For any PPT A there exists a negligible function negl such that:*

$$\left| \mathbb{P} \left[\text{Exp}_{\mathcal{COM}, A}^{\text{comp-hiding}}(\kappa) = 1 \right] - 1/2 \right| \leq \text{negl}(\kappa).$$

Proof. The trapdoor hiding property implies that for any $(\vartheta, \tau) \leftarrow_{\$} \text{ESetup}(1^\kappa)$ and any $(m_0, m_1) \in \mathcal{M}$ the following two probability distributions are indistinguishable:

$$\{com \mid com \leftarrow_{\$} \text{Commit}(\vartheta, m_0)\} \text{ and } \{com \mid com \leftarrow_{\$} \text{Commit}(\vartheta, m_1)\}. \quad (1)$$

Since Eq. (1) holds for any ϑ and any m_0, m_1 we can allow the messages to be chosen adaptively after seeing the verification key. Hence, for any $(\vartheta, \tau) \leftarrow_{\$} \text{ESetup}(1^\kappa)$ and for any A :

$$\left| \mathbb{P} \left[b = b' \mid \begin{array}{l} m_0, m_1 \leftarrow_{\$} A(\vartheta); b \leftarrow_{\$} \{0, 1\}; \\ com \leftarrow_{\$} \text{Commit}(\vartheta, m_b); b' \leftarrow_{\$} A(\vartheta, com); \end{array} \right] - 1/2 \right| \leq \text{negl}(\kappa). \quad (2)$$

The hybridness property and Eq. (2) imply the lemma. \square

2.4 NIWI and NIZK Arguments

For a relation $\mathfrak{R} \subseteq \{0, 1\}^* \times \{0, 1\}^*$, the language associated with \mathfrak{R} is $\mathcal{L}_{\mathfrak{R}} = \{x : \exists w \text{ s.t. } (x, w) \in \mathfrak{R}\}$. A non-interactive argument system $(\text{Init}, \text{Prove}, \text{Ver})$ is a tuple of algorithms specified as follows: (1) Init takes as input the security parameter and outputs a common reference string $\text{crs} \leftarrow \text{Init}(1^\kappa)$; (2) Prove takes as input a pair $(x, w) \in \mathfrak{R}$, and outputs a proof $\pi \leftarrow \text{Prove}(\text{crs}, x, w)$; (3) Ver takes as input a statement x and a proof π , and outputs a bit $b \leftarrow \text{Ver}(\text{crs}, x, \pi)$. We require that for all pairs $(x, w) \in \mathfrak{R}$, and all $\text{crs} \leftarrow \text{Init}(1^\kappa)$, the following holds: $\mathbb{P}[\text{Ver}(\text{crs}, x, \text{Prove}(\text{crs}, x, w))] \geq 1 - \text{negl}(\kappa)$ (over the randomness of the prover and verifier algorithms).

We review the notions of non-interactive witness indistinguishable and zero-knowledge argument systems for a relation \mathfrak{R} .

NIWI arguments. A non-interactive witness indistinguishable argument system satisfies two properties known as adaptive soundness and statistical witness indistinguishability.

Definition 2.2 (Adaptive soundness). *Let \mathcal{NIWI} be an argument system for a language \mathcal{L} . We say that \mathcal{NIWI} satisfies adaptive soundness, if for all PPT adversaries A there exists a negligible function $\text{negl}(\cdot)$ such that*

$$\mathbb{P}[\text{Ver}(\text{crs}, x, \pi) = 1 \wedge x \notin \mathcal{L} : \text{crs} \leftarrow \text{Init}(1^\kappa), (x, \pi) \leftarrow A(1^\kappa, \text{crs})] \leq \text{negl}(\kappa).$$

Definition 2.3 (Statistical witness indistinguishability). *Let \mathcal{NIWI} be an argument system for a relation \mathfrak{R} . We say that \mathcal{NIWI} satisfies statistical witness indistinguishability if for any triplet (x, w, w') such that $(x, w) \in \mathfrak{R}$ and $(x, w') \in \mathfrak{R}$, the distributions $\{(\text{crs}, \pi) \mid \text{crs} \leftarrow \text{Init}(1^\kappa), \pi \leftarrow \text{Prove}(\text{crs}, x, w)\}$ and $\{(\text{crs}, \pi) \mid \text{crs} \leftarrow \text{Init}(1^\kappa), \pi \leftarrow \text{Prove}(\text{crs}, x, w')\}$ are statistically indistinguishable.*

NIZK arguments. We will need a non-interactive zero-knowledge argument system satisfying two properties, known as adaptive multi-theorem zero-knowledge and true-simulation extractability [22].

Definition 2.4 (Adaptive multi-theorem zero-knowledge). *Let \mathcal{NIZK} be an argument system for a relation \mathfrak{R} . We say that \mathcal{NIZK} satisfies adaptive multi-theorem zero-knowledge with labels if the following holds:*

- (i) There exists an algorithm $\overline{\text{Init}}$ that outputs a CRS crs , and a simulation trapdoor τ_{sim} .
- (ii) There exists a PPT simulator Sim such that, for all PPT adversaries \mathbf{A} , we have that

$$\left| \mathbb{P} \left[\mathbf{A}^{\text{Prove}(\text{crs}, \cdot)}(\text{crs}) = 1 \mid \text{crs} \leftarrow \overline{\text{Init}}(1^\kappa) \right] - \mathbb{P} \left[\mathbf{A}^{\text{SIM}(\tau_{\text{sim}}, \cdot)}(\text{crs}) = 1 \mid (\text{crs}, \tau_{\text{sim}}) \leftarrow \overline{\text{Init}}(1^\kappa) \right] \right|$$

is negligible in κ . The simulation oracle $\text{SIM}(\tau_{\text{sim}}, \cdot)$ takes as input a pair (x, w) and a label L , checks if $(x, w) \in \mathfrak{R}$, and, if true, ignores w and outputs a simulated argument $\text{Sim}(\tau_{\text{sim}}, (x, L))$, and otherwise outputs \perp .

Definition 2.5 (True-simulation extractability). Let \mathcal{NIZK} be an argument system for a relation \mathfrak{R} . We say that \mathcal{NIZK} is true-simulation extractable (tSE) with labels if the following holds:

- (i) There exists an algorithm $\overline{\text{Init}}$ that outputs a CRS crs , a simulation trapdoor τ_{sim} , and an extraction trapdoor τ_{ext} .
- (ii) There exists a PPT algorithm $\mathbf{K}(\tau_{\text{ext}}, (x, L), \pi)$ such that, for all PPT adversaries \mathbf{A} , we have $\mathbb{P}[\mathbf{A} \text{ wins}] \leq \text{negl}(\kappa)$ in the following game:
- The challenger runs $(\text{crs}, \tau_{\text{sim}}, \tau_{\text{ext}}) \leftarrow \overline{\text{Init}}(1^\kappa)$, and gives crs to \mathbf{A} .
 - \mathbf{A} is given access to the simulation oracle SIM , which it can access adaptively.
 - \mathbf{A} outputs a tuple $((x^*, L^*), \pi^*)$.
 - The challenger runs $w \leftarrow \mathbf{K}(\tau_{\text{ext}}, (x^*, L^*), \pi)$. We say that \mathbf{A} wins if: (a) the pair (x^*, L^*) was not part of a query to the simulation oracle; (b) $\text{Ver}(\text{crs}, (x^*, L^*), \pi^*) = 1$; and (c) $(x^*, w) \notin \mathfrak{R}$.

Oblivious sampling. For our application we will need the additional property that the Init algorithm supports oblivious sampling (a.k.a. reference string uniformity in [56]), meaning that the distribution of the common reference string is statistically-close to the uniform distribution. Such property will be satisfied in our concrete instantiations based on the Groth-Sahai proof system [36].

3 Fully-Leakage One-More Unforgeability

A signature scheme is a triple of algorithms $\mathcal{SS} = (\text{KGen}, \text{Sign}, \text{Verify})$ defined as follows: (1) The key generation algorithm takes as input the security parameter κ and outputs a verification key/signing key pair $(vk, sk) \leftarrow \text{KGen}(1^\kappa)$; (2) The signing algorithm takes as input a message $m \in \mathcal{M}$ and the signing key sk and outputs a signature $\sigma \leftarrow \text{Sign}(sk, m)$; (3) The verification algorithm takes as input the verification key vk and a pair (m, σ) and outputs a bit $\text{Verify}(vk, (m, \sigma)) \in \{0, 1\}$. We say that \mathcal{SS} satisfies correctness if for all messages $m \in \mathcal{M}$ and for all pairs of keys (vk, sk) generated via KGen , we have that $\text{Verify}(vk, (m, \text{Sign}(sk, m)))$ returns 1 with overwhelming probability over the randomness of the signing algorithm.

In this section we extend the notion of one-more unforgeability [51] to a setting where all intermediate values generated within the lifetime of the system (and not just the secret key) are subject to leakage. As outlined in the introduction, we consider two cases depending whether one is willing to assume perfect erasures or not.

3.1 Definition without Erasures

Given a signature scheme \mathcal{SS} , consider the following experiment $\text{Exp}_{\mathcal{SS}, \mathbf{A}}^{\text{one-more}}(\kappa, \ell, q_s, \gamma)$ running with a PPT adversary \mathbf{A} and parametrized by the security parameter $\kappa \in \mathbb{N}$, the leakage parameter $\ell \in \mathbb{N}$, and the slack parameter $\gamma := \gamma(\kappa)$:

1. Sample $r_0 \in \{0,1\}^*$, run the key generation algorithm to obtain a pair $(vk, sk) := \text{KGen}(1^\kappa; r_0)$, and return vk to \mathbf{A} ; let $st = \{r_0\}$.
2. The adversary \mathbf{A} can adaptively issue signing queries. Upon input a message $m \in \mathcal{M}$, the adversary is given a signature $\sigma := \text{Sign}(sk, m; r)$ computed using fresh coins r . The state is updated to $st := st \cup \{r\}$.
3. The adversary \mathbf{A} can adaptively issue leakage queries. Upon input an arbitrary function f , the adversary is given $f(st)$ where st is the current state.
4. Let \mathcal{Q} be the set of signing queries issued by \mathbf{A} , and let $\Lambda \in \{0,1\}^\lambda$ be the total amount of information leaked by the adversary. \mathbf{A} outputs n pairs $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$.
5. The experiment outputs 1 if and only if the following conditions are satisfied:
 - (a) $\text{Verify}(vk, (m_i^*, \sigma_i^*)) = 1$ and $m_i^* \notin \mathcal{Q}$, for all $i \in [n]$.
 - (b) The messages m_1^*, \dots, m_n^* are pairwise distinct.
 - (c) $n \geq \lfloor \lambda / (\gamma \cdot s) \rfloor + 1$, where $s := |\sigma|$ is the size of a signature.
 - (d) $\lambda \leq \ell$ and $|\mathcal{Q}| \leq q_s$.

Definition 3.1. We say that $\mathcal{SS} = (\text{KGen}, \text{Sign}, \text{Verify})$ is $(\ell, q_s, \gamma, \varepsilon)$ -fully-leakage one-more unforgeable if for every PPT adversary \mathbf{A} asking q_s signature queries we have that $\mathbb{P}[\text{Exp}_{\mathcal{SS}, \mathbf{A}}^{\text{one-more}}(\kappa, \ell, q_s, \gamma) = 1] \leq \varepsilon$. Whenever $\varepsilon = \text{negl}(\kappa)$, $q_s = \text{poly}(\kappa)$, and γ does not depend on q_s , we simply say that \mathcal{SS} is (ℓ, γ) -fully-leakage one-more unforgeable.

Note that the number of signatures the adversary has to forge depends on the length of the leakage he asks to see. In particular (ℓ, γ) -fully-leakage one-more unforgeability implies standard unforgeability for any adversary asking no leakage ($\lambda = 0$).

As pointed out in [51], the slack parameter γ specifies how close to optimal security \mathcal{SS} is. In particular, in case $\gamma = 1$ one-more unforgeability requires that \mathbf{A} cannot forge even a single signature more than what it could have leaked via leakage queries. As γ decreases, so does the strength of the signature scheme (the extreme case being $\gamma = |\mathcal{M}|^{-1}$, where we have no security).

For some of our schemes, the slack parameter will be dependent on the security parameter, and in some case even on the number of signature queries asked by the adversary. We stress, however, that as long as this dependency is polynomial, one-more unforgeability is still meaningful for some applications (cf. [51, Section 5]).

3.2 Definition with Erasures

In the setting of Definition 3.1 the leakage is allowed to depend on the entire randomness produced within the lifetime of the system. A more restricted (yet natural) setting is to assume that the state is erased after each signature is generated. Although perfect erasures might be hard to implement [12], we believe that such weakening might still be useful in practice. Moreover, looking ahead, the ideas behind the schemes we construct assuming perfect erasures (cf. Section 4) form the core for the scheme which does not consider erasures (cf. Section 5).

To define fully-leakage one-more unforgeability with erasures one simply modifies the experiment $\text{Exp}_{\mathcal{SS}, \mathbf{A}}^{\text{one-more}}(\kappa, \ell, q_s, \gamma)$ by requiring that the state variable st is always equal to sk . Since this is clearly equivalent to a setting where the adversary asks up to one leakage query after each signature query, we can equivalently⁴ give \mathbf{A} access to the following machine $\widetilde{\text{Sign}}(sk, \cdot)$:

⁴It is, in fact, easy to see that an adversary that makes multiple queries between two consecutive signature queries can be reduced to an adversary in our model.

- Upon input a value $m \in \mathcal{M}$, compute $\sigma \leftarrow \text{Sign}(sk, m)$ and return σ . We call such invocation of the signature oracle a *leak-free query*.
- Upon input a pair (m, f) , where $m \in \mathcal{M}$ and f is an arbitrary function, sample randomness $r \leftarrow \{0, 1\}^*$, compute $\sigma := \text{Sign}(sk, m; r)$ and return $(\sigma, f(sk, r))$. We call such invocation of the signature oracle a *leaky query*.

A formal definition follows. Given a signature scheme \mathcal{SS} , consider the following experiment $\widetilde{\text{Exp}}_{\mathcal{SS}, \mathbf{A}}^{\text{one-more}}(\kappa, \ell, q_s, \gamma)$ running with a PPT adversary \mathbf{A} and parametrized by the security parameter $\kappa \in \mathbb{N}$, the leakage parameters $\ell \in \mathbb{N}$, and the slack parameter $\gamma := \gamma(\kappa)$:

1. The adversary can issue an initial leakage query. Upon input an arbitrary function f_0 , sample $r_0 \in \{0, 1\}^*$, run the key generation algorithm to obtain a pair $(vk, sk) := \text{KGen}(1^\kappa; r_0)$, and return $(vk, f_0(r_0))$ to \mathbf{A} .
2. The adversary \mathbf{A} can adaptively access oracle $\widetilde{\text{Sign}}(sk, \cdot)$.
3. Let \mathcal{Q} be the set of signing queries issued by \mathbf{A} , and denote with $\Lambda \in \{0, 1\}^\lambda$ the total amount of information leaked by the adversary, including the initial leakage $\Lambda_0 \in \{0, 1\}^{\lambda_0}$ (with $\lambda \geq \lambda_0$). \mathbf{A} outputs n pairs $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$.
4. The experiment outputs 1 if and only if the following conditions are satisfied:
 - (a) $\text{Verify}(vk, (m_i^*, \sigma_i^*)) = 1$ and $m_i^* \notin \mathcal{Q}$, for all $i \in [n]$.
 - (b) The messages m_1^*, \dots, m_n^* are pairwise distinct.
 - (c) $n \geq \lfloor \lambda / (\gamma \cdot s) \rfloor + 1$, where $s := |\sigma|$ is the size of a signature.
 - (d) $\lambda \leq \ell$, and $|\mathcal{Q}| \leq q_s$.

Definition 3.2. We say that $\mathcal{SS} = (\text{KGen}, \text{Sign}, \text{Verify})$ is $(\ell, q_s, \gamma, \varepsilon)$ -fully-leakage one-more unforgeable with perfect erasures if for every PPT adversary \mathbf{A} we have that $\mathbb{P}[\widetilde{\text{Exp}}_{\mathcal{SS}, \mathbf{A}}^{\text{one-more}}(\kappa, \ell, q_s, \gamma) = 1] \leq \varepsilon$. Whenever $\varepsilon = \text{negl}(\kappa)$, $q_s = \text{poly}(\kappa)$, and γ does not depend on q_s , we simply say that \mathcal{SS} is (ℓ, γ) -fully-leakage one-more unforgeable with perfect erasures.

4 A Scheme with Erasures

4.1 The Basic Construction

We give a construction of a fully-leakage one-more unforgeable signature scheme (cf. Definition 3.2) based on the following building blocks:

- A statistical non-interactive witness-indistinguishable argument system $\mathcal{NWI} = (\text{Init}, \text{Prove}, \text{Ver})$ (cf. Section 2.4).
- A trapdoor hiding, linearly homomorphic commitment $\mathcal{COM}_1 = (\text{ESetup}_1, \text{Commit}_1, \text{ECommit}_1, \text{Equiv}_1)$ equipped with an obviously sampleable key generation algorithm Setup and with message space equal to \mathbb{F}^μ (for a finite field \mathbb{F} and a parameter $\mu \in \mathbb{N}$) and randomness space equal to \mathbb{F} (cf. Section 2.3). Let ρ be defined as $\rho := \frac{\log |\mathbb{F}^\mu|}{\log(|\mathbb{F}^\mu| \times |\mathbb{F}|)} = \frac{\mu}{\mu+1}$.
- A hybrid, linearly homomorphic commitment scheme $\mathcal{COM}_2 = (\text{Setup}_2, \text{Commit}_2, \text{ESetup}_2, \text{ECommit}_2, \text{Equiv}_2)$, with message space equal to \mathbb{F}^μ and randomness space equal to \mathbb{F}^ν for a parameter $\nu \geq \mu$.
- A pseudorandom generator (PRG) $G : \mathcal{X} \rightarrow \mathcal{Y}$.⁵

Our scheme $\mathcal{SS}_1 = (\text{KGen}, \text{Sign}, \text{Verify})$ has message space equal to \mathbb{F} and is described below:

⁵A PRG is a deterministic, efficiently computable, function such that the ensemble $\{G(x) \mid x \in \{0, 1\}^\kappa\}_{\kappa \in \mathbb{N}}$ is pseudorandom.

Key Generation. Let $t, d, \mu, \nu \in \mathbb{N}$ be parameters. Run $\text{crs} \leftarrow \text{Init}(1^\kappa)$, sample a uniform $y \in \mathcal{Y}$ from the range of G , run $\vartheta \leftarrow \text{Setup}_1(1^\kappa)$, and $\vartheta_i \leftarrow \text{Setup}_2(1^\kappa)$ for all $i \in [t]$. Sample $\Delta \leftarrow_{\$} (\mathbb{F}^\mu)^{d+1}$ and $\mathbf{r} = (r_0, \dots, r_d) \leftarrow_{\$} \mathbb{F}^{d+1}$, and compute commitments $\text{com}_i = \text{Commit}_1(\vartheta, \delta_i; r_i)$ for $i \in [0, d]$, where $\delta_i \in \mathbb{F}^\mu$ is the i -th column of Δ . Output

$$sk = (\Delta, \mathbf{r}) \quad vk = (\text{crs}, y, \vartheta, \{\vartheta_i\}_{i=1}^t, \{\text{com}_i\}_{i=0}^d).$$

Signature. For $j \in [\mu]$, let $\delta_j(X)$ be the degree d polynomial having as coefficients the elements in the j -th row of Δ , i.e., $\delta_j(X) := \sum_{i=0}^d \delta_{j,i} \cdot X^i$; similarly let $r(X)$ be the degree d polynomial having as coefficients the elements of \mathbf{r} , i.e., $r(X) := \sum_{i=0}^d r_i \cdot X^i$. Define

$$\Delta(X) := \begin{pmatrix} \delta_1(X) \\ \vdots \\ \delta_\mu(X) \end{pmatrix}.$$

Consider the following polynomial-time relation:

$$\mathfrak{R} := \left\{ (\vartheta, \vartheta', \tilde{com}, \text{com}); (\tilde{m}, \tilde{r}, r) \mid \begin{array}{l} \tilde{com} = \text{Commit}_1(\vartheta, \tilde{m}; \tilde{r}) \\ \text{com} = \text{Commit}_2(\vartheta', \tilde{m}; r) \end{array} \right\}.$$

together with the polynomial-time OR-relation

$$\mathfrak{R}_{OR} := \{(\alpha, y); (\beta, x) \mid (\alpha, \beta) \in \mathfrak{R} \vee G(x) = y\}.$$

To sign a message $m \in \mathbb{F}$ compute $\tilde{m} = \Delta(m)$ and $\tilde{r} = r(m)$, and let $\tilde{com} := \text{Commit}_1(\vartheta, \tilde{m}; \tilde{r})$. Notice that both values \tilde{m}, \tilde{r} can be computed efficiently as a function of the signing key (Δ, \mathbf{r}) and the message to be signed. Pick a random index $j \leftarrow_{\$} [t]$, and compute $\text{com} := \text{Commit}_2(\vartheta_j, \tilde{m}; r)$ where $r \leftarrow_{\$} \mathbb{F}^\nu$. Using crs as common reference string, generate a NIWI argument π for $(\vartheta, \vartheta_j, \tilde{com}, \text{com}, y) \in \mathfrak{L}_{\mathfrak{R}_{OR}}$, the language generated by the above relation \mathfrak{R}_{OR} . Output $\sigma = (\pi, j, \text{com})$.

Verification. Given a pair (m, σ) , parse σ as (π, j, com) . Output the same as:

$$\text{Ver}(\text{crs}, \pi, (\vartheta, \vartheta_j, \prod_{i=0}^d (\text{com}_i)^{m^i}, \text{com}, y)).$$

Let us first argue that the signature scheme satisfies the correctness property. This follows from the correctness of the NIWI argument system, and from the fact that \mathcal{COM} is linearly homomorphic (cf. Section 2.3), as

$$\tilde{com} = \prod_{i=0}^d (\text{com}_i)^{m^i} = \prod_{i=0}^d \text{Commit}_1(\vartheta, \delta_i \cdot m^i; r_i \cdot m^i) = \text{Commit}_1\left(\vartheta, \underbrace{\sum_{i=0}^d \delta_i \cdot m^i}_{\tilde{m}}; \underbrace{\sum_{i=0}^d r_i \cdot m^i}_{\tilde{r}}\right).$$

Theorem 4.1. *Let $\mu \in \mathbb{N}$, and let \mathbb{F} be a finite field of size $|\mathbb{F}| = \kappa$ for security parameter $\kappa \in \mathbb{N}$. For any $0 < \xi \leq \frac{\mu}{\mu+1}$ and $t = O(\kappa)$, whenever $\gamma = O(1/\kappa)$ the above signature scheme is $((\frac{\mu}{\mu+1} - \xi)|sk|, \gamma)$ -fully-leakage one-more unforgeable with perfect erasures.*

Notice that the slack parameter γ is not optimal, as $\gamma = O(1/\kappa)$. However, as pointed out in Section 3, this is fine for some applications. The proof of Theorem 4.1 proceeds in two steps. Let \mathbf{A} be a PPT adversary breaking fully-leakage one-more unforgeability with probability $\varepsilon(\kappa)$. In the first step we define a series of hybrid experiments that are indistinguishable from the original experiment defining security of the signature scheme. The purpose of these hybrids is to reach a state in which signature queries reveal (information theoretically) as little information

as possible on the secret matrix Δ . In the second step we analyse the uncertainty of Δ given the adversary's view; in particular we show two facts: (1) The conditional min-entropy of Δ is high with high probability; (2) There exists a predictor able to recover Δ with probability depending on $\varepsilon(\kappa)$. Finally one observes that fact (2) contradicts fact (1) and, whenever the advantage is noticeable, the probability of fact (2) is high, which cannot be true, since also the probability of fact (1) is, and thus the scheme has to be secure.

The main technical challenge to overcome is that although simulated signatures should not leak information about the secret key, the forgeries output by the adversary need to. It is easy to see that these two requirements are in contrast. To overcome this difficulty we construct the simulator in a way that, on average, simulated signatures do not provide information about Δ , meaning that sometimes the simulator will have to retrieve and reveal some information about the secret key. We obtain a contradiction as soon as the forgeries reveal more information about the secret key than the information that the adversary has leaked. Setting the slack parameter to $O(1/\kappa)$ provides us with enough signatures to overcome the bound set by the leakage parameter and reach a contradiction.

4.2 Proof of Theorem 4.1

Let A be an adversary making $q = \text{poly}(\kappa)$ *leaky* signature queries, such that

$$\mathbb{P}[\widetilde{\text{Exp}}_{\mathcal{SS}_1, A}^{\text{one-more}}(\kappa, \ell, q_s, \gamma) = 1] = \varepsilon(\kappa).$$

Experiment \mathcal{H}_0 : This experiment is exactly the same as the original experiment defining fully-leakage one-more unforgeability of \mathcal{SS}_1 . Note that the initial state (after key generation), is equal to $st := ((\Delta, \mathbf{r}), \text{crs}, \vartheta, \{\vartheta_i\}_{i=1}^t)$. (Recall that we assume oblivious sampling for both the NIWI and the commitment scheme.) This state is used to answer the initial leakage query f_0 ; afterwards everything but (Δ, \mathbf{r}) is erased. Additionally at each invocation of oracle $\widetilde{\text{Sign}}(sk, \cdot)$ corresponding to a *leaky query* (m, f) , the experiment produces a signature $\sigma = (\pi, j, \text{com})$ following the signing algorithm, and defines the current state to be $st := ((\Delta, \mathbf{r}), r, r_{\text{niwi}})$, where r is the randomness used to generate the commitment com and r_{niwi} is the randomness used to generate the argument π .

Denote with $((m_1^*, (\pi_1^*, j_1^*, \text{com}_1^*)), \dots, (m_n^*, (\pi_n^*, j_n^*, \text{com}_n^*)))$ the set of forgeries returned by the adversary at the end of the experiment. Let Forge_0 be the event that \mathcal{H}_0 returns 1, so that $\mathbb{P}[\text{Forge}_0] = \varepsilon$. Define False_0 to be the event that at least one of the proofs contained in the adversary's forgeries is relative to a false statement, i.e., False_0 is verified if in \mathcal{H}_0 there exists some index $i \in [1, n]$ for which com_i^* is a well-formed commitment to a value $m_i' \neq \Delta(m_i^*)$. Let $\varepsilon_0 := \mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0]$.

Claim 4.1. $\varepsilon_0 \geq \varepsilon - \text{negl}(\kappa)$.

Proof. By adaptive computational soundness of the NIWI argument system, we must have that $\mathbb{P}[\text{False}_0] \leq \text{negl}(\kappa)$. In fact, from an adversary A provoking False_0 , we can easily construct an adversary B breaking adaptive soundness by simply emulating the entire experiment for A and outputting one of the n proofs π_i^* contained in A 's forgeries (chosen uniformly at random). In case $\mathbb{P}[\text{False}_0]$ is noticeable, B breaks soundness with noticeable advantage, so False_0 must happen with negligible probability. Hence,

$$\begin{aligned} \varepsilon_0 &= \mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0] = \mathbb{P}[\text{Forge}_0] - \mathbb{P}[\text{Forge}_0 \wedge \text{False}_0] \\ &\geq \mathbb{P}[\text{Forge}_0] - \mathbb{P}[\text{False}_0] \geq \varepsilon - \text{negl}(\kappa), \end{aligned}$$

as desired. □

We now define a series of hybrid experiments. For each hybrid \mathcal{H}_i , we write ε_i for the probability of the event $\text{Forge}_i \wedge \neg \text{False}_i$, and View_i for the view of the adversary at the end of the experiment. The first sequence of hybrids are made to switch the public key of the signature scheme to a simulated public key. The simulated public key is generated together with a trapdoor, and moreover does not reveal any information about the signing key.

Experiment $\mathcal{H}_{1.a}$: This experiment is the same as \mathcal{H}_0 , except that during key generation the verification key for the commitment scheme \mathcal{COM}_1 is computed using the equivocable setup algorithm ESetup_1 , a random index $i^* \in [t]$ is sampled, and all verification keys for the commitment scheme \mathcal{COM}_2 but the i^* -th are computed using the equivocable setup algorithm ESetup_2 (whereas ϑ_{i^*} is computed as before). More precisely,

$$\begin{cases} (\vartheta, \tau) \leftarrow \text{ESetup}_1(1^\kappa) \\ (\vartheta_i, \tau_i) \leftarrow \text{ESetup}_2(1^\kappa) & \text{for } i \neq i^* \\ \vartheta_i \leftarrow \text{Setup}_2(1^\kappa) & \text{for } i = i^*. \end{cases}$$

Note that from the point of view of the adversary, the initial state remains equal to $st := ((\Delta, \mathbf{r}), \text{crs}, \vartheta, \{\vartheta_i\}_{i=1}^t)$.

Claim 4.2. $\varepsilon_{1.a} \geq \varepsilon_0 - \text{negl}(\kappa)$.

Proof. First note the distributions of Setup_1 and ESetup_1 are the same. The claim follows by a standard hybrid argument relying on the hybridness property of the commitment scheme. \square

Experiment $\mathcal{H}_{1.b}$: This experiment is the same as the previous hybrid, except that now the commitments $\{com_i\}_{i=0}^d$ to the columns δ_i are replaced by equivocable commitments, i.e. $(com_i, r'_i) \leftarrow \text{ECommit}_1(\vartheta, \tau)$ for all $i \in [0, d]$. Notice that the actual randomness r_i , used to produce com_i in $\mathcal{H}_{1.a}$, can be recovered efficiently as a function of the coefficients δ_i and the fake randomness r'_i , as $r_i(\Delta) := \text{Equip}_1(\tau, \delta_i, r'_i)$. Given r_i the signature computation is identical in $\mathcal{H}_{1.a}$ and $\mathcal{H}_{1.b}$; in the following we write $\mathbf{r}(\Delta)$ for the vector $(r_0(\Delta), \dots, r_d(\Delta))$.

Claim 4.3. $\varepsilon_{1.b} \geq \varepsilon_{1.a} - \text{negl}(\kappa)$.

Proof. The trapdoor hiding property of the commitment scheme implies that the distribution of each pair (com_i, r_i) in the two hybrids are statistically close. The claim follows by a standard hybrid argument. \square

Experiment $\mathcal{H}_{1.c}$: This experiment is identical to the previous hybrid, except that the value y contained in the verification key is now computed by sampling a random $x \in \mathcal{X}$ and by outputting $y := G(x)$ (instead of sampling $y \in \mathcal{Y}$ uniformly at random).

Claim 4.4. $\varepsilon_{1.c} \geq \varepsilon_{1.b} - \text{negl}(\kappa)$.

Proof. It follows from the fact that G is a secure pseudorandom generator. \square

The next sequence of hybrids is made in order to replace all signatures returned by the signing oracle to simulated signatures that, on average, do not reveal information about the secret key. Looking ahead, the trapdoor information that allows the above transition to go through consists of the tuple $(i^*, \tau, \{\tau_i\}_{i \neq i^*}, \{r'_i\}_{i=0}^d, x)$, i.e. the index i^* corresponding to the perfectly binding verification key, the trapdoors for all other verification keys, the randomness r'_i used to equivocate the values com_i , and the seed x of the pseudorandom generator.

Experiment $\mathcal{H}_{2.a}$: This hybrid differs from the previous one in the way *leak-free* signature queries are dealt with. Specifically, for each query $m \in \mathbb{F}$ the proof π is computed using x as witness for $y = G(x)$ instead of $\beta = (\tilde{m}, \tilde{r}, r)$ as witness for $\alpha = (\vartheta, \vartheta_j, c\tilde{m}, com) \in \mathcal{L}_{\mathfrak{R}}$.

Claim 4.5. $\varepsilon_{2.a} \geq \varepsilon_{1.c} - \text{negl}(\kappa)$.

Proof. Follows from the statistical witness indistinguishability property of the argument system, via a standard hybrid argument on the number of *leak-free* signature queries. \square

Experiment $\mathcal{H}_{2.b}$: For each *leak-free* signature query of the form $m \in \mathbb{F}$, the experiment checks whether the sampled index $j \leftarrow_{\$} [t]$ is such that $j = i^*$. If this is the case, such query is answered exactly as in the previous hybrid, and otherwise the value com in the signature σ is computed as $(com, r') \leftarrow_{\$} \text{ECommit}_2(\vartheta_j, \tau_j)$.

Claim 4.6. $\varepsilon_{2.b} \geq \varepsilon_{2.a} - \text{negl}(\kappa)$.

Proof. Let $\sigma = (\pi, j, com)$ be the answer to a leak-free signature query $m \in \mathbb{F}$. The distribution of the proof π and of the index j is identical in the two hybrids. The trapdoor hiding property implies that the distribution of the commitment com in hybrid $\mathcal{H}_{2.b}$ is statistically close to the distribution of com in hybrid $\mathcal{H}_{2.a}$. The claim follows. \square

Experiment $\mathcal{H}_{2.c}$: This experiment behaves in the same way as the previous hybrid, with the exception that for each *leak-free* query such that $j = i^*$ the value com is replaced with a commitment to zero, i.e. $com := \text{Commit}_2(\vartheta_{i^*}, 0^\mu; r)$ for $r \leftarrow_{\$} \mathbb{F}^\nu$.

Claim 4.7. $\varepsilon_{2.c} \geq \varepsilon_{2.b} - \text{negl}(\kappa)$.

Proof. Recall that there are three types of signature queries handled by the two hybrids $\mathcal{H}_{2.b}$ and $\mathcal{H}_{2.c}$. Specifically:

1. Both hybrids handle in the same way *leaky* signature queries.
2. In case of a *leak-free* signature query on a message m both hybrids randomly select the index $j \leftarrow_{\$} [t]$ and:
 - (a) if $j \neq i^*$, both hybrids handle the signature in the same way;
 - (b) if $j = i^*$, hybrid $\mathcal{H}_{2.b}$ computes com as an honestly generated commitment to $\Delta(m)$, whereas hybrid $\mathcal{H}_{2.c}$ computes com by committing to zero. We call such a query *binding*.

Let $q' = \text{poly}(\kappa)$ be the total number of *leak-free* signature queries and let $K \in [q']$ be a random variable representing the number of such queries that are binding. Note that in case no signature query happens to be binding the two hybrids are identically distributed, and thus there is nothing to prove. On the other hand, we show that for any possible realization $k \in [q']$ of the random variable K the two hybrids are computationally close. This implies the claim.

Denote with $\mathcal{J} = \{j_1, \dots, j_k\} \subseteq [q']$ the set of indexes corresponding to the binding signature queries. For an index $0 \leq w \leq k$, consider the following series of hybrids experiments. Hybrid $\mathcal{H}_{2.b}^w$ treats the first w binding signature queries (i.e., the ones up to index j_w) as $\mathcal{H}_{2.b}$ does, whereas the remaining queries (i.e., the ones going from index j_{w+1} to j_k) are treated as in $\mathcal{H}_{2.c}$. Note that $\mathcal{H}_{2.c} = \mathcal{H}_{2.b}^0$, and $\mathcal{H}_{2.b} = \mathcal{H}_{2.b}^k$.

By contradiction, assume that there exists an efficient distinguisher D that for some $w \in [0, k]$ can distinguish the distribution of $\mathcal{H}_{2.b}^w$ and $\mathcal{H}_{2.b}^{w+1}$ with noticeable probability. Consider the following adversary B (using D) attacking the adaptive computational hiding property of the hybrid commitment scheme (cf. Lemma 2.2).

- At the beginning B receives a verification key ϑ^* and sets-up the verification key for the signature scheme with the same distribution as in $\mathcal{H}_{2.b}$ and $\mathcal{H}_{2.c}$, and setting $\vartheta_{i^*} = \vartheta^*$.
- All *leaky* signature queries, and all *non-binding* signature queries are treated in the same way as in $\mathcal{H}_{2.b}$ and $\mathcal{H}_{2.c}$.

- For the first $w - 1$ binding signature queries (corresponding to indexes j_1, \dots, j_{w-1}), \mathbf{B} computes com as in $\mathcal{H}_{2,b}$.
- Let m be the message corresponding to the w -th binding signature query. At this point \mathbf{B} outputs $(0^\mu, \Delta(m))$ and sets $com = com^*$, where com^* is the challenge commitment sent to \mathbf{B} .
- For all remaining binding signature queries (corresponding to indexes j_{w+1}, \dots, j_k), \mathbf{B} computes com as in $\mathcal{H}_{2,c}$.
- In case \mathbf{D} outputs a set of forgeries $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$, \mathbf{B} performs the same checks as done in $\mathcal{H}_{2,b}$ and $\mathcal{H}_{2,c}$.
- Finally, \mathbf{B} outputs whatever \mathbf{D} does.

Note that the verification key simulated by \mathbf{B} has the correct distribution; in fact, all public keys are set to be equivocable (as in both $\mathcal{H}_{2,b}$ and $\mathcal{H}_{2,c}$) but the one corresponding to coordinate i^* , which is taken from the reduction (and thus has the right distribution). Additionally, in case com^* is a commitment to $\Delta(m)$ the reduction perfectly emulates the view of $\mathcal{H}_{2,b}^w$, whereas in case com^* is a commitment to 0^μ the reduction perfectly emulates the view of $\mathcal{H}_{2,b}^{w+1}$. Thus, \mathbf{B} breaks the adaptive computational hiding property with noticeable probability contradicting Lemma 2.2. This concludes the proof. \square

Experiment $\mathcal{H}_{2,d}$: In this experiment we modify the way *leaky* signature queries are answered.

Given such a query (m, f) , the experiment samples the index $j \leftarrow_{\$} [t]$ as before and, in case $j \neq i^*$, computes the commitment com as $(com, r') \leftarrow_{\$} \text{ECommit}_2(\vartheta_j, \tau_j)$. Notice that the randomness r used to compute com in the previous hybrid can be obtained as a function of the matrix Δ by running the equivocation algorithm. In particular:

$$r(\Delta) := \begin{cases} \text{Equiv}_2(\tau_j, \Delta(m), r') & \text{if } j \neq i^* \\ r \leftarrow_{\$} \mathbb{F}^\nu & \text{if } j = i^*. \end{cases}$$

The state used to reply the current leaky query is set to $st = \{(\Delta, \mathbf{r}), r(\Delta), r_{niwi}\}$, where r_{niwi} is the randomness used to generate the NIWI argument π .

Claim 4.8. $\varepsilon_{2,d} \geq \varepsilon_{2,c} - \text{negl}(\kappa)$.

Proof. Let us consider the view of the adversary in the two hybrids, consisting of the verification key, and the answers to both leaky and leak-free signature queries. The two hybrids differ only in the way leaky queries are treated.

Consider one such query. In case $j = i^*$ the two hybrids distributions are identical; in case $j \neq i^*$ the trapdoor hiding property implies that the distribution of the pair (com, r) in $\mathcal{H}_{2,d}$ is statistically close to the distribution of the same pair in $\mathcal{H}_{2,c}$. Note that in both experiments, the verification key vk , the argument π , the leakage Λ are computed as a function of Δ , $r(\Delta)$ and independent randomness. It follows that the answers to each leaky query are statistically close in the two experiments. By a standard hybrid argument, it must be that the joint distributions $((\Delta, \mathbf{r}(\Delta)), \text{View}_{2,c})$ and $((\Delta, \mathbf{r}(\Delta)), \text{View}_{2,d})$ are statistically close. The claim follows. \square

From now on the proof is purely information theoretic, so in particular the next experiments might no longer be efficient. Before defining the next hybrid, we state a lemma implicit in [8].

Lemma 4.1. *Let $\mathfrak{R} \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be an NP relation, and consider a statistical NIWI argument system for \mathfrak{R} . There exists a (possibly not efficient) reconstruction algorithm Rand that, given as input a proof π for an instance x under a witness w , another witness w' for x , and the randomness r_{niwi} used to generate π , outputs a value r'_{niwi} such that $\pi = \text{Prove}(\text{crs}, x, w'; r'_{niwi})$ (with overwhelming probability over the choice of r_{niwi}). Additionally, the distribution of r'_{niwi} is statistically close to the uniform distribution.*

Proof. For any $(x, w) \in \mathfrak{R}$, and $\text{crs} \leftarrow \text{Init}(1^\kappa)$, let $\Pi_w := \{\pi \mid \pi = \text{Prove}(\text{crs}, x, w; r_{niwi})\}$. The statistical witness indistinguishability property implies:

$$\mathbb{P}_{\pi \leftarrow \text{Prove}(\text{crs}, x, w)} [\pi \in \Pi_{w'}] \geq 1 - \text{negl}(\kappa).$$

Otherwise the event $\pi \in \Pi_{w'}$ can be used to distinguish the ensembles Π_w and $\Pi_{w'}$. In case the above condition is satisfied, algorithm **Rand** samples r'_{niwi} from the weighted distribution $\{r \mid \pi = \text{Prove}(\text{crs}, x, w'; r)\}$ (otherwise it aborts). \square

Experiment $\mathcal{H}_{2,e}$: This experiment is identical to the previous hybrid, except that (whenever possible) it uses a different witness β' to compute the argument corresponding to a *leaky* signature query. In particular, given such a query (m, f) , the experiment samples $j \leftarrow [t]$ and, in case $j \neq i^*$ it generates the argument π by running

$$\text{Prove}(\text{crs}, \underbrace{(\vartheta, \vartheta_j, \tilde{c}om, com)}_{\alpha}, \underbrace{(0, \text{Equiv}_1(\tau, 0^\mu, r'(m)), \text{Equiv}_2(\tau_j, 0^\mu, r'))}_{\beta'}; r'_{niwi}),$$

where $r'(m) = \sum_{i=0}^d r'_i \cdot m^i$ is computed using the randomness $\{r'_i\}_{i=0}^d$. Notice that the randomness r_{niwi} used to generate the NIWI argument in the previous experiment can be sampled (inefficiently) as a function of the witness $\beta := (\Delta(m), \text{Equiv}_1(\tau, \Delta(m), r'(m)), \text{Equiv}_2(\tau_j, \Delta(m), r'))$ and r'_{niwi} . In particular:

$$r_{niwi}(\Delta) := \begin{cases} r_{niwi} \text{ (as sampled by the prover)} & \text{if } j = i^* \\ \text{Rand}(\pi, r'_{niwi}, \beta) & \text{if } j \neq i^*, \end{cases}$$

where **Rand** is the reconstruction algorithm of Lemma 4.1. The state used to answer the leakage query f is set to $st(\Delta) = ((\Delta, \mathbf{r}(\Delta)), r(\Delta), r_{niwi}(\Delta))$.

Claim 4.9. $\varepsilon_{2,e} \geq \varepsilon_{2,d} - \text{negl}(\kappa)$.

Proof. The linear homomorphic property of the hybrid commitment scheme (cf. Section 2.3) ensures that the value $r'(m)$ is the right randomness to equivocate $\tilde{c}om$.

By statistical witness indistinguishability, for each leaky query such that $j \neq i^*$ the distribution of the argument π is statistically close in the two hybrids. The same holds for the reconstructed state, by Lemma 4.1. It follows that $((\Delta, \mathbf{r}(\Delta)), \text{View}_{2,d})$ and $((\Delta, \mathbf{r}(\Delta)), \text{View}_{2,e})$ are statistically close, which concludes the proof. \square

The next hybrid that we define has no direct access to the matrix Δ , but instead depends on a leakage oracle $\text{Leak}_\Delta(\cdot)$ which takes as input a function f and returns $f(\Delta)$.

Experiment $\mathcal{H}_3^{\text{Leak}_\Delta(\cdot)}$: This experiment is the same as the previous hybrid, with the difference that Δ is not sampled by the hybrid as part of the signing key, but can instead be accessed via $\text{Leak}_\Delta(\cdot)$. In particular, all *leak-free* signature queries and *leaky* signature queries where $j \neq i^*$, are handled as in $\mathcal{H}_{2,e}$. For a *leaky* signature query (m, f) where the index j happens to be equal to i^* , the experiment defines the leakage function $f'_{m, (r'_0, \dots, r'_d), r', r'_{niwi}} := (f(st(\Delta)), \Delta(m))$ and queries $\text{Leak}_\Delta(\cdot)$ on f' .⁶ We call such query a *bad* query and we let **Bad** denote the corresponding event. Given the value $\Delta(m)$, the signature σ can be computed as in $\mathcal{H}_{2,e}$.

Claim 4.10. $\varepsilon_3 = \varepsilon_{2,e}$.

⁶Note that the function hard-wires the message m , and the randomness needed to reconstruct the current state as a function of Δ (as specified in the previous hybrids).

Proof. The only difference between the two hybrids stems from the fact that Δ is only accessed via a leakage oracle in \mathcal{H}_3 . Note that in $\mathcal{H}_{2.e}$ the matrix Δ is needed only to answer leaky signature queries; in case of a bad query, \mathcal{H}_3 accesses the leakage oracle to additionally obtain $\Delta(m)$ in addition to the usual leakage on the state. Given $\Delta(m)$, the signature computation is identical in the two hybrids. \square

In what follows, let View_3 be the entire view of experiment \mathcal{H}_3 ; this includes the verification key vk , and the answers to all leaky and leak-free signature queries. The two lemmas below characterize the conditional min-entropy of Δ given View_3 . The first lemma says that, with high probability over the randomness in the experiment, the only information revealed on Δ comes from the leakage and the bad queries. Recall that we assume the adversary \mathbf{A} makes exactly q leaky queries.

Lemma 4.2. *For any $\beta > 0$ we have that*

$$\mathbb{P}[\mathbb{H}_\infty(\Delta | \text{View}_3 = v) \geq |\Delta| - (2eq/t)\mu \log |\mathbb{F}| - \beta - \lambda] > 1 - 2^{-\beta} - 2^{-2eq/t}, \quad (3)$$

where the probability is taken over the randomness of the experiment.

Proof. Let q_s be the total number of signature queries made by the adversary. We write $\mathcal{L} \subseteq [q_s]$ for the set of indexes corresponding to the leaky signature queries. Furthermore, we let $\mathcal{Z} := \{i \in \mathcal{L} \wedge \text{Bad}_i\}$ be the set of indexes corresponding to the bad queries. Note that the cardinality Z of \mathcal{Z} is a random variable, depending on the event that the index $j \leftarrow_s [t]$ (sampled independently for each leaky query) happens to hit the index i^* .

The random variable View_3 consists of the simulated verification key, the answers to all leak-free signature queries and the answers to all leaky signature queries made by the adversary. Let $\text{View}'_3 := (\Lambda_i)_{i \in \mathcal{L}} \| (\Delta(m_i))_{i \in \mathcal{Z}}$ the total information leaked in \mathcal{H}_3 . Recall that in \mathcal{H}_3 both the verification key, the answers to leak-free queries and the signatures corresponding to leaky queries where the event Bad does not happen, are completely independent on Δ . Hence,

$$\mathbb{H}_\infty(\Delta | \text{View}_3 = v) = \mathbb{H}_\infty(\Delta | \text{View}'_3 = v'),$$

where for any realization of the adversary's view v , we write v' for the “striped-off” view that considers only the total information retrieved via the leakage oracle.

Observe that $|\mathcal{Z}| = Z = \sum_{i \in [q_s]} \text{Bad}_i$, and in particular $|\text{View}'_3| \leq \lambda + Z\mu \log |\mathbb{F}|$. Furthermore, note that the number of leaky queries is equal to q , so that $\mathbb{E}[Z] = q/t$. Since the events Bad_i are independent, we can apply a Chernoff bound and write:

$$\mathbb{P}[Z > 2eq/t] < 2^{-2eq/t}.$$

Setting $\lambda' := \lambda + (2eq/t)\mu \log |\mathbb{F}|$ and $\Delta := \lambda' + \beta$, we obtain

$$\begin{aligned} & \mathbb{P}[\mathbb{H}_\infty(\Delta | \text{View}'_3 = v') \geq |\Delta| - \Delta] \\ & \geq \mathbb{P}[\mathbb{H}_\infty(\Delta | \text{View}'_3 = v') \geq |\Delta| - \Delta \mid Z \leq 2eq/t] (1 - 2^{-2eq/t}) = \\ & = \mathbb{P}[\mathbb{H}_\infty(\Delta | \text{View}'_3 = v') \geq |\Delta| - \Delta \mid |\text{View}'_3| \leq \lambda'] (1 - 2^{-2eq/t}). \end{aligned} \quad (4)$$

Applying Lemma 2.1 to the right hand side of the last inequality, we have:

$$\mathbb{P}[\mathbb{H}_\infty(\Delta | \text{View}'_3 = v') \geq |\Delta| - \Delta \mid |\text{View}'_3| \leq \lambda'] \geq 1 - 2^{\lambda' - \Delta} = 1 - 2^{-\beta}. \quad (5)$$

The lemma follows by combining inequality (4) and (5). \square

Lemma 4.3. *For any adversary \mathbf{A} such that $\mathbb{P}[\widetilde{\text{Exp}}_{\mathcal{SS}_1, \mathbf{A}}^{\text{one-more}}(\kappa, \ell, q_s, \gamma) = 1] = \varepsilon$, there exists a constant $c > 0$ such that:*

$$\mathbb{E}[\mathbb{H}_\infty(\Delta | \text{View}_3 = v)] \leq \left(d + 1 - \frac{n - 1 + q}{t} \right) \mu \log |\mathbb{F}| - c \log \varepsilon(\kappa),$$

where the expectation is taken over the randomness of the experiment.

Proof. To prove the lemma we define a predictor strategy P (running A) that outputs Δ with “high enough” probability.

Let \mathcal{K} be the set of indexes such that $i \in \mathcal{K}$ if and only if the i -th forged signature σ_i^* is of the form $(\pi_i^*, i^*, c\tilde{m}_i^*)$. W.l.o.g. let us assume that $\mathcal{K} := \{1, \dots, K\}$ for a random variable $K \in [n]$. The predictor P proceeds as follows:

1. At the beginning P interacts with A , playing the role of the challenger in \mathcal{H}_3 . Let $\{m_i\}_{i \in \mathcal{Z}}$ be the set of messages corresponding to the *bad* leaky signature queries, i.e. to the queries for which \mathcal{H}_3 retrieved the vector $\Delta(m_i)$ via the leakage oracle $\text{Leak}_\Delta(\cdot)$.
2. Whenever A outputs a set of forgeries $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$, P finds, for all $i \in [K]$, the unique vector $\Delta(m_i^*)$ such that $\text{com}_i^* = \text{Commit}_2(\vartheta_{i^*}, \Delta(m_i^*); r_i^*)$ (for some randomness $r_i^* \in \mathbb{F}^\nu$).
3. For each $j \in [\mu]$, P solves the following linear system:

$$\begin{pmatrix} 1 & m_1 & \dots & m_1^d \\ & & \ddots & \\ 1 & m_Z & \dots & m_Z^d \\ 1 & m_1^* & \dots & m_1^{*d} \\ & & \ddots & \\ 1 & m_K^* & \dots & m_K^{*d} \\ 1 & \bar{m}_1 & \dots & \bar{m}_1^d \\ & & \ddots & \\ 1 & \bar{m}_{d+1-K-Z} & \dots & \bar{m}_{d+1-Z-K}^d \end{pmatrix} \cdot \begin{pmatrix} \delta_{j,0} \\ \vdots \\ \delta_{j,d} \end{pmatrix} = \begin{pmatrix} \delta_j(m_1) \\ \vdots \\ \delta_j(m_Z) \\ \delta_j(m_1^*) \\ \vdots \\ \delta_j(m_K^*) \\ y_{j,1} \\ \vdots \\ y_{j,d+1-Z-K} \end{pmatrix}, \quad (6)$$

where $\delta_j(X)$ is the polynomial defined by the coefficients in the j -th row of Δ . The values $y_1, \dots, y_{d+1-Z-K}$ are chosen uniformly at random from \mathbb{F} , whereas the values $\bar{m}_1, \dots, \bar{m}_{d+1-Z-K}$ are chosen uniformly in $\mathbb{F} \setminus (\{m_i\}_{i \in [Z]} \cup \{m_i^*\}_{i \in [K]})$.

4. Output the matrix Δ' having $(\delta_{j,i})_{j \in [\mu], i \in [0,d]}$ as elements.

Recall that the common reference string corresponding to index i^* is perfectly binding. Hence P always succeeds in finding the values $\{\Delta(m_i^*)\}_{i \in [K]}$, and these values are unique. Moreover, since we are conditioning on False_3 not happening, the extracted values contain the evaluation of the polynomials $\delta_1(X), \dots, \delta_\mu(X)$ at m_i^* . Also note that the system of Eq. (6) is full-rank.

We conclude that the matrix Δ' is equal to Δ , provided that P guessed correctly the values $y_{j,1}, \dots, y_{j,d-Z-K}$ for all $j \in [\mu]$, and that the forgeries output by A are such that $\text{Forge}_3 \wedge \neg \text{False}_3$ happens. The probability of the first event is equal to $|\mathbb{F}|^{-\mu(d+1-Z-K)}$. The probability of the second event is equal to ε_3 ; combining Claims 4.1–4.10 there exists a constant c such that $\varepsilon_3(\kappa) \geq \varepsilon(\kappa)^c$. Thus,

$$\mathbb{H}_\infty(\Delta \mid \text{View}_3 = v) \leq \mu(d+1-Z-K) \log |\mathbb{F}| - c \log \varepsilon(\kappa).$$

The lemma now follows by noticing that $\mathbb{E}[K] > \frac{n-1}{t}$, which implies

$$\mathbb{E}[-Z-K] \leq -\frac{n-1+q}{t}.$$

□

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. We start by computing the leakage bound. To this end, let us set the parameters such that in Lemma 4.2 the probability that the conditional min-entropy of Δ given

the view in the last hybrid is at least 1. Let $\beta > 0$, for concreteness say $\beta = 2$; furthermore let $t := \frac{2e\rho q}{\xi(d+1)} - 1$ for any $0 < \xi \leq \rho$. Lemma 4.2 requires that:

$$\begin{aligned}\lambda &\leq |\Delta| - (2eq/t)\mu \log |\mathbb{F}| - 3 \\ &= |\Delta| - (\xi/\rho(d+1) + 1)\mu \log |\mathbb{F}| - 3,\end{aligned}$$

which is implied by:

$$\begin{aligned}\lambda &\leq |\Delta| - (\xi/\rho)(d+1)\mu \log |\mathbb{F}| = \\ &= (1 - \xi/\rho)|\Delta| = (\rho - \xi)|sk|.\end{aligned}$$

W.l.o.g. let us assume $\lambda = (\rho - \xi)|sk|$; since $q \leq \lambda$ we get

$$t = \frac{2e\rho q}{\xi(d+1)} - 1 \leq \frac{2e\rho\lambda}{\xi(d+1)} = \frac{2e\rho(1 - \xi/\rho)(d+1)\mu\kappa}{\xi(d+1)} = \frac{2e\rho(1 - \xi/\rho)\mu}{\xi}\kappa = O(\kappa).$$

Lemma 4.3 implies:

$$\mathbb{P} \left[\mathbb{H}_\infty(\Delta | \text{View}_3 = v) \leq \frac{\left(d + 1 - \frac{n-1+q}{t}\right) \mu \log |\mathbb{F}| - c \log \varepsilon(\kappa)}{1 - 2^{-\beta} - 2^{-2eq/t}} \right] \geq 2^{-\beta} + 2^{-2eq/t}. \quad (7)$$

Notice that Eq. (3) and Eq. (7) give, respectively, a lower bound and an upper bound on the conditional min-entropy of the random variable Δ given the entire view at the end of the experiment. We show that, whenever $\varepsilon(\kappa)$ is noticeable, by setting the parameters as in the statement of the theorem the lower bound is strictly bigger than the upper bound. It follows that the event of Eq. (3) implies the negation of the event of Eq. (7). We obtain a contradiction by noticing that the sum of the two probabilities is strictly bigger than 1.

Hence, we need to show that the following holds:

$$\frac{\left(d + 1 - \frac{n-1+q}{t}\right) \mu \log |\mathbb{F}| - c \log \varepsilon(\kappa)}{1 - 2^{-\beta} - 2^{-2eq/t}} < (d+1)\mu \log |\mathbb{F}| - (2eq/t)\mu \log |\mathbb{F}| - 2 - \lambda. \quad (8)$$

Substituting $\log |\mathbb{F}| = \kappa$, and rearranging, we obtain:

$$\frac{(n-1)\mu\kappa}{t} > (2^{-\beta} + 2^{-2eq/t} - 1) \cdot ((d+1 - 2eq/t)\mu\kappa - 2 - \lambda) + (d+1 - q/t)\mu\kappa - c \log \varepsilon(\kappa).$$

For the last inequality to hold it suffices that

$$\frac{(n-1)\mu\kappa}{t} > (d+1 - q/t)\mu\kappa + (2e-1)q/t + 4 + 2\lambda - c \log \varepsilon.$$

Recall that $q/t \geq \frac{\xi(d+1)}{2e\rho}(d+1)$ and $(d+1) = \frac{\lambda}{\mu\kappa(\rho-\xi)}$, hence, for the last inequality to hold, it suffices that

$$\frac{(n-1)\mu\kappa}{t} > \left(2 + \frac{1 - \xi/2e\rho}{\rho - \xi} + \frac{2e-1}{2e\rho}\right) \lambda + 4 - c \log \varepsilon(\kappa).$$

Thus, in order for Eq. (8) to hold, it suffices that

$$\frac{n\kappa}{t} = \Omega(\lambda - \log \varepsilon(\kappa)).$$

If $\varepsilon(\kappa)$ is noticeable, recalling that $t = O(\kappa)$ and $\lambda = \Omega(\kappa)$, we get that n is linear in λ . Since, by definition, $n \geq \lfloor \frac{\lambda}{\gamma \cdot s} \rfloor + 1$ and a signature consists of a constant number of group elements, it suffices to set the slack parameter γ as $\gamma(\kappa) = O(1/\kappa)$ in order for Eq. (8) to hold. \square

4.3 A Variant

The construction of Section 4.1 implicitly defines an *unbounded simulation sound* NIZK. In fact, the OR-proof technique resembles the unbounded simulation sound NIZK argument systems given in [36, 35, 11]. The main drawback, as noticed in [22, 45], is that the arguments for the relation \mathfrak{R}_{OR} instantiated using the Groth-Sahai proof system [37] need to use quadratic equations, and thus can be significantly less efficient than the arguments for \mathfrak{R} .

In this section we show how to avoid the expensive OR-proof trick. The technical tools we use are:

- A trapdoor hiding, linearly homomorphic commitment scheme \mathcal{COM}_1 as in the scheme from Section 4.1.
- A hybrid commitment \mathcal{COM}_2 as in the scheme from Section 4.1, with the additional property that the public key space is a group \mathbb{G} of cardinality p with $\log p = O(\kappa)$, and the equivocable keys form a subgroup \mathbb{H} of \mathbb{G} (with cardinality p'). We also require that: (a) for any two elements $g_1 \in \mathbb{G} \setminus \mathbb{H}, g_2 \in \mathbb{H}$ the element $g_1 \cdot g_2 \in \mathbb{G} \setminus \mathbb{H}$, and (b) the ESetup_2 algorithm produces an equivocable key that is uniformly distributed in \mathbb{H} .
- A statistical NIWI argument system $\mathcal{NIWI} = (\text{Init}_1, \text{Prove}_1, \text{Ver}_1)$ for the same relation $\mathfrak{R}_1 := \mathfrak{R}$ of the construction from Section 4.1.
- A true-simulation extractable NIZK argument system $\mathcal{NIZK} = (\text{Init}_2, \text{Prove}_2, \text{Ver}_2)$ (see Section 2.4) supporting labels, and with obvious sampling of the common reference string, for the relation

$$\mathfrak{R}_2 := \{(\vartheta, \vartheta_1, \vartheta_2); (\alpha_1, \alpha_2) \mid \vartheta = \vartheta_1^{\alpha_1} \cdot \vartheta_2^{\alpha_2}\} \subseteq \mathbb{G}^3 \times \mathbb{Z}_p^2. \quad (9)$$

For brevity, given a vector $\mathbf{g} \in \mathbb{G}^n$ and $\boldsymbol{\alpha} \in \mathbb{Z}_p^n$ we write $\langle \mathbf{g}, \boldsymbol{\alpha} \rangle = \prod_{i=1}^n g_i^{\alpha_i}$.

Our scheme $\mathcal{SS}_2 = (\text{KGen}, \text{Sign}, \text{Verify})$ has message space equal to \mathbb{F} and is described below:

Key Generation. Let $t, d, \mu, \nu \in \mathbb{N}$ be parameters. Run $\text{crs}_1 \leftarrow \text{Init}_1(1^\kappa)$ and $\text{crs}_2 \leftarrow \text{Init}_2(1^\kappa)$, and sample the values $\vartheta, \{\vartheta_i\}_{i=1}^t, \boldsymbol{\Delta} \in (\mathbb{F}^\mu)^{d+1}, \mathbf{r} \in \mathbb{F}^{d+1}$, and $\{\text{com}_i\}_{i=1}^d$ as in the scheme from Section 4.1. Output

$$sk = (\boldsymbol{\Delta}, \mathbf{r}) \quad vk = (\text{crs}_1, \text{crs}_2, \vartheta, \{\vartheta_i\}_{i=1}^t, \{\text{com}_i\}_{i=0}^d).$$

Signature. For $j \in [\mu]$, let $\delta_j(X), r(X)$, and $\boldsymbol{\Delta}(X)$ be defined as in the scheme from Section 4.1. Consider the following polynomial-time relation:

$$\mathfrak{R}_1 := \left\{ (\vartheta, \vartheta', \text{c}\tilde{m}, \text{com}); (\tilde{m}, \tilde{r}, r) \mid \begin{array}{l} \text{c}\tilde{m} = \text{Commit}_1(\vartheta, \tilde{m}; \tilde{r}) \\ \text{com} = \text{Commit}_2(\vartheta', \tilde{m}; r) \end{array} \right\}.$$

To sign a message $m \in \mathbb{F}$ compute $\tilde{m} = \boldsymbol{\Delta}(m)$ and $\tilde{r} = r(m)$, and let $\text{c}\tilde{m} := \text{Commit}_1(\vartheta, \tilde{m}; \tilde{r})$. Notice that both values \tilde{m}, \tilde{r} can be computed efficiently as a function of the signing key $(\boldsymbol{\Delta}, \mathbf{r})$ and the message to be signed. Then proceed as follows:

1. Pick random indexes $j_1, j_2 \leftarrow [t]$ such that $j_1 \neq j_2$, sample $\alpha_1, \alpha_2 \leftarrow \mathbb{Z}_p$ and compute $\vartheta' := \vartheta_{j_1}^{\alpha_1} \cdot \vartheta_{j_2}^{\alpha_2}$ and $\text{com} := \text{Commit}_2(\vartheta', \tilde{m}; r)$ where $r \leftarrow \mathbb{F}^\mu$.
2. Using crs_1 as common reference string, generate a NIWI argument π_1 for $(\vartheta, \vartheta', \text{c}\tilde{m}, \text{com}) \in \mathfrak{L}_{\mathfrak{R}_1}$.
3. Using crs_2 as common reference string, generate a NIZK argument π_2 for $(\vartheta', \vartheta_{j_1}, \vartheta_{j_2}) \in \mathfrak{L}_{\mathfrak{R}_2}$ using m as a label.
4. Output $\sigma = (j_1, j_2, \vartheta', \text{com}, \pi_1, \pi_2)$.

Verification Given a pair (m, σ) , parse σ as $(j_1, j_2, \vartheta', \text{com}, \pi_1, \pi_2)$. Output:

$$\text{Ver}_1(\text{crs}_1, \pi_1, (\vartheta, \vartheta', \prod_{i=1}^d \text{com}_i^{m_i}, \text{com})) \wedge \text{Ver}_2(\text{crs}_2, \pi_2, m, (\vartheta', \vartheta_{j_1}, \vartheta_{j_2})).$$

Theorem 4.2. *Let $\mu \in \mathbb{N}$, and let \mathbb{F} be a finite field of size $\log |\mathbb{F}| = \kappa$ for security parameter $\kappa \in \mathbb{N}$. For any constant $0 < \xi \leq \frac{\mu}{\mu+1}$ and $t = O(\kappa)$, whenever $\gamma = O(1/\kappa)$ the above signature scheme is $((\frac{\mu}{\mu+1} - \xi)|sk|, \gamma)$ -fully-leakage one-more unforgeable with perfect erasures.*

The main difference with the proof of Theorem 4.1 is in the way we simulate leak-free queries. In the previous proof the simulator knows a value y , pre-image of x under the PRG G , and therefore it can produce a valid argument for the OR-relation and simulate correctly leak-free queries thanks to the WI property of \mathcal{NIZK} . Note, however, that whenever the sampled index $j \leftarrow^s [t]$ is such that $j \neq i^*$ (i.e. both the verification keys for the commitment scheme are equivocable) a valid NIWI argument can be produced by opening the commitments to 0.

Unfortunately the above strategy cannot be always used by the simulator, as the index j is part of the signature (and thus known to the adversary). In other words, although the simulator knows where the equivocable keys are, it cannot exploit this extra knowledge, as it needs to sample j using the same distribution used by the real signer.

The idea of re-randomizing the verification keys each time we generate a signature, allows the simulator to exploit the above extra knowledge. In fact, whenever the index j corresponding to the binding verification key is chosen, the simulator can force the re-randomized verification key ϑ' to always be equivocable. The zero-knowledge and the hybridness property ensure that the view of the adversary is computationally indistinguishable from the view in the real experiment; using the message m as label avoids the simple “cut-and-paste” attack in which an adversary re-uses a pair (ϑ', π_2) from a leak-free signature query in its forgeries.⁷

4.4 Proof Sketch of Theorem 4.2

The proof of Theorem 4.2 uses a hybrid argument, along the same lines of the proof of Theorem 4.1. Since many of the hybrids are identical, we only give a proof sketch here.

Let A be an adversary making $q = \text{poly}(\kappa)$ leaky signature queries, with advantage ε in the fully-leakage one-more unforgeability experiment for \mathcal{SS}_2 . Let $\mathcal{H}_0, \mathcal{H}_{1,a}, \mathcal{H}_{1,b}, \text{Forge}_0$ be the same as in the proof of Theorem 4.1. Let False_0^1 be the same as False_0 , and define $\varepsilon_i := \mathbb{P}[\text{Forge}_i \wedge \neg \text{False}_i^1]$. A reasoning similar to the one in Claim 4.1–4.3 shows that $\varepsilon_{1,b} \geq \varepsilon - \text{negl}(\kappa)$.

Experiment $\mathcal{H}_{2,a}$: This hybrid differs from the previous one in the way leak-free signature queries are dealt with. Specifically, during key generation the common reference string crs_2 is sampled by running $(\text{crs}_2, \tau_{\text{sim}}, \tau_{\text{ext}}) \leftarrow \text{Init}(1^\kappa)$ which additionally returns a simulation trapdoor τ_{sim} and an extraction trapdoor τ_{ext} . Moreover, for each query $m \in \mathbb{F}$ the NIZK argument π_2 is computed by running $\pi_2 \leftarrow \text{Sim}(\tau_{\text{sim}}, (m, (\vartheta', \vartheta_{j_1}, \vartheta_{j_2})))$.

Let $\pi_{2,i}^*$ be the NIZK argument contained in the i -th forgery returned by the adversary. Also, let $(\alpha_{1,i}^*, \alpha_{2,i}^*) \leftarrow \text{K}(\tau_{\text{ext}}, m_i^*, \pi_{2,i}^*)$ be the witness extracted from such argument. Define $\text{False}_{2,a}^2$ to be the event that there exists an index $i \in [n]$ such that the extracted witness $(\alpha_{1,i}^*, \alpha_{2,i}^*)$ is not valid, i.e.,

$$\vartheta'^* \neq \langle (\vartheta_{j_1}^*, \vartheta_{j_2}^*), (\alpha_{1,i}^*, \alpha_{2,i}^*) \rangle.$$

For the hybrid $\mathcal{H}_{2,a}$, and for each of the hybrids below, we write ε'_i for the probability of the event $\text{Forge}_i \wedge \neg \text{False}_i^1 \wedge \neg \text{False}_i^2$.

Claim 4.11. $\varepsilon'_{2,a} \geq \varepsilon_{1,b} - \text{negl}(\kappa)$.

⁷In this case the predictor of Lemma 4.3 does not work.

Proof. Note that, by true-simulation extractability, we have that $\mathbb{P}[\text{False}_{2,a}^2] \leq \text{negl}(\kappa)$ which implies $\varepsilon'_{2,a} \geq \varepsilon_{2,a} - \text{negl}(\kappa)$. Furthermore, by the adaptive multi-theorem zero-knowledge property of the NIZK argument system, we have that the views $\text{View}_{2,a}$ and $\text{View}_{1,b}$ are computationally indistinguishable, and thus $\varepsilon_{2,a} \geq \varepsilon_{1,b} - \text{negl}(\kappa)$ which implies the claim. \square

Experiment $\mathcal{H}_{2,b}$: The same as the previous hybrid, but now for each *leak-free* signature query of the form $m \in \mathbb{F}$, in case one of the two random indexes (say j_2) is equal to i^* then set $\alpha_2 := 0$ and compute $\vartheta' = \vartheta_{j_1}^{\alpha_1}$ where $\alpha_1 \leftarrow_{\mathcal{S}} \mathbb{Z}_{p'}$.

Claim 4.12. $\varepsilon'_{2,b} \geq \varepsilon'_{2,a} - \text{negl}(\kappa)$.

Proof. Note that for each *leak-free* signature query the verification key ϑ' is uniformly distributed in \mathbb{H} , and thus it has the same distribution as a verification key output by ESetup_2 . The claim follows by the hybridness property of the commitment scheme. \square

Experiment $\mathcal{H}_{2,c}$: The same as the previous hybrid, but now for each *leak-free* signature query of the form $m \in \mathbb{F}$, the value com of the signature σ is computed as $(com, r') \leftarrow \text{ECommit}_2(\vartheta', \tau')$.

Claim 4.13. $\varepsilon'_{2,c} \geq \varepsilon'_{2,b} - \text{negl}(\kappa)$.

Proof. Follows by a hybrid argument relying on the trapdoor hiding property of the hybrid commitment scheme. \square

Hybrids $\mathcal{H}_{2,d}$ and $\mathcal{H}_{2,e}$ are identical to the corresponding hybrids in the proof of Theorem 4.1. Putting it all together we have shown $\varepsilon'_{2,e} \geq \varepsilon - \text{negl}(\kappa)$. One can show that the equivalents of Lemma 4.2 and Lemma 4.3 hold true with exactly the same parameters. The only difference is in the way we show that $\mathbb{E}[K] \geq \frac{n-1}{t}$ (in the proof of Lemma 4.3). In particular we have to evaluate (in expectation) the number of forged signatures for which the commitment is binding. Towards this, for a value $i \in [n]$, let $\text{Binding}_{2,e}^i$ be the event that the verification key contained in the j -th forged signature of hybrid $\mathcal{H}_{2,e}$ is an element of $\mathbb{G} \setminus \mathbb{H}$. It is easy to see that, for all $i \in [n]$,

$$\mathbb{P}[\text{Binding}_{2,e}^i] = \mathbb{P}[\text{Binding}_{2,c}^i] \geq \mathbb{P}[\text{Binding}_{2,b}^i] - \text{negl}(\kappa),$$

where the inequality comes from the hybridness property of the commitment scheme. Furthermore, $\mathbb{P}[\text{Binding}_{2,b}^i] = \mathbb{P}[\text{Binding}_{2,a}^i]$ so we are left to bound the probability of $\text{Binding}_{2,a}^i$. We have that:

$$\begin{aligned} \mathbb{P}[\text{Binding}_{2,a}^i] &\geq \mathbb{P}[\text{Binding}_{2,a}^i \mid \neg \text{False}_{2,a}^2] \cdot \mathbb{P}[\neg \text{False}_{2,a}^2] \\ &\geq \mathbb{P}[\text{Binding}_{2,a}^i \mid \neg \text{False}_{2,a}^2] - \text{negl}(\kappa) \end{aligned} \tag{10}$$

$$\geq \underbrace{\mathbb{P}[\exists k \in \{1, 2\} : j_k = i^* \wedge \alpha_{k,i}^* \neq 0]}_{\text{Hit}_{2,a}^i} - \text{negl}(\kappa), \tag{11}$$

where i^* is the index corresponding to the only binding verification key ϑ_{i^*} . Eq. (10) follows by true-simulation extractability of the NIZK argument system.⁸ Eq. (11) follows by the fact that, given that the extracted witness $(\alpha_{1,i}^*, \alpha_{2,i}^*)$ is valid, the key ϑ' is binding only if the event $\text{Hit}_{2,a}^i$ is verified.

One observes that $\mathbb{P}[\text{Hit}_{2,a}^i] = \mathbb{P}[\text{Hit}_{1,a}^i]$. Furthermore, by the hybridness property of the commitment scheme,

$$\begin{aligned} \mathbb{P}[\text{Hit}_{1,a}^i] &\geq \mathbb{P}_{i^* \leftarrow_{\mathcal{S}} [t]}^{\mathcal{H}_0} [\exists k \in \{1, 2\} : j_k = i^* \wedge \alpha_{k,i}^* \neq 0] - \text{negl}(\kappa) \\ &\geq \frac{1}{2} \left(\frac{1}{t} + \frac{1}{t-1} \right) - \text{negl}(\kappa) \geq \frac{1}{t} - \text{negl}(\kappa). \end{aligned}$$

⁸The forged messages were not queried to the signature oracle which means that for each forged signature (m^*, σ^*) no simulated argument was given for the pair $((\vartheta^*, \vartheta_{j_1}^*, \vartheta_{j_2}^*), m^*)$.

where the last inequality follows from the fact that the random index i^* is independent from \mathcal{H}_0 . This concludes the proof sketch.

5 A Scheme without Erasures

In this section we construct a signature scheme that is fully-leakage one-more unforgeable, without assuming secure erasures (cf. Definition 3.1). The drawback of the scheme is that the graceful degradation provided is not tight, as the slack parameter γ depends linearly on the number q_s of signature queries asked by the adversary. As mentioned in Section 3.1, since q_s is polynomial in the security parameter, this flavour of one-more unforgeability is still suitable for some applications.

Our scheme, and its security analysis, is similar to the constructions of Section 4. One additional difficulty is that we cannot distinguish anymore between leaky and leak-free signature queries, as now a leakage query might depend on the entire state of the signer (which is never erased).

To deal with this difficulty, we define a new simulation strategy that only rarely requires to access the signing key in order to simulate signature queries; moreover the probability of accessing the signing key can be controlled by the simulator, such that, on average, the conditional min-entropy of the signing key remains high. This feature allows the simulation to go through.

Before coming to the scheme, we abstract away a special type of hybrid commitment scheme whose properties will turn useful in the security proof. We call this tool a *secret sharing* hybrid commitment, and we define it in Section 5.1. The signature scheme is described in Section 5.2, and its security analysis can be found in Section 5.3.

5.1 Secret Sharing Hybrid Commitment

Let \mathbb{F} be a finite field. For parameters $\mu, \nu \in \mathbb{N}$, let $\mathcal{COM} = (\text{Setup}, \text{Commit}, \text{ESetup}, \text{ECommit}, \text{Equiv})$ be a hybrid, linearly homomorphic commitment scheme with message space equal to \mathbb{F}^μ and randomness space equal to \mathbb{F}^ν . Fix parameters $p := \log \kappa$ and $t := \kappa$, for security parameter $\kappa \in \mathbb{N}$. A *secret sharing* hybrid commitment scheme is a tuple of algorithms $\mathcal{COM}_{ss} = (\text{Setup}_{ss}, \text{Commit}_{ss}, \text{ESetup}_{ss}, \text{ECommit}_{ss}, \text{Equiv}_{ss})$ specified as follow.

Algorithm $\text{Setup}_{ss}(1^\kappa)$: For all $i \in [t]$ run $\vartheta_i \leftarrow_s \text{Setup}(1^\kappa)$ and output the verification key $\bar{\vartheta} := (\vartheta_1, \dots, \vartheta_t)$.

Algorithm $\text{ESetup}_{ss}(1^\kappa)$: For all $i \in [t]$ generate $(\vartheta_i, \tau_i) \leftarrow_s \text{ESetup}(1^\kappa)$ and output the verification key $\bar{\vartheta} := (\vartheta_1, \dots, \vartheta_t)$ and the trapdoor key $\bar{\tau} := (\tau_1, \dots, \tau_t)$.

Algorithm $\text{Commit}_{ss}(\bar{\vartheta}, m)$: Given a verification key $\bar{\vartheta}$ and a message $m \in \mathbb{F}^\mu$ do the following:

1. Choose random shares $s_1, \dots, s_{p-1} \leftarrow_s \mathbb{F}^\mu$ and set $s_p := m - \sum_{i=1}^{p-1} s_i$;
2. Choose indexes $\mathbf{j} := (j_1, \dots, j_p) \leftarrow_s [t]^p$ and for each $i \in [p]$ compute $com_i \leftarrow_s \text{Commit}(\vartheta_{j_i}, s_i; r_i)$ where $r_i \leftarrow_s \mathbb{F}^\nu$;
3. Output $c\bar{m} := (\mathbf{j}, \mathbf{com})$ where $\mathbf{com} = (com_1, \dots, com_p)$, and define the randomness (i.e., the opening) as $\bar{\tau} := (\mathbf{j}, \mathbf{s}, \mathbf{r})$ where $\mathbf{s} = (s_1, \dots, s_{p-1})$ and $\mathbf{r} = (r_1, \dots, r_p)$.

Algorithm $\text{ECommit}_{ss}(\bar{\vartheta}, \bar{\tau})$: Given a verification key $\bar{\vartheta}$ and the trapdoor $\bar{\tau}$, do the following:

1. Choose indexes $\mathbf{j} = (j_1, \dots, j_p) \leftarrow_s [t]^p$;
2. Run $(com_p, r'_p) \leftarrow_s \text{ECommit}(\vartheta_{j_p}, \tau_{j_p})$ and, for each $i \in [1, p-1]$, compute $com_i \leftarrow_s \text{Commit}(\vartheta_{j_i}, s_i; r_i)$ where $s_i \in \mathbb{F}^\mu$ and $r_i \leftarrow_s \mathbb{F}^\nu$;
3. Output commitment $c\bar{m} := (\mathbf{j}, \mathbf{com})$ where $\mathbf{com} = (com_1, \dots, com_p)$, and randomness $\bar{r}' := (\mathbf{j}, \mathbf{s}, \mathbf{r}')$ where $\mathbf{s} = (s_1, \dots, s_{p-1})$ and $\mathbf{r}' = (r_1, \dots, r_{p-1}, r'_p)$.

Algorithm $\text{Equiv}_{ss}(\bar{\tau}, m, \bar{r}')$: Given a trapdoor $\bar{\tau}$, a message m and randomness \bar{r}' , do the following:

1. Parse the randomness as $\bar{r}' = (\mathbf{j}, \mathbf{s}, \mathbf{r}')$, where $\mathbf{r}' = (r_1, \dots, r_{p-1}, r'_p)$, and the trapdoor as $\bar{\tau} = (\tau_1, \dots, \tau_t)$;
2. Compute $s_p := m - \sum_{i=1}^{p-1} s_i$, and run $r_p := \text{Equiv}(\vartheta_{j_p}, s_p, r'_p)$;
3. Output randomness $\bar{r} := (\mathbf{j}, \mathbf{s}, \mathbf{r})$, for $\mathbf{r} = (r_1, \dots, r_{p-1}, r_p)$.

It is easy to see that \mathcal{COM}_{ss} is a hybrid linearly homomorphic commitment scheme with message space \mathbb{F}^μ and randomness space equal to $\mathcal{R}_{ss} := [t]^p \times \mathbb{F}^{(p-1)\mu} \times \mathbb{F}^{p\nu}$.

Probable binding commitment. Let \mathcal{COM}_{ss} be a secret sharing hybrid commitment. We equip \mathcal{COM}_{ss} with a special tuple of algorithms that allow, once instantiated the verification key, to generate commitments that are perfectly binding with some fixed probability. The algorithms are specified below.

Algorithm pbSetup($1^\kappa, c$): Given the security parameter κ , and a constant $c \in \mathbb{N}$ such that $1 \leq c \leq \log \kappa$, output a verification key $\bar{\vartheta} := (\vartheta_1, \dots, \vartheta_t)$ and a trapdoor $\bar{\tau} := (\tau_1, \dots, \tau_t)$ defined as follows:

1. Choose I^* random subset of $[t]$ with cardinality $n := 2^{-c}t$ and for all $i \in [t]$ let:

$$\begin{cases} \vartheta_i \leftarrow \text{Setup}(1^\kappa) & \text{if } i \in I^* \\ (\vartheta_i, \tau_i) \leftarrow \text{ESetup}(1^\kappa) & \text{else;} \end{cases}$$

2. Let $\bar{\tau} := (I^*, \{\tau_i\}_{i \in [t] \setminus I^*})$.

Algorithm pbECommit($\bar{\vartheta}, \bar{\tau}$): Choose indexes $\mathbf{j} := (j_1, \dots, j_p) \leftarrow [t]^p$; if \mathbf{j} is a subset of I^* return \perp . Otherwise w.l.o.g. let j_p be the index not in I^* ; the algorithm is identical to ECommit_{ss} .

Algorithm pbEquiv($\bar{\tau}, m, \bar{r}'$): Parse $\bar{r}' = (\mathbf{j}, \mathbf{s}, \mathbf{r}')$; in case \mathbf{j} is a subset of I^* return \perp . Otherwise the algorithm is identical to Equiv_{ss} .

Note that whenever pbECommit and pbEquiv do not return \perp , the trapdoor key $\bar{\tau}$ contains the trapdoor information τ_{j_p} which is needed by the algorithms ECommit_{ss} and Equiv_{ss} .

We say that a commitment $c\bar{om} = (\mathbf{j}, \text{com})$ is *binding* iff $\mathbf{j} \subseteq I^*$.

Lemma 5.1. *Let \mathcal{COM}_{ss} and $(\text{pbSetup}, \text{pbECommit}, \text{pbEquiv})$ be defined as above. For any $c \in \mathbb{N}$ such that $0 < c \leq \log \kappa$ the following holds.*

- (a) *The following probability distributions are computationally indistinguishable:*

$$\{\bar{\vartheta} \mid (\bar{\vartheta}, \bar{\tau}) \leftarrow \text{pbSetup}(1^\kappa)\} \left\{ \bar{\vartheta} \mid \bar{\vartheta} \leftarrow \text{Setup}_{ss}(1^\kappa) \right\}.$$

- (b) *For all $(\bar{\vartheta}, \bar{\tau}) \leftarrow \text{pbSetup}(1^\kappa)$ and for all $m \in \mathbb{F}^\mu$ the following probability distributions are statistically indistinguishable:*

$$\left\{ (c\bar{om}, \bar{r}) \mid \begin{array}{l} \bar{r} \leftarrow \mathcal{R}_{ss}, \\ c\bar{om} := \text{Commit}_{ss}(\bar{\vartheta}, m; \bar{r}), \\ c\bar{om} \text{ is not binding} \end{array} \right\}$$

$$\left\{ (c\bar{om}, \bar{r}) \mid \begin{array}{l} (c\bar{om}, \bar{r}') \leftarrow \text{pbECommit}(\bar{\vartheta}, \bar{\tau}), \\ \bar{r} := \text{pbEquiv}(\bar{\tau}, m, \bar{r}') \\ \text{pbECommit doesn't return } \perp \end{array} \right\}$$

(c) For all PPT adversaries A there exists a negligible function $\text{negl}(\kappa)$ such that the following holds:

$$\left| \mathbb{P} \left[\begin{array}{l} (\bar{\vartheta}, \bar{\tau}) \leftarrow_{\$} \text{pbSetup}(1^\kappa, c); \\ \text{com} \text{ is binding s.t. } (\text{com}, m, \bar{r}) \leftarrow_{\$} A(\bar{\vartheta}); \\ \text{com} = \text{Commit}_{ss}(\bar{\vartheta}, m, \bar{r}) \end{array} \right] - \frac{1}{\kappa^c} \right| \leq \text{negl}(\kappa).$$

(d) For any $(\bar{\vartheta}, \bar{\tau}) \leftarrow_{\$} \text{pbSetup}(1^\kappa, c)$ the probability over the randomness \bar{r}' that $\text{pbECommit}(\bar{\vartheta}; \bar{r}')$ aborts is $1/\kappa^c$.

Proof. Properties (a) and (b) follow, respectively, from the hybridness and the trapdoor hiding property of \mathcal{COM}_{ss} . To prove property (c), consider the following probability:

$$\mathbb{P} \left[\begin{array}{l} \bar{\vartheta} \leftarrow_{\$} \text{Setup}_{ss}(1^\kappa), I^* \leftarrow_{\$} \binom{[t]}{n}; \\ \text{com} \text{ is binding s.t. } (\text{com}, m, \bar{r}) \leftarrow_{\$} A(\bar{\vartheta}); \\ \text{com} = \text{Commit}_{ss}(\bar{\vartheta}, m; \bar{r}); \end{array} \right].$$

Recall that com is binding iff $\mathbf{j} \subseteq I^*$. Since I^* and the verification key $\bar{\vartheta}$ are uniformly and independently sampled, the above probability is upper bounded by

$$\left(\frac{n}{t}\right)^p = 2^{-c \log \kappa} = \frac{1}{\kappa^c}. \quad (12)$$

Property (c) easily follows by the previous equation and Property (a).

Property (d) follows directly from Eq. (12), as the algorithm ECommit_{ss} chooses the indexes \mathbf{j} uniformly at random. \square

5.2 The Signature Scheme

Let $\mathcal{COM} := \mathcal{COM}_1$ be a trapdoor hiding, linearly homomorphic commitment scheme (like the one used in the construction from Section 4.1). Let \mathcal{COM}_2 be a hybrid, linearly homomorphic commitment scheme, and \mathcal{COM}_{ss} be the corresponding secret sharing hybrid commitment as described in Section 5.1. We modify the scheme \mathcal{SS}_1 from Section 4.1 in two ways: (i) we strip off the OR-construction; and (ii) we replace the commitment com with a secret sharing hybrid commitment. Our scheme $\mathcal{SS}_3 = (\text{KGen}, \text{Sign}, \text{Verify})$ has message space equal to \mathbb{F} and is described below:

Key Generation. Let $t, d, \mu \in \mathbb{N}$ be parameters. Run $\text{crs} \leftarrow \text{Init}(1^\kappa)$, sample $\vartheta \leftarrow \text{Setup}(1^\kappa)$, and $\bar{\vartheta} := (\vartheta_1, \dots, \vartheta_t) \leftarrow_{\$} \text{Setup}_{ss}(1^\kappa)$. Sample $\Delta \leftarrow_{\$} (\mathbb{F}^\mu)^{d+1}$ and $\mathbf{r} = (r_0, \dots, r_d) \leftarrow_{\$} \mathbb{F}^{d+1}$, and compute commitments $\text{com}_i = \text{Commit}(\vartheta, \delta_i; r_i)$ for $i \in [0, d]$, where $\delta_i \in \mathbb{F}^\mu$ is the i -th column of Δ . Output

$$sk = (\Delta, \mathbf{r}) \quad vk = (\text{crs}, \vartheta, \{\vartheta_i\}_{i=1}^t, \{\text{com}_i\}_{i=0}^d).$$

Signature. For $j \in [\mu]$, let $\delta_j(X)$, $r(X)$, and $\Delta(X)$ be defined as in the scheme from Section 4.1. Consider the following polynomial-time relation:

$$\mathfrak{R} := \left\{ (\vartheta, \bar{\vartheta}, \text{com}, \text{com}); (\tilde{m}, \tilde{r}, \bar{r}) \left| \begin{array}{l} \text{com} = \text{Commit}(\vartheta, \tilde{m}; \tilde{r}) \\ \text{com} = \text{Commit}_{ss}(\bar{\vartheta}, \tilde{m}; \bar{r}) \end{array} \right. \right\}.$$

To sign a message $m \in \mathbb{F}$ compute $\tilde{m} = \Delta(m)$ and $\tilde{r} = r(m)$, and let $\text{com} := \text{Commit}(\vartheta, \tilde{m}; \tilde{r})$ and $\text{com} := \text{Commit}_{ss}(\bar{\vartheta}, \tilde{m}; \bar{r})$ where $\bar{r} \leftarrow_{\$} \mathcal{R}_{ss}$. Using crs as common reference string, generate a NIWI argument π for $(\vartheta, \bar{\vartheta}, \text{com}, \text{com}) \in \mathcal{L}_{\mathfrak{R}}$, the language generated by the above relation \mathfrak{R} . Output $\sigma = (\pi, \text{com})$.

Verification. Given a pair (m, σ) , parse σ as $\sigma = (\pi, \text{com})$. Output the same as:

$$\text{Ver}(\text{crs}, \pi, (\vartheta, \bar{\vartheta}, \text{com}, \text{com})).$$

Theorem 5.1. *Let $\mu \in \mathbb{N}$, and let \mathbb{F} be a finite field of size $\log |\mathbb{F}| = \kappa$ for security parameter $\kappa \in \mathbb{N}$. For any constant $0 < \xi \leq \mu/(\mu + 1)$ and $t = O(\kappa)$ the above signature scheme is $((\frac{\mu}{\mu+1} - \xi)|sk|, q_s, O(\frac{1}{q_s \cdot \log \kappa}), \text{negl}(\kappa))$ -fully-leakage one-more unforgeable.*

Notice that, when μ grows, the leakage rate asymptotically goes to 1. The slack parameter $\gamma(\kappa, q_s)$ is linear in the number of signature queries, but still polynomial in the security parameter as $q_s = \text{poly}(\kappa)$. As shown in [51, Section 5] this is good enough for some applications.

5.3 Proof Sketch

The proof of Theorem 5.1 requires an hybrid argument, along the same lines of the proof of Theorem 4.1. Since many of the hybrids are the same, we only give a sketch of the proof here.

Let A be an adversary asking $q_s = \kappa^c$ (for some constant c) signature queries, such that

$$\mathbb{P}[\text{Exp}_{\mathcal{SS}_3, A}^{\text{one-more}}(\kappa, \ell, q_s, \gamma) = 1] = \varepsilon(\kappa).$$

Experiment \mathcal{H}_0 : This experiment is exactly the same as the original experiment defining fully-leakage one-more unforgeability of \mathcal{SS}_3 . Note that the initial state (after key generation), is equal to $st := ((\Delta, \mathbf{r}), \text{crs}, \vartheta, \bar{\vartheta})$. (Recall that we assume oblivious sampling for both the NIWI and the commitment scheme.) At each invocation of the signing oracle $\text{Sign}(sk, \cdot)$ upon input a message m , the experiment produces a signature $\sigma = (\pi, c\bar{om})$ following the signing algorithm, and appends the randomness used to generate σ to the current state, i.e. $st := st \cup (\bar{r}, r_{niwi})$, where \bar{r} is the randomness used to generate the commitment $c\bar{om}$ and r_{niwi} is the randomness used to generate the argument π .

Denote with $((m_1^*, (\pi_1^*, com_1^*)), \dots, (m_n^*, (\pi_n^*, com_n^*)))$ the set of forgeries returned by the adversary at the end of the experiment. Let Forge_0 and False_0 be defined as in the proof of Theorem 4.1, and define $\varepsilon_0 := \mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0]$.

Claim 5.1. $\varepsilon_0 \geq \varepsilon - \text{negl}(\kappa)$.

Proof. The proof is similar to the proof of Claim 4.1 and is therefore omitted. \square

For each hybrid \mathcal{H}_i , described below, we write ε_i for the probability of the event $\text{Forge}_i \wedge \neg \text{False}_i$, and View_i for the view of the adversary at the end of the experiment.

Experiment $\mathcal{H}_{1,a}$: This experiment is the same as \mathcal{H}_0 , except that the verification key is computed using the probable binding key generation. Namely, we run $(\bar{\vartheta}, \bar{\tau}) \leftarrow_{\text{s}} \text{pbSetup}(1^\kappa, c)$ (recall that the number of signature queries made by A is $q_s = \kappa^c$). Note that from the point of view of the adversary, the initial state remains equal to $st := ((\Delta, \mathbf{r}), \text{crs}, \vartheta, \bar{\vartheta})$.

Claim 5.2. $\varepsilon_{1,a} \geq \varepsilon_0 - \text{negl}(\kappa)$.

Proof. Follows readily from Property (a) of Lemma 5.1. \square

Experiment $\mathcal{H}_{1,b}$: Identical to the same hybrid in the proof of Theorem 4.1.

Claim 5.3. $\varepsilon_{1,b} \geq \varepsilon_{1,a} - \text{negl}(\kappa)$.

Proof. Similar to the proof of Claim 4.3, and therefore omitted. \square

Experiment $\mathcal{H}_{2,d}$: In this experiment we modify the way signature queries are answered. Given such a query m , the experiment runs $\text{pbECommit}(\bar{\vartheta})$; in case the algorithm returns \perp the query is answered as before. Otherwise let $(c\bar{om}, \bar{r}')$ be the output of $\text{pbECommit}(\bar{\vartheta})$; the answer to the signature query uses $c\bar{om}$ as commitment. Observe that in the latter

case the randomness \bar{r} used to compute $c\bar{0}m$ in the previous hybrid can be obtained as a function of the matrix Δ by running the equivocation algorithm. In particular:

$$\bar{r}(\Delta) := \begin{cases} \bar{r} \leftarrow \mathcal{R}_{ss} & \text{if } \text{ECommit}_{ss}(\bar{\vartheta}) \text{ returns } \perp \\ \text{Equiv}_{ss}(\bar{r}, \Delta(m), \bar{r}') & \text{otherwise.} \end{cases}$$

The state is set to $st := st \cup \{(\Delta, \mathbf{r}), \bar{r}(\Delta), r_{niwi}\}$, where r_{niwi} is the randomness used to generate the NIWI argument π .

Claim 5.4. $\varepsilon_{2,d} \geq \varepsilon_{2,b} - \text{negl}(\kappa)$.

Proof. Follows readily from Property (b) of Lemma 5.1. \square

Experiment $\mathcal{H}_{2,e}$: This experiment is identical to the previous hybrid, except that (whenever possible) it uses a different witness to compute the argument corresponding to a signature query. In particular, given such a query $m \in \mathbb{F}$ for which $\text{pbECommit}(\bar{\vartheta})$ does not return \perp , it generates the argument π by running

$$\text{Prove}(\text{crs}, \underbrace{(\vartheta, \bar{\vartheta}, c\bar{0}m, c\bar{0}m)}_{\alpha}, \underbrace{(0, \text{Equiv}(\tau, 0^\mu, r'(m)), \text{pbEquiv}(\bar{r}, 0^\mu, \bar{r}'))}_{\beta'}, r'_{niwi}),$$

where $r'(m) = \sum_{i=0}^d r'_i \cdot m^i$ is computed using the randomness $\{r'_i\}_{i=0}^d$. Notice that the randomness r_{niwi} used to generate the NIWI argument in the previous experiment can be sampled (inefficiently) as a function of the witness $\beta := (\Delta(m), \text{Equiv}(\tau, \Delta(m), r'(m)), \text{pbEquiv}(\bar{r}, \Delta(m), \bar{r}'))$ and r'_{niwi} . In particular:

$$r_{niwi}(\Delta) := \begin{cases} r_{niwi} \text{ (as sampled by the prover)} & \text{if } \text{pbECommit}(\bar{\vartheta}) \text{ returns } \perp \\ \text{Rand}(\pi, r'_{niwi}, \beta) & \text{otherwise,} \end{cases}$$

where Rand is the reconstruction algorithm of Lemma 4.1. The state is updated as in $st(\Delta) = st \cup \{r(\Delta), r_{niwi}(\Delta)\}$.

Claim 5.5. $\varepsilon_{2,e} \geq \varepsilon_{2,d} - \text{negl}(\kappa)$.

Proof. Similar to the proof of Claim 4.9, and therefore omitted. \square

Experiment $\mathcal{H}_3^{\text{Leak}\Delta(\cdot)}$: This experiment is the same as the previous hybrid, with the difference that Δ is not sampled by the hybrid as part of the signing key, but can instead be accessed via $\text{Leak}\Delta(\cdot)$. Note that each signature query $m \in \mathbb{F}$ for which $\text{pbECommit}(\bar{\vartheta})$ happens to not return \perp can be answered without knowing Δ . Moreover:

- For each signature query $m \in \mathbb{F}$ where $\text{pbECommit}(\bar{\vartheta})$ returns \perp , the experiment defines the leakage function $f'_m := \Delta(m)$ and queries $\text{Leak}\Delta(\cdot)$ on f'_m . We call such query a *bad* query and we let Bad denote the corresponding event. Given the value $\Delta(m)$, the signature σ can be computed as in $\mathcal{H}_{2,e}$.
- For each leakage query f the experiment defines a function f'' and queries $\text{Leak}\Delta(\cdot)$ on f'' . The function f'' hard-wires all values necessary to reconstruct the current state $st = st(\Delta)$; this includes the trapdoor information needed in order to reconstruct \mathbf{r} , and all random coins used to simulate previous signature queries.

Claim 5.6. $\varepsilon_3 = \varepsilon_{2,e}$.

Proof. The only difference between the two hybrids stems from the fact that Δ is only accessed via a leakage oracle in \mathcal{H}_3 . Note that in $\mathcal{H}_{2,e}$ the matrix Δ is needed only to answer bad signature queries and leakage queries. In case of a bad query, \mathcal{H}_3 accesses the leakage oracle to additionally obtain $\Delta(m)$ via f'_m ; given $\Delta(m)$ the signature computation is identical in the two hybrids. The distribution of the leakage is also identical in the two hybrids, as the function f'' perfectly reconstructs the current state $st(\Delta)$. \square

Lemma 5.2. *For any $\beta > 0$ we have that*

$$\mathbb{P}[\mathbb{H}_\infty(\Delta | \text{View}_3 = v) \geq |\Delta| - 2e\mu \log |\mathbb{F}| - \beta - \lambda] \geq 2^{-\beta} + 2^{-2e}, \quad (13)$$

where the probability is taken over the randomness of the experiment.

Proof. The proof is a straightforward adaptation of the proof of Lemma 4.2. The only difference is in the computation of the probability of the events Bad_i , which is now upper bounded by $1/\kappa^c = 1/q_s$ by Property (d) of Lemma 5.1. Note that this gives $\mathbb{E}[Z] = 1$. \square

Lemma 5.3. *For any adversary A such that $\mathbb{P}[\text{Exp}_{\mathcal{SS}, A}^{\text{one-more}}(\kappa, \ell, q_s, \gamma) = 1] = \varepsilon$, there exists a constant $c' > 0$ such that:*

$$\mathbb{E}[\mathbb{H}_\infty(\Delta | \text{View}_3 = v)] \leq \left(d + 1 - \frac{n}{q_s}\right) \mu \log |\mathbb{F}| - c' \log \varepsilon(\kappa),$$

where the expectation is taken over the randomness of the experiment.

Proof. The proof is a straightforward adaptation of the proof of Lemma 4.3. The only difference is in the computation of $\mathbb{E}[X + Z]$. Applying Property (c) and Property (d) of Lemma 4.2, we get $\mathbb{E}[X + Z] \geq n(1/q_s - \text{negl}(\kappa)) + 1 \geq n/q_s$. \square

Given the two lemmas above the proof of Theorem 5.1 can be derived following the same strategy as for the proof of Theorem 4.1. In particular, one can verify that to obtain a contradiction we need that:

$$\frac{n\kappa}{q_s} = \Omega(\lambda - \log \varepsilon(\kappa)). \quad (14)$$

Whenever $\varepsilon(\kappa)$ is noticeable, since $n \geq \lfloor \frac{\lambda}{\gamma \cdot s} \rfloor + 1$, and $s = \Theta(\kappa \log \kappa)$, in order for Eq. (14) to hold it suffices to set $\gamma = O(\frac{1}{q_s \cdot \log \kappa})$. This concludes the proof sketch.

6 A Scheme in the Bounded Retrieval Model

In this section we describe a fully-leakage one-more unforgeable signature scheme in the Bounded Retrieval Model (BRM) [15, 26, 27]. The BRM imposes additional efficiency constraints on the size of the public key and the computational time of the signer and verifier algorithm. The scheme is in the random oracle model [4]. To define fully-leakage one-more unforgeability in the BRM, we let the key generation, signature and verification algorithms additionally be parametrized by the leakage parameter ℓ .

Definition 6.1 (Fully-leakage one-more unforgeability in the BRM). *A signature scheme $\mathcal{SS} = (\text{KGen}, \text{Sign}, \text{Verify})$ is (ℓ, γ) -fully-leakage one-more unforgeable in the BRM if the following conditions are met:*

- *The verification key size, signature size, signing-time and verification-time (and the number of secret-key bits read during signature computation) are independent of the leakage-bound ℓ . More formally, there exist polynomials vksize , sgsize , sigT , verT , such that, for any polynomial ℓ , and any $(vk, sk) \leftarrow \text{KGen}(1^\kappa, 1^{\ell(\kappa)})$, $m \in \mathcal{M}$, $\sigma \leftarrow \text{Sign}(sk, m)$, we have that:*
 - *Verification key size is $|vk| = O(\text{vksize}(\kappa))$, signature size is $|\sigma| = O(\text{sgsize}(\kappa, |m|))$.*
 - *Run-time of $\text{Sign}(sk, m)$ and the number of bits of sk accessed is $O(\text{sigT}(\kappa, |m|))$.*
 - *Run-time of $\text{Verify}(vk, m, \sigma)$ is $O(\text{verT}(\kappa, |m|))$.*
- *The scheme \mathcal{SS} is (ℓ, γ) -fully-leakage one-more unforgeable according to Definition 3.1.*

6.1 Ingredients

Let \mathbb{F} be a finite field with additive identity 0. For two parameter d and t , with $d > t$, let us define a distribution $\text{Sparse}(\mathbb{F}, d, t)$ by sampling a vector $\mathbf{c} = (c_1, \dots, c_d)$ such that for all $i \in [d]$: (i) $c_i \leftarrow_s \mathbb{F}$ with probability t/d , and (ii) $c_i = 0$ with probability $1 - t/d$. We call t the locality parameter.

We will rely on the following lemma.

Lemma 6.1. *Let $\kappa = \log |\mathbb{F}|$, for security parameter $\kappa \in \mathbb{N}$. For any subspace $A \subset \mathbb{F}^d$ of dimension less than d the probability that a vector \mathbf{c} randomly chosen over $\text{Sparse}(\mathbb{F}, d, t)$ lies in the subspace A is at most $\frac{1}{|\mathbb{F}|} + e^{-t}$.*

Proof. Note that the bigger A is, the more likely $\mathbf{c} \in A$, therefore the worst case is $\dim(A) = d - 1$. If $\mathbf{c} \in A$ then there exist $\alpha_1, \dots, \alpha_d \in \mathbb{F}$ with at least one coordinate i s.t. $\alpha_i \neq 0$, such that $\mathbf{c} \in A$ if and only if $E_d = 1$ where E_d is the event defined as

$$E_d := \begin{cases} 1 & \text{if } \sum_{i=0}^d \alpha_i \cdot c_i = 0 \\ 0 & \text{otherwise.} \end{cases}$$

For any $d' \leq d$ we have:

$$\begin{aligned} \mathbb{P}[E_{d'}] &\leq \sum_{x \in \mathbb{F}} \mathbb{P}[c_{d'} = x] \mathbb{P}\left[\sum_{i=1}^{d'-1} \alpha_i \cdot c_i = -\alpha_{d'} \cdot x\right] \\ &\leq \mathbb{P}[c_{d'} = 0] \mathbb{P}[E_{d'-1}] + \sum_{x \in \mathbb{F} \setminus \{0\}} \mathbb{P}[c_{d'} = x] \cdot \mathbb{P}\left[\sum_{i=1}^{d'-1} \alpha_i \cdot c_i = -\alpha_{d'} \cdot x\right] \\ &\leq \left(\frac{t \cdot 2^{-\kappa}}{d} + 1 - \frac{t}{d}\right) \mathbb{P}[E_{d'-1}] + \frac{t \cdot 2^{-\kappa}}{d} \cdot \sum_{x \in \mathbb{F} \setminus \{0\}} \mathbb{P}\left[\sum_{i=1}^{d'-1} \alpha_i \cdot c_i = -\alpha_{d'} \cdot x\right] \\ &\leq \left(\frac{t \cdot 2^{-\kappa}}{d} + 1 - \frac{t}{d}\right) \mathbb{P}[E_{d'-1}] + \frac{t \cdot 2^{-\kappa}}{d} \cdot (1 - \mathbb{P}[E_{d'-1}]) \\ &\leq \left(1 - \frac{t}{d}\right) \mathbb{P}[E_{d'-1}] + \frac{t \cdot 2^{-\kappa}}{d}. \end{aligned}$$

From the last inequality, by setting $d' := d$, it follows that

$$\begin{aligned} \mathbb{P}[E_d] &\leq \left(1 - \frac{t}{d}\right)^d \mathbb{P}[E_0] + \frac{t \cdot 2^{-\kappa}}{d} \left(\sum_{i=0}^{d-1} \left(1 - \frac{t}{d}\right)^i\right) \\ &\leq e^{-t} \cdot \mathbb{P}[E_0] + \frac{t \cdot 2^{-\kappa}}{d} \left(\frac{1 - (1 - \frac{t}{d})^d}{t/d}\right) \\ &\leq e^{-t} + \frac{1}{|\mathbb{F}|}. \end{aligned}$$

□

Our construction in the BRM is based on the following ingredients:

- Two random oracles $\mathcal{H}_0 : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and $\mathcal{H}_1^{d,t} : \{0, 1\}^* \rightarrow \text{Sparse}(\mathbb{F}, d, t)$.⁹

⁹The two random oracles can be derived by a single random oracle \mathcal{H} , e.g., by letting $\mathcal{H}_0 := \mathcal{H}(0, \cdot)$ and $\mathcal{H}_1 := \mathcal{H}(1, \cdot)$.

- A perfectly hiding, linearly homomorphic commitment scheme $\mathcal{COM} = (\text{Setup}, \text{Commit})$, with message space equal to \mathbb{F}^μ and randomness space equal to \mathbb{F} (for a finite field \mathbb{F} and a parameter $\mu \in \mathbb{N}$).
- A perfect zero-knowledge proof of knowledge system $\mathcal{NIZK} = (\text{Prove}, \text{Ver})$ with an on-line extractor (w.r.t. the random oracle \mathcal{H}_0), for proving knowledge of the opening of a commitment.

We review the definition of a NIZK PoK with an online extractor below (see, e.g., [52, 32]).

Definition 6.2 (Online Extractability). *We say that $\mathcal{NIZK} = (\text{Prove}, \text{Ver})$ is a NIZK PoK for a relation \mathfrak{R} (w.r.t. a random oracle \mathcal{H}) if the following holds:*

Perfect Zero-Knowledge. *There exist a pair of PPT algorithms $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ (a.k.a. the zero-knowledge simulator), such that for any pair of PPT distinguishers $\text{D} = (\text{D}_0, \text{D}_1)$ the following distributions are identical:*

- Let $(x, w, st_{\text{D}}) \leftarrow_{\$} \text{D}_0^{\mathcal{H}}(1^\kappa)$, and $\pi \leftarrow_{\$} \text{Prove}^{\mathcal{H}}(x, w)$ if $(x, w) \in \mathfrak{R}$ and $\pi := \perp$ otherwise. Output $\text{D}_1^{\mathcal{H}}(\pi, st_{\text{D}})$.
- Let $(\mathcal{H}_0, st_{\text{Sim}}) \leftarrow_{\$} \text{Sim}_0(1^\kappa)$, $(x, w, st_{\text{D}}) \leftarrow_{\$} \text{D}_0^{\mathcal{H}_0}(1^\kappa)$, and $(\mathcal{H}_1, \pi) \leftarrow_{\$} \text{Sim}_1(st_{\text{Sim}}, x, \text{yes})$ if $(x, w) \in \mathfrak{R}$ and $(\mathcal{H}_1, \pi) \leftarrow_{\$} \text{Sim}_1(st_{\text{Sim}}, x, \text{no})$ otherwise. Output $\text{D}_1^{\mathcal{H}_1}(\pi, st_{\text{D}})$.

Online Extractability. *There exist a probabilistic polynomial-time algorithm K (a.k.a. the on-line extractor), such that the following holds for any algorithm A . Let $(x, \pi) \leftarrow_{\$} \text{A}^{\mathcal{H}}(1^\kappa)$ and $\mathcal{Q}_{\mathcal{H}}(\text{A})$ be the sequence of queries of A to \mathcal{H} and \mathcal{H} 's answers. Let $w \leftarrow_{\$} \text{K}(x, \pi, \mathcal{Q}_{\mathcal{H}}(\text{A}))$. Then, as a function of κ ,*

$$\mathbb{P}[(x, w) \notin \mathfrak{R} \wedge \text{Ver}^{\mathcal{H}}(x, \pi) = 1] \leq \text{negl}(\kappa).$$

6.2 Basic Construction

We first describe a scheme that meets partially the efficiency requirements of the BRM; namely, the signing and verification algorithms have running time independent of the leakage parameter ℓ , but the length of the verification key still depends on ℓ . Later we adopt the technique of Alwen *et al.* [2] to obtain a full-fledged scheme in the BRM.

Our scheme $\mathcal{SS}_4 = (\text{KGen}, \text{Sign}, \text{Verify})$ has $\mathcal{M} = \{0, 1\}^*$ and is described below:

Key Generation. Let $t, d, \mu \in \mathbb{N}$ be parameters; choose d to be large enough such that $\ell(\kappa) < d\mu \cdot \kappa$. Run $\vartheta \leftarrow \text{Setup}(1^\kappa)$. Sample $\Delta \leftarrow_{\$} \mathbb{F}^{\mu \times d}$ and $\mathbf{r} = (r_1, \dots, r_d) \leftarrow_{\$} \mathbb{F}^d$, and compute commitments $com_i = \text{Commit}(\vartheta, \delta_i; r_i)$ for $i \in [d]$, where $\delta_i \in \mathbb{F}^\mu$ is the i -th column of Δ . Output

$$sk = (\Delta, \mathbf{r}) \quad vk = (\vartheta, \{com_i\}_{i=1}^d).$$

Signature. Consider the following polynomial-time relation:

$$\mathfrak{R} := \{(\vartheta, c\tilde{m}); (\tilde{m}, \tilde{r}) \mid c\tilde{m} = \text{Commit}(\vartheta, \tilde{m}; \tilde{r})\}.$$

To sign a message $m \in \{0, 1\}^*$, let $\mathbf{c} := \mathcal{H}_1^{d,t}(m)$ and compute $\Delta(m) := \Delta \cdot \mathbf{c}^\top$ and $r(m) := \mathbf{r} \cdot \mathbf{c}^\top$. Generate a proof $\pi \leftarrow_{\$} \text{Prove}^{\mathcal{H}_0}((\vartheta, c\tilde{m}), (\Delta(m), r(m)))$, where $c\tilde{m} := \text{Commit}(\vartheta, \Delta(m); r(m))$. Output $\sigma := \pi$.

Verification. Given a pair (m, σ) , parse σ as $\sigma = \pi$ and let $\mathbf{c} := \mathcal{H}_1^{d,t}(m) = (c_1, \dots, c_d)$. Output the same as:

$$\text{Ver}^{\mathcal{H}_0}(\pi, (\vartheta, \prod_{i=1}^d (com_i)^{c_i})).$$

Notice that for $t = O(\kappa)$ the efficiency constraint on the running time of signer and verifier are met only *on average*. If the signature scheme needs to be *always* local, the signer can just re-sign the message whenever the number of non-zero locations in \mathbf{c} is not in the range $\{(1 \pm \varepsilon)t\}$, for a constant ε .

Theorem 6.1. *Let \mathbb{F} be a finite field such that $\log |\mathbb{F}| = \kappa$, for security parameter $\kappa \in \mathbb{N}$. For any $\xi = \omega(\log \kappa)$, whenever $\gamma = O(1)$ the above signature scheme is $((\frac{\mu}{\mu+1})|sk| - \xi, \gamma)$ -fully-leakage one-more unforgeable in the random oracle model.*

We reduce to the computational binding property of the commitment scheme. The reduction simulates the entire experiment by sampling a legitimate secret key (Δ, \mathbf{r}) . Whenever the adversary outputs its forgeries $(m_1^*, \pi_1^*), \dots, (m_n^*, \pi_n^*)$, the reduction chooses an index $i \leftarrow_s [n]$ uniformly at random and extracts the proof π_i^* obtaining a valid witness $(\tilde{m}_i^*, \tilde{r}_i^*)$. Finally the reduction outputs $(\tilde{c}m_i, (\tilde{m}_i^*, \tilde{r}_i^*), (\Delta(m_i^*), r(m_i^*)))$.

It remains to argue that the conditional average min-entropy of Δ given the view of the adversary is high. By a counting argument, this implies that there exists an index $i \in [n]$ for which the min-entropy of the witness sitting behind each of the forged signatures is larger than 1, so that the reduction breaks the computational binding property with noticeable probability.

Proof. Let \mathbf{A} be an adversary in the fully-leakage one-more unforgeability experiment. Define with Forge the event that \mathbf{A} makes the experiment output 1, and assume $\mathbb{P}[\text{Forge}] = \varepsilon$.

Consider the following reduction algorithm \mathbf{B} (based on \mathbf{A}), attacking the computational binding property of \mathcal{COM} . Adversary \mathbf{B} works as follows:

1. At the beginning \mathbf{B} receives the public key $\vartheta \leftarrow \text{Setup}(1^\kappa)$ for the commitment scheme. Hence, it generates a legitimate signing/verification key pair (sk, vk) for \mathcal{SS}_4 , where $sk = (\Delta, \mathbf{r})$ and $vk = (\vartheta, \{com_i\}_{i=1}^d)$.
2. \mathbf{B} takes care of \mathbf{A} 's queries to the random oracle, and keeps a list $\mathcal{Q}_{\mathcal{H}}(\mathbf{A})$ of such queries and the relative answers.
3. Given the above setup, \mathbf{B} answers signature queries and leakage queries asked by \mathbf{A} in the natural way. Every time a signature is computed, the state is updated as $st := st \cup \{r_{nizk}\}$, where r_{nizk} is the randomness required to generate each signature.
4. Eventually \mathbf{A} returns a set of forgeries $(m_1^*, \pi_1^*), \dots, (m_n^*, \pi_n^*)$. At this point \mathbf{B} runs $\text{K}(\vartheta, \tilde{c}m_i^*, \mathcal{Q}(\mathbf{A}))$ for all $i \in [n]$, where $\tilde{c}m_i^* = \text{Commit}(\vartheta, \Delta(m_i^*), r(m_i^*))$, obtaining a witness $(\tilde{m}_i^*, \tilde{r}_i^*)_{i \in [n]}$. In case there exists an index $i \in [n]$ such that $(\tilde{c}m_i^*, (\tilde{m}_i^*, \tilde{r}_i^*)) \notin \mathfrak{R}$, the reduction \mathbf{B} aborts.
5. \mathbf{B} finds an index i such that $(\Delta(m_i^*), r(m_i^*)) \neq (\tilde{m}_i^*, \tilde{r}_i^*)$. If no such index is found, \mathbf{B} aborts. Otherwise \mathbf{B} outputs $(\tilde{c}m_i^*, (\tilde{m}_i^*, \tilde{r}_i^*), (\Delta(m_i^*), r(m_i^*)))$.

We say that \mathbf{B} *wins* if it breaks the computational binding property of \mathcal{COM} . Let Abort_1 (resp. Abort_2) be the event that \mathbf{B} aborts in step 4 (resp. step 5); define $\text{Abort} := \text{Abort}_1 \vee \text{Abort}_2$. By definition of \mathbf{B} :

$$\begin{aligned} \mathbb{P}[\mathbf{B} \text{ wins}] &= \mathbb{P}[\text{Forge} \wedge \neg \text{Abort}_1 \wedge \neg \text{Abort}_2] \\ &\geq \mathbb{P}[\text{Forge} \wedge \neg \text{Abort}_1] - \mathbb{P}[\text{Abort}_2] \\ &\geq \mathbb{P}[\text{Forge}] - \mathbb{P}[\text{Forge} \wedge \text{Abort}_1] - \mathbb{P}[\text{Abort}_2]. \end{aligned}$$

Claim 6.1. $\mathbb{P}[\text{Forge} \wedge \text{Abort}_1] \leq \text{negl}(\kappa)$.

Proof. Event Abort_1 holds true in case there is at least one index $i \in [n]$ such that K fails to extract a valid witness for $\tilde{c}m_i^*$. However, since Forge occurs, all proofs π_i^* are valid and thus there is at least one index $i \in [n]$ for which:

$$\text{Ver}(\vartheta, \tilde{c}m_i^*, \pi_i^*) = 1 \wedge ((\vartheta, \tilde{c}m_i^*), (\tilde{m}_i^*, \tilde{r}_i^*)) \notin \mathfrak{R},$$

which would contradict simulation extractability of the NIZK argument system. We conclude that $\mathbb{P}[\text{Forge} \wedge \text{Abort}_1]$ must be negligible. \square

Claim 6.2. $\mathbb{P}[\text{Abort}_2] \leq \text{negl}(\kappa)$.

Proof. Event Abort_2 holds true if for all $i \in [n]$ we have that $(\Delta(m_i^*), r(m_i^*)) = (\tilde{m}_i^*, \tilde{r}_i^*)$ for all values extracted by \mathbf{B} . Consider the matrix \mathbf{M} for which the i -th column is the vector $\mathbf{c}_i = \mathcal{H}_1^{d,t}(m_i^*)$. Let FullRank be the event that \mathbf{M} is full rank, and suppose that \mathbf{A} makes q_s signature queries. Note that the view of \mathbf{A} is a random variable View consisting of: (i) the verification key vk (ii) the leakage Λ , (iii) the signatures $(m_1, \sigma_1), \dots, (m_{q_s}, \sigma_{q_s})$, (iv) the set $\mathcal{Q}_{\mathcal{H}}(\mathbf{A})$ of random oracle calls and corresponding answers. We can write:

$$\mathbb{P}[\text{Abort}_2] \leq \mathbb{P}[\neg \text{FullRank}] + 2^{-\tilde{\mathbb{H}}_{\infty}(\Delta \mathbf{M}, \mathbf{rM} | \text{View}; \text{FullRank})},$$

where $\tilde{\mathbb{H}}_{\infty}(X|Y; E)$ is the average conditional min-entropy of the random variable X given the random variable Y and conditioned on the event E .

Let q_h be the number of random oracle queries made by \mathbf{A} , we claim that:

$$\mathbb{P}[\neg \text{FullRank}] \leq q_h \cdot \left(\frac{1}{|\mathbb{F}|} + e^{-t} \right).$$

This is because the first $d-1$ columns of \mathbf{M} form a subspace of dimension less than d , and for any subspace $A \subset \mathbb{F}^d$ of dimension less than d the probability that a vector \mathbf{c} randomly chosen over $\text{Sparse}(\mathbb{F}, d, t)$ lies in the subspace A is at most $\frac{1}{|\mathbb{F}|} + e^{-t}$ (by Lemma 6.1), and so the above inequality follows by the union bound. Hence,

$$\tilde{\mathbb{H}}_{\infty}(\Delta \mathbf{M}, \mathbf{rM} | \text{View}; \text{FullRank}) = \tilde{\mathbb{H}}_{\infty}(\Delta, \mathbf{r} | vk, \Lambda, \{m_i, \sigma_i\}_{i=1}^{q_s}, \mathcal{Q}_{\mathcal{H}}(\mathbf{A})) \quad (15)$$

$$= \tilde{\mathbb{H}}_{\infty}(sk | vk, \Lambda, \{m_i, \sigma_i\}_{i=1}^{q_s})$$

$$= \tilde{\mathbb{H}}_{\infty}(sk | vk, \Lambda) \quad (16)$$

$$\geq \tilde{\mathbb{H}}_{\infty}(sk | vk) - \lambda \quad (17)$$

$$\geq \rho |sk| - \lambda. \quad (18)$$

Where Eq. (15) holds by the fact that \mathbf{M} is full rank, Eq. (16) follows by perfect zero-knowledge of \mathcal{NIZK} , Eq. (17) follows by the chain rule on the average conditional min-entropy (see, e.g., [25, Lemma 2.2]), and Eq. (18) holds by the fact that sk is uniform over $\mathbb{F}^{d(\mu+1)}$. We conclude that:

$$\mathbb{P}[\text{Abort}_2] \leq \text{negl}(\kappa) + 2^{\rho |sk| - \lambda} \leq \text{negl}(\kappa) + 2^{-\omega(\log \kappa)} = \text{negl}(\kappa),$$

as desired. \square

Combining Claim 6.1 and Claim 6.2 we get that $\mathbb{P}[\mathbf{B} \text{ wins}] \geq \varepsilon - \text{negl}(\kappa)$, and thus ε must be negligible, finishing the proof. \square

6.3 The Full-Fledged Scheme

In order to reduce the size of the verification key in the scheme from the previous section, we rephrase the idea of Alwen et al. [2, Section 4.3] using the notion of homomorphic linear authenticator introduced in [3]. Let \mathbb{F} be a field equipped with a function $\text{Combine}(\mathbf{c}, \mathbf{f})$ that takes as input two vector $\mathbf{c} \in \mathbb{Z}^d$ and $\mathbf{f} \in \mathbb{F}^d$, and outputs a field element, concretely, we can think of it as the scalar multiplication $\sum_i f_i \cdot c_i$ or the scalar multiplication “at the exponents” $\prod_i f_i^{c_i}$.

Definition 6.3 (Homomorphic linear authenticator). *A public-key homomorphic linear authenticator for a field \mathbb{F} equipped with a function $\text{Combine}(\mathbf{c}, \mathbf{f})$ is a tuple of four PPT algorithms $\mathcal{HCA} = (\text{KGen}, \text{Tag}, \text{Auth}, \text{Ver})$ such that:*

$(PK, SK) \leftarrow \text{KGen}(1^\kappa)$ is a probabilistic algorithm used to set up the scheme. It takes as input the security parameter and the field description \mathbb{F} (where $\log |\mathbb{F}| = \kappa$), and outputs a public and private key pair (PK, SK) .

$(\mathbf{tag}, st) \leftarrow \text{Tag}(SK, \mathbf{f})$ is a probabilistic algorithm that is run in order to tag a vector. It takes as input a secret key SK and a vector $\mathbf{f} \in \mathbb{F}^d$, and outputs a vector of tags $\mathbf{tag} \in \mathbb{F}^d$ and state information st .

$\tau := \text{Auth}(PK, \mathbf{f}, \mathbf{tag}, \mathbf{c})$ is a deterministic algorithm that is run to generate tags. It takes as input a public key PK , a vector $\mathbf{f} \in \mathbb{F}^d$, a tag vector \mathbf{tag} , and a challenge vector $\mathbf{c} \in \mathbb{Z}^d$, and outputs a tag τ .

$b := \text{Ver}(PK, st, \gamma, \mathbf{c}, \tau)$ is a deterministic algorithm that is used to verify a tag. It takes as input a public key PK , state information st , an element $\gamma \in \mathbb{F}$, a challenge vector $\mathbf{c} \in \mathbb{Z}^d$, and a tag τ . It outputs a bit b , where $b = 1$ indicates acceptance and $b = 0$ indicates rejection.

For correctness, we require that for all $\kappa \in \mathbb{N}$, all (PK, SK) output by $\text{KGen}(1^\kappa)$, all $\mathbf{f} \in \mathbb{F}^d$, all (\mathbf{tag}, st) output by $\text{Tag}(SK, \mathbf{f})$, and all $\mathbf{c} \in \mathbb{Z}^d$, it holds that

$$\text{Ver}(PK, st, \text{Combine}(\mathbf{c}, \mathbf{f}), \mathbf{c}, \text{Auth}(PK, \mathbf{f}, \mathbf{tag}, \mathbf{c})) = 1.$$

Definition 6.4 (Unforgeability for public-key HLAs). Let \mathbf{A} be an adversary for the following experiment:

1. The challenger computes $(PK, SK) \leftarrow \text{KGen}(1^\kappa)$.
2. Given PK and oracle access to $\text{Tag}(SK, \cdot)$, adversary \mathbf{A} outputs a vector $\mathbf{f} \in \mathbb{F}^d$.
3. The challenger tags \mathbf{f} by computing $(\mathbf{tag}, st) \leftarrow \text{Tag}(SK, \mathbf{f})$.
4. Given \mathbf{tag} and st , the adversary \mathbf{A} outputs a challenge vector $\mathbf{c} \in \mathbb{Z}^d$, an element $\gamma' \in \mathbb{F}$, and a tag τ .
5. The adversary succeeds if $\gamma' \neq \text{Combine}(\mathbf{c}, \mathbf{f})$ and $\text{Ver}(PK, st, \gamma', \mathbf{c}, \tau) = 1$.

We say that \mathcal{HLA} is unforgeable if the success probability of every PPT adversary \mathbf{A} in the above experiment is negligible.

We are now ready to describe our signature scheme $\mathcal{SS}_5 = (\text{KGen}, \text{Sign}, \text{Verify})$, based on our construction from the previous section $\mathcal{SS}_4 = (\text{KGen}, \text{Sign}, \text{Verify})$ and a public-key homomorphic linear authenticator $\mathcal{HLA} = (\text{KGen}, \text{Tag}, \text{Auth}, \text{Ver})$.

Key Generation. Run $(vk', sk') \leftarrow \mathcal{SS}_4.\text{KGen}(1^\kappa)$ and $(PK, SK) \leftarrow \mathcal{HLA}.\text{KGen}(1^\kappa)$; parse vk' as $(\vartheta, \{com_i\}_{i=1}^d)$ and define $\mathbf{f} = (com_1, \dots, com_d)$. Generate $(\mathbf{tag}, st) \leftarrow \text{Tag}(SK, \mathbf{f})$ and output:

$$sk := (sk', \mathbf{tag}) \quad pk := (\vartheta, PK, st).$$

Signature. Run $\pi \leftarrow \mathcal{SS}_4.\text{Sign}(sk, m)$, and compute $c\tilde{om} = \prod_{i=1}^d com_i^{c_i}$ and $\tau := \text{Auth}(PK, \mathbf{f}, \mathbf{tag}, \mathbf{c})$ where $\mathbf{c} = (c_1, \dots, c_d) = \mathcal{H}_1^{d,t}(m)$. Output $\sigma := (c\tilde{om}, \tau, \pi)$.

Verification Given a pair (m, σ) , parse σ as $(c\tilde{om}, \tau, \pi)$. Compute $\mathbf{c} := \mathcal{H}_1^{d,t}(m)$ and output:

$$\mathcal{HLA}.\text{Ver}(PK, st, c\tilde{om}, \mathbf{c}, \tau) \wedge \text{Ver}^{\mathcal{H}_0}((\text{crs}, c\tilde{om}), \pi).$$

Theorem 6.2. Let \mathbb{F} be a finite field such that $\log |\mathbb{F}| = \kappa$, for security parameter $\kappa \in \mathbb{N}$. For any $\xi = \omega(\log \kappa)$, whenever $\gamma = O(1)$ the above signature scheme is $((\frac{\mu}{\mu+2})|sk| - \xi, \gamma)$ -fully-leakage one-more unforgeable in the bounded retrieval model under the condition that the adversary does not leak from the randomness used during key generation.

Note that the scheme does not allow leakage on the coins of the key generation; these coins have to be securely erased (together with SK). However, after key generation is over no more erasures are needed. We give only a proof sketch here:

Proof. Let A be an adversary in the fully-leakage one-more unforgeability experiment for \mathcal{SS}_5 .

We construct an adversary B that either breaks fully-leakage one-more unforgeability of \mathcal{SS}_4 (contradicting Theorem 6.1) or the unforgeability of \mathcal{HLA} . Adversary B plays both games at the same time, with challenger C_1 (for \mathcal{SS}_4) and C_2 (for \mathcal{HLA}). A description of B follows:

1. At the beginning B receives as input the verification key $vk' = (\vartheta, \{com_i\}_{i=1}^d)$ (from C_1), and the public key for the HLA PK (from C_2). Recall that there is no leakage during the key generation phase. Thus, B queries oracle $\text{Tag}(SK, \cdot)$ upon input the value $\mathbf{f} = (com_1, \dots, com_d)$, obtaining a pair (\mathbf{tag}, st) , and returns $vk = (\vartheta, PK, st)$ to A .
2. Upon input a signature query m from A , adversary B forwards this query to C_1 receiving back a signature π . It then computes $\mathbf{c} = (c_1, \dots, c_d) = \mathcal{H}_1^{d,t}(m)$, $c\tilde{om} = \prod_{i=1}^d com_i^{c_i}$, $\tau = \text{Auth}(PK, \mathbf{f}, \mathbf{tag}, \mathbf{c})$, and returns $(c\tilde{om}, \tau, \pi)$ to A .
3. Upon input a leakage query $f(\cdot)$ from A , adversary B defines a function $f'(\cdot) := f(\mathbf{tag}, \cdot)$ and forwards such query to C_1 . Note that this is a perfect simulation, as the random coins for signature generation consists only of the coins for the signing algorithm of \mathcal{SS}_4 .
4. Eventually A returns n forgeries $(m_1^*, (c\tilde{om}_1^*, \tau_1^*, \pi_1^*), \dots, m_n^*, (c\tilde{om}_n^*, \tau_n^*, \pi_n^*))$. At this point B checks, for $\mathbf{c}_i^* = \mathcal{H}_1^{d,t}(m_i^*)$, that there exists an index $i \in [n]$ such that

$$\text{Ver}(PK, st, c\tilde{om}_i^*, \mathbf{c}_i^*, \tau_i^*) = 1 \wedge c\tilde{om}_i^* \neq \text{Combine}(\mathbf{c}_i^*, \mathbf{f}). \quad (19)$$

In case such index is found, B outputs $(c\tilde{om}_i^*, \mathbf{c}_i^*, \tau_i^*)$ to C_2 . Otherwise, B outputs the forgeries $(m_1^*, \pi_1^*, \dots, m_n^*, \pi_n^*)$ to C_1 .

Let Fresh be the event that Eq. (19) holds for some $i \in [n]$. Moreover, let Forge be the event that A wins in the fully-leakage one-more unforgeability experiment. We can write:

$$\mathbb{P}[\text{Forge}] = \mathbb{P}[\text{Forge} \wedge \text{Fresh}] + \mathbb{P}[\text{Forge} \wedge \neg \text{Fresh}].$$

The two claims below conclude the proof.

Claim 6.3. $\mathbb{P}[\text{Forge} \wedge \text{Fresh}] \leq \text{negl}(\kappa)$.

Proof. One can easily see that in case A provokes both Forge and Fresh , adversary B is successful against C_2 contradicting unforgeability of the HLA. This is because $(c\tilde{om}_i^*, \mathbf{c}_i^*, \tau_i^*)$ is verified correctly, and moreover $c\tilde{om}_i^*$ is a fresh value different than $\text{Combine}(\mathbf{c}_i^*, \mathbf{f})$. \square

Claim 6.4. $\mathbb{P}[\text{Forge} \wedge \neg \text{Fresh}] \leq \text{negl}(\kappa)$.

Proof. One can easily see that in case A provokes both Forge and $\neg \text{Fresh}$, the forgeries output by B against C_1 are all accepting and with the right distribution, as, in particular, $c\tilde{om}_i^* = \text{Combine}(\mathbf{c}_i^*, \mathbf{f}) = \prod_{j=1}^d com_j^{c_i[j]}$. This contradicts fully-leakage one-more unforgeability of \mathcal{SS}_4 . \square

\square

7 Instantiations

In the following let \mathbb{G} and \mathbb{G}_T be groups of prime order q , supporting oblivious sampling of group elements, and let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a non-degenerate efficiently computable bilinear map. Set the field \mathbb{F} to be \mathbb{Z}_q . For ease of notation, given three vectors $\mathbf{g}, \mathbf{h} \in \mathbb{G}^n$ and $\boldsymbol{\alpha} \in \mathbb{F}^n$ and an element $\alpha \in \mathbb{F}$ we write $\langle \mathbf{g}, \boldsymbol{\alpha} \rangle$ for the group element $\prod_i g_i^{\alpha_i}$ (as already done in Section 4.3). We also write $\mathbf{g} \cdot \mathbf{h}$ for the vector $(g_1 \cdot h_1, \dots, g_n \cdot h_n)$, \mathbf{g}^α for the vector $(g_1^{\alpha_1}, \dots, g_n^{\alpha_n})$, \mathbf{g}^α for the vector $(g_1^\alpha, \dots, g_n^\alpha)$ and $\hat{e}(\mathbf{g}, \mathbf{h})$ for $(e(g_1, h), \dots, e(g_n, h))$.

Definition 7.1 (DLIN assumption [6]). *Let $g_1, g_2, g_3 \leftarrow^s \mathbb{G}$ and $r, s, z \leftarrow^s \mathbb{F}$, the Decisional Linear Assumption states that the following two distributions are computationally indistinguishable:*

$$(\mathbb{G}, g_1, g_2, g_3, g_1^r, g_2^s, g_3^{r+s}) \quad \text{and} \quad (\mathbb{G}, g_1, g_2, g_3, g_1^r, g_2^s, g_3^z).$$

Trapdoor hiding commitment. We use the generalized Pedersen commitment [53]. Algorithm `ESetup` outputs $\vartheta = \mathbf{g}^{\mathbf{r}}$ for random $\mathbf{r} \leftarrow_{\$} \mathbb{F}^{\mu+1}$, and sets the trapdoor to $\tau := \mathbf{r}$. The commitment to an element $m \in \mathbb{F}^{\mu}$ using randomness $\mathbf{r} \leftarrow \mathbb{F}$ is computed as $c = \text{Commit}(\vartheta, m; r) := \langle \vartheta, (m, r) \rangle$.

Note that Pedersen commitment is linearly homomorphic: For $com_1 = \text{Commit}(m_1; r_1)$ and $com_2 = \text{Commit}(m_2; r_2)$ it holds that

$$com_1 \cdot com_2 = \langle \vartheta, (m_1, r_1) + (m_2, r_2) \rangle = \text{Commit}(m_1 + m_2; r_1 + r_2).$$

Moreover, for all constants $c \in \mathbb{F}$ we have that $com_1^c = \langle \vartheta, (m_1, r) \rangle^c = \text{Commit}(c \cdot m_1; c \cdot r_1)$. Algorithm `Setup` obviously samples an element from $\mathbb{G}^{\mu+1}$; algorithm `ECommit` picks $r' \leftarrow_{\$} \mathbb{F}$ and returns $com = g^r$; algorithm `Equiv` upon input a message $m \in \mathbb{F}^{\mu}$, trapdoor $\tau = \mathbf{r}$, and randomness r' , returns $r := r' - \sum_i r_i m_i$.

Hybrid commitment. We use the homomorphic commitment of [37], that is secure under the DLIN assumption. The trapdoor hiding property is derived applying the general transformation of Damgaard and Nielsen in [18].

Let $\psi : \mathbb{F}^2 \rightarrow \mathbb{F}^3$ be such that $\psi(r_1, r_2) = (r_1, r_2, r_1 + r_2)$. The algorithms depend on global parameters (\mathbb{G}, \mathbf{g}) , where $\mathbf{g} \leftarrow_{\$} \mathbb{G}^3$. Algorithm `Setup` chooses $\bar{\mathbf{r}} \leftarrow_{\$} \mathbb{F}^3$, and sets the verification key as $\vartheta := \mathbf{g}^{\bar{\mathbf{r}}}$; note that oblivious sampling requires that the group \mathbb{G} itself supports oblivious sampling. Algorithm `ESetup` chooses $\bar{\mathbf{r}} \leftarrow_{\$} \mathbb{F}^2$, and sets the verification and trapdoor keys as $\vartheta := \mathbf{g}^{\psi(\bar{\mathbf{r}})}$ and $\tau := \bar{\mathbf{r}}$. The hybridness property follows by the DLIN assumption.

To commit to a value $m \in \mathbb{F}^{\mu}$, algorithm `Commit` takes the verification key $\vartheta = \mathbf{h}$, chooses randomness $\mathbf{r}_1, \dots, \mathbf{r}_{\mu} \leftarrow_{\$} \mathbb{F}^2$, and outputs:

$$com := (\mathbf{h}^{m_1} \cdot \mathbf{g}^{\psi(\mathbf{r}_1)}, \dots, \mathbf{h}^{m_{\mu}} \cdot \mathbf{g}^{\psi(\mathbf{r}_{\mu})}).$$

Algorithm `ECommit` takes the verification key $\vartheta = (\mathbf{g}^{\psi(\bar{\mathbf{r}}_1)}, \dots, \mathbf{g}^{\psi(\bar{\mathbf{r}}_{\mu})})$ and trapdoor key $\tau = \bar{\mathbf{r}}$, chooses randomness $\mathbf{r}'_1, \dots, \mathbf{r}'_{\mu} \leftarrow_{\$} \mathbb{F}^2$, and outputs:

$$com := (\mathbf{g}^{\psi(\mathbf{r}'_1)}, \dots, \mathbf{g}^{\psi(\mathbf{r}'_{\mu})}) \text{ and } \mathbf{r}' := (\mathbf{r}'_1, \dots, \mathbf{r}'_{\mu}).$$

To equivocate algorithm `Equiv` takes as input \mathbf{r}' and a trapdoor key $\tau = \bar{\mathbf{r}}$, and returns $\mathbf{r}_i := \mathbf{r}'_i - m_i \cdot \bar{\mathbf{r}}$ for each $i \in [\mu]$. We note that for any message m , any i , and any verification key ϑ , the distribution of \mathbf{r}_i is a bijective function of \mathbf{r}'_i , therefore \mathbf{r}_i is uniformly distributed on \mathbb{F}^2 as desired.

In Section 4.2 we require the extra property that given the randomness \mathbf{r}'_1 and \mathbf{r}'_2 for two equivocal commitments com_1 and com_2 (under verification key ϑ , with trapdoor information $\tau = \bar{\mathbf{r}}$), the randomness $\mathbf{r}'_c := \alpha \cdot \mathbf{r}'_a + \mathbf{r}'_b$ can equivocate the equivocal commitment $com_3 := com_1^{\alpha} \cdot com_2$. This is indeed the case, as

$$\begin{aligned} com_3 &= \left(\mathbf{g}^{\psi(\mathbf{r}'_{a,1})}, \dots, \mathbf{g}^{\psi(\mathbf{r}'_{a,\mu})} \right)^{\alpha} \cdot \left(\mathbf{g}^{\psi(\mathbf{r}'_{b,1})}, \dots, \mathbf{g}^{\psi(\mathbf{r}'_{b,\mu})} \right) = \\ &= \left(\mathbf{g}^{\alpha \cdot \psi(\mathbf{r}'_{a,1}) + \psi(\mathbf{r}'_{b,1})}, \dots, \mathbf{g}^{\alpha \cdot \psi(\mathbf{r}'_{a,\mu}) + \psi(\mathbf{r}'_{b,\mu})} \right) = \\ &= \left(\mathbf{g}^{\psi(\alpha \cdot \mathbf{r}'_{a,1} + \mathbf{r}'_{b,1})}, \dots, \mathbf{g}^{\psi(\alpha \cdot \mathbf{r}'_{a,\mu} + \mathbf{r}'_{b,\mu})} \right) = \\ &= \left(\mathbf{g}^{\psi(\mathbf{r}'_{c,1})}, \dots, \mathbf{g}^{\psi(\mathbf{r}'_{c,\mu})} \right) \end{aligned}$$

In Section 4.3 we require the following extra properties: (i) $\mathbb{H} = \{\mathbf{g}^{\psi(\mathbf{r})} : \mathbf{r} \in \mathbb{F}^2\}$ forms a proper subgroup of $\mathbb{G} = \{\mathbf{g}^{\mathbf{r}} : \mathbf{r} \in \mathbb{F}^3\}$; (ii) given $\vartheta_1 \in \mathbb{G} \setminus \mathbb{H}$ and $\vartheta_2 \in \mathbb{H}$, the product $\vartheta_1 \cdot \vartheta_2$ (defined as the element-wise multiplication) is in $\mathbb{G} \setminus \mathbb{H}$; (iii) $\mathbf{g}_1^{\psi(\mathbf{r})}$ is uniformly distributed in \mathbb{H} , for $\mathbf{r} \leftarrow_{\$} \mathbb{F}^2$. The above properties hold by simple inspection.

NIWI arguments. We instantiate the NIWI argument system using the powerful Groth-Sahai proofs system [37]. Using the so-called simulated common reference string key generator, the system is a statistical NIWI argument system.¹⁰ As pointed out in [8], the `Init` algorithm admits oblivious sampling. Following [37], the number of group elements for the common reference string is 12 group elements. The relation \mathfrak{R} in Section 4.1 can be expressed using $3\mu + 1$ equations with $3\mu + 1$ variables; the relation \mathfrak{R} in Section 5.2 can be expressed using $(3\mu + 1)\log \kappa + 1$ equations with $4\mu \log \kappa + 1$ variables. According to Libert [45] we have:

- A NIWI argument for the scheme \mathcal{SS}_1 requires $O(\mu)$ group elements, while the number of pairings operations for verification is $\Omega(\mu^2)$; the exact number of group elements is comparable with the OR-construction of [11]. The difference is that [11] employs the strong one-time signature of [35], which requires 2 group elements for the public key; instead our PRG (based on DLIN) uses 3 group elements. According to [45], we estimate that the NIWI argument of scheme \mathcal{SS}_1 requires approximately $24\mu + 90$ group elements.
- A NIWI argument for the scheme \mathcal{SS}_2 requires $15\mu + 5$ group elements, while the number of pairings operations for verification is $27\mu^2 + 36\mu + 12$.
- A NIWI argument for the scheme \mathcal{SS}_3 requires $(12 \log \kappa + 4)\mu + 4\mu + 5$ group elements, while the number of pairings operations for verification is $\Omega((\mu \log \kappa)^2)$.

NIZK arguments. We use the tSE-NIZK based on DLIN from Dodis [22]. Their instantiation uses a multi-message version of the Cramer-Shoup encryption scheme [14], and the Groth-Sahai proof system. One can check that Cramer-Shoup admits oblivious sampling of the public key.

Roughly, the argument system encrypts the witness and proves that the encrypted value is a valid witness for the relation. For the scheme \mathcal{SS}_2 presented in Section 4.3 the witness space for the relation \mathfrak{R}_2 (cf. Eq. (9)) is \mathbb{F}^2 ; however the Cramer-Shoup encryption scheme has message space \mathbb{G}^n for some n .

In order to solve the above mismatch, we instantiate the tSE-NIZK with the following relation:

$$\mathfrak{R}'_2 := \{(\vartheta, \vartheta_1, \vartheta_2); (g^{\alpha_1}, g^{\alpha_2}) \mid \hat{e}(\vartheta, g) = \hat{e}(\vartheta_1, g^{\alpha_1}) \cdot \hat{e}(\vartheta_2, g^{\alpha_2})\} \subseteq (\mathbb{G}^3)^3 \times \mathbb{G}^2.$$

Following [22], the size of a tSE-NIZK argument is of 124 group elements and 6 elements in \mathbb{F} , while the common reference string size is of 18 group elements.

Instantiating the scheme in the BRM. We use Fischlin’s transformation [32] for the NIZK with online extraction. To obtain reasonable security parameters the size of a proof needs to be $\Omega(\log \kappa)$ times the prover’s transcript in the Sigma protocol. Hence, by using the generalized Pedersen commitment scheme and the corresponding Sigma-protocol for proving knowledge of an opening, we obtain that the size of a signature is $\log \kappa$ group elements and $(\mu + 1)\log \kappa$ elements in \mathbb{F} .

For the homomorphic linear authenticator we use the scheme implicitly defined by Shacham and Waters in [57], based on BLS signature of Boneh *et al.* [7].

8 Extension to Noisy Leakage

Our definitions of Section 3 are in the bounded leakage model, where the adversary is allowed to query a leakage oracle with arbitrary polynomial-time functions, provided that the total amount

¹⁰Strictly speaking, to instantiate the NIWI argument using [37] one would have to slightly adapt Definition 2.3 in order to account for the fact that the witness indistinguishability property only holds when the CRS is generated using an alternative initialization algorithm `Init` (producing a computationally indistinguishable distribution); this is easily taken into account in the analysis by defining a further hybrid in the proofs (see Theorem 4.1, Theorem 4.2, and Theorem 5.1), where `Init` is replaced by `Init`.

of leakage is smaller than a global parameter ℓ . In this section, we explain how to extend our definitions and some of our results to a more general setting (that strictly implies the case of bounded leakage) where the leakage functions can have a much longer output, as long as the entropy of the secret key does not decrease by too much.

Following [21], we define a notion of a function being ℓ -leaky.

Definition 8.1. *A (possibly randomized) function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is ℓ -leaky, if for all $\kappa \in \mathbb{N}$ we have that $\mathbb{H}_\infty(U_\kappa | f(U_\kappa)) \geq \kappa - \ell$, where U_κ is the uniform distribution over $\{0, 1\}^\kappa$.*

As shown by [21][Lemma L.3] the above notion composes nicely, meaning that if two functions are respectively ℓ_1 - and ℓ_2 -leaky, their concatenation is $(\ell_1 + \ell_2)$ -leaky.

We explain how to adapt Definition 3.1 to the setting of noisy-leakage. Essentially we allow the adversary to query the leakage oracle on arbitrary polynomial-time functions, provided that each function is ℓ_i -leaky and $\sum_i \ell_i \leq \ell$, where ℓ is the leakage parameter. We stress that one-more unforgeability still requires that an adversary should produce a number of forgeries strictly larger than the ones he could have leaked (up-to the slack factor γ), however the number of leaked signatures now might depend on the amount of information that a signature actually carries. We distinguish two cases, depending on whether the signature algorithm (statistically) reveals partial information on the secret key or not. In the first case, the parameter n in the winning condition is related to the leakage parameter ℓ , since a forgery is de-facto a leaky function of the secret; in the second case we need to be more pessimistic, and let the parameter n be related to the actual leakage λ performed by the adversary.

Definition 8.2. *We say that a signature scheme $\mathcal{SS} = (\text{KGen}, \text{Sign}, \text{Verify})$ is (ℓ, γ) -fully-leakage one-more unforgeable w.r.t. noisy leakage if it satisfies Definition 3.1 with the following modification to the security experiment:*

5. *The experiment outputs 1 if and only if conditions (a)-(b) are met and moreover:*

(c) *Let Sign be \tilde{s} -leaky and $s := |\sigma|$ be the size of a signature, then:*

$$\begin{cases} n \geq \lfloor \lambda / (\gamma \cdot s) \rfloor + 1 & \text{if } \tilde{s} = 0 \\ n \geq \lfloor \ell / (\gamma \cdot \tilde{s}) \rfloor + 1 & \text{else} \end{cases} \quad (20)$$

where λ is the total length of the leakage.

(d) *Each function submitted to the leakage oracle is ℓ_i -leaky, $\sum_i \ell_i \leq \ell$, and $|\mathcal{Q}| \leq q_s$.*

The definitions for the case of perfect erasures, and in the BRM can be updated accordingly. Note that the above formulation offers a graceful degradation of security also for schemes where the secret key is potentially shorter than a signature, and signatures do not constrain the secret key. In fact, for such schemes, leaking a signature is allowed, but does not directly constitute a security breach as the adversary is required to produce at least one more forgery.

It is easy to see that the above formulation strictly implies our definitions in the bounded leakage model. In fact, an ℓ -bit output function is in particular ℓ -leaky.

The theorem below shows that our signature scheme \mathcal{SS}_3 (cf. Section 4) and our schemes \mathcal{SS}_4 and \mathcal{SS}_5 (cf. Section 6) are secure also w.r.t. noisy leakage. Unfortunately we are not able to prove a similar statement for the schemes \mathcal{SS}_1 and \mathcal{SS}_2 (cf. Section 5). The reason for this is that in the erasure case the simulator could be forced to leak too much information, as the number of leaky signature queries is not anymore upper-bounded by λ , but is instead an arbitrary polynomial. In such a case, we cannot upper bound the probability of the events Bad_i in the proof of Lemma 4.2. This issue does not apply to the non-erasure case, as the simulator can arbitrarily control the probability of the events Bad_i .¹¹

Theorem 8.1. *The statement of Theorem 5.1, Theorem 6.1, and Theorem 6.2 also hold w.r.t. noisy-leakage security, with the same leakage and efficiency parameters.*

¹¹We do not have a concrete attack on the schemes.

8.1 Proof Sketch of Theorem 8.1

We will rely on the following technical lemma (see [25, Lemma 2.2] and [21, Lemma L.3]).

Lemma 8.1. *Let X be a random variable over some domain \mathcal{X} . The following holds:*

1. *For any random variable $Y \in \mathcal{Y}$, and for any $\beta > 0$, we have that*

$$\mathbb{P} \left[\mathbb{H}_\infty(X|Y = y) \geq \tilde{\mathbb{H}}_\infty(X|Y) - \beta \right] \geq 1 - 2^{-\beta}$$

(where the probability is over the choice of y).

2. *For all sequences of ℓ_i -leaky functions $f_i : \{0, 1\}^* \rightarrow \{0, 1\}^*$, such that $\sum_i \ell_i \leq \ell$, we have that $\tilde{\mathbb{H}}_\infty(X|f_1(X), f_2(X), \dots) \geq \mathbb{H}_\infty(X) - \ell$.*

We start with the statement of Theorem 5.1. We need to show the following twist of Lemma 4.2, where the bound now depends on ℓ (the amount of min-entropy left in the secret key) instead of λ (the total length of the output of the leakage functions).

Lemma 8.2. *For any $\beta > 0$ we have that*

$$\mathbb{P} [\mathbb{H}_\infty(\Delta | \text{View}_3 = v) \geq |\Delta| - (1 + B)\mu \log |\mathbb{F}| - \ell - \beta] \geq 1 - 2^{-\beta} - 2^{-2e}. \quad (21)$$

Proof. Let $\text{View}'_3 := (f_i((\Delta, \mathbf{r}), st_i))_{i \in [q_s]} || (\Delta(m_i))_{i \in \mathcal{Z}}$ the total information leaked in \mathcal{H}_3 , where \mathcal{Z} , as before, is the set of bad queries that force the simulator to leak from Δ . For any adversary's view v let v' the “stripped-off” view where we consider only the total information leaked via the leakage oracle.

$$\mathbb{H}_\infty(\Delta | \text{View}_3 = v) = \mathbb{H}_\infty(\Delta | \text{View}'_3 = v')$$

We condition on the number of bad queries $Z = |\mathcal{Z}|$ being less than $2e$, then we apply Property (1) of Lemma 8.1:

$$\mathbb{P} \left[\mathbb{H}_\infty(\Delta | \text{View}'_3 = v') \geq \tilde{\mathbb{H}}_\infty(\Delta | \text{View}'_3) - \beta \mid Z \leq 2e \right] \geq 1 - 2^{-\beta}.$$

Now we can apply Property (2) of Lemma 8.1, using the fact that the function $f(\Delta) := (\Delta(m_i))_{i \in \mathcal{Z}}$ is $(Z\mu \log |\mathbb{F}|)$ -leaky:

$$\mathbb{P} \left[\mathbb{H}_\infty(\Delta | \text{View}'_3 = v') \geq |\Delta| - (2e\mu \log |\mathbb{F}| + \ell) - \beta \mid Z \leq 2e \right] \geq 1 - 2^{-\beta}.$$

Since $\mathbb{P}[Z \leq 2e] > 1 - 2^{-2e}$ by the Chernoff bound, we obtain the statement of the lemma. \square

To conclude the proof, one can verify that, similarly to Theorem 5.1, in order to obtain a contradiction we need that:

$$\frac{n\kappa}{q_s} = \Omega(\ell - \log \varepsilon(\kappa)). \quad (22)$$

For $\varepsilon(\kappa)$ noticeable, since $n \geq \lfloor \frac{\ell}{\gamma \cdot \kappa} \rfloor + 1$ (as the signing algorithm is κ -leaky), it suffices to set $\gamma = O(\frac{1}{q_s})$ in order for Eq. (22) to hold.

We turn to the proof of the noisy-variant of Theorem 6.1 and Theorem 6.2. It suffices to prove Theorem 6.1, as the other proof is essentially identical. The proof proceeds as before, except that now Eq. (17) follows by a direct application of Property (2) of Lemma 8.1 (and not by the chain rule of average conditional min-entropy). The signature algorithm is 0-leaky, so condition (c) in the definition of fully-leakage one-more unforgeability stays the same (see Definition 8.2). The proof follows.

9 Conclusions and Open Problems

We have shown new constructions of fully leakage resilient signatures in the bounded and noisy leakage model. Our schemes enjoy a graceful degradation of security in situations where the size of a single signature is independent of the size of the secret key. This notion is still meaningful for security, and moreover allows to obtain shorter signature tolerating a $1 - o(1)$ fraction of leakage on the secret key.

Our main schemes are efficient, in the standard model, and can be instantiated under fairly standard cryptographic assumptions. (We also built a scheme in the BRM, relying on a random oracle.)

We conclude by highlighting a few open problems left open by our work. As for the standard model constructions, it would be interesting to construct more efficient schemes, both in the erasure and non-erasure model, with optima slack parameter $\gamma = O(1)$. As for the scheme in the BRM, it is a challenging open question whether it is possible to remove the random oracle. Also it would be interesting to make the proof work without assuming online extractability for the NIZK (e.g., using the Fiat-Shamir transform [31]). This case requires a more careful analysis, dealing with rewinding an adversary that outputs n forgeries, and seems to require a more general version of the forking lemma [54].

Finally, it remains open to generalize our schemes to the setting of continuous and hard-to-invert leakage (possibly without relying on hardware assumptions, like the Floppy model of [2, 17]).

References

- [1] Joël Alwen, Yevgeniy Dodis, Moni Naor, Gil Segev, Shabsi Walfish, and Daniel Wichs. Public-key encryption in the bounded-retrieval model. In *EUROCRYPT*, pages 113–134, 2010.
- [2] Joël Alwen, Yevgeniy Dodis, and Daniel Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In *CRYPTO*, pages 36–54, 2009.
- [3] Giuseppe Ateniese, Seny Kamara, and Jonathan Katz. Proofs of storage from homomorphic identification protocols. In *ASIACRYPT*, pages 319–333, 2009.
- [4] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
- [5] Dan Boneh and Xavier Boyen. Secure identity based encryption without random oracles. In *CRYPTO*, pages 443–459, 2004.
- [6] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *CRYPTO*, pages 41–55, 2004.
- [7] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *ASIACRYPT*, pages 514–532, 2001.
- [8] Elette Boyle, Gil Segev, and Daniel Wichs. Fully leakage-resilient signatures. In *EUROCRYPT*, pages 89–108, 2011.
- [9] Zvika Brakerski and Shafi Goldwasser. Circular and leakage resilient public-key encryption under subgroup indistinguishability - (or: Quadratic residuosity strikes back). In *CRYPTO*, pages 1–20, 2010.
- [10] Zvika Brakerski, Yael Tauman Kalai, Jonathan Katz, and Vinod Vaikuntanathan. Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In *FOCS*, pages 501–510, 2010.
- [11] Jan Camenisch, Nishanth Chandran, and Victor Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *EUROCRYPT*, pages 351–368, 2009.

- [12] Ran Canetti, Dror Eiger, Shafi Goldwasser, and Dah-Yoh Lim. How to protect yourself without perfect shredding. In *ICALP*, pages 511–523, 2008.
- [13] Dario Catalano and Ivan Visconti. Hybrid trapdoor commitments and their applications. In *ICALP*, pages 298–310, 2005.
- [14] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO*, pages 13–25, 1998.
- [15] Giovanni Di Crescenzo, Richard J. Lipton, and Shabsi Walfish. Perfectly secure password protocols in the bounded retrieval model. In *TCC*, pages 225–244, 2006.
- [16] Özgür Dagdelen and Daniele Venturi. A second look at Fischlin’s transformation. In *AFRICACRYPT*, pages 356–376, 2014.
- [17] Ivan Damgaard, Sebastian Faust, Pratyay Mukherjee, and Daniele Venturi. Bounded tamper resilience: How to go beyond the algebraic barrier. In *ASIACRYPT*, pages 140–160, 2013.
- [18] Ivan Damgaard and Jesper Buus Nielsen. Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In *CRYPTO*, pages 581–596, 2002.
- [19] Francesco Davì, Stefan Dziembowski, and Daniele Venturi. Leakage-resilient storage. In *SCN*, pages 121–137, 2010.
- [20] Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Public-key encryption schemes with auxiliary inputs. In *TCC*, pages 361–381, 2010.
- [21] Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs. Cryptography against continuous memory attacks. In *FOCS*, pages 511–520, 2010.
- [22] 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.
- [23] Yevgeniy Dodis, Yael Tauman Kalai, and Shachar Lovett. On cryptography with auxiliary input. In *STOC*, pages 621–630, 2009.
- [24] Yevgeniy Dodis, Allison B. Lewko, Brent Waters, and Daniel Wichs. Storing secrets on continually leaky devices. In *FOCS*, pages 688–697, 2011.
- [25] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- [26] Stefan Dziembowski. Intrusion-resilience via the bounded-storage model. In *TCC*, pages 207–224, 2006.
- [27] Stefan Dziembowski. On forward-secure storage. In *CRYPTO*, pages 251–270, 2006.
- [28] Sebastian Faust, Carmit Hazay, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Angela Zottarel. Signature schemes secure against hard-to-invert leakage. In *ASIACRYPT*, pages 98–115, 2012.
- [29] Sebastian Faust, Eike Kiltz, Krzysztof Pietrzak, and Guy N. Rothblum. Leakage-resilient signatures. In *TCC*, pages 343–360, 2010.
- [30] Sebastian Faust, Markulf Kohlweiss, Giorgia Azzurra Marson, and Daniele Venturi. On the non-malleability of the fiat-shamir transform. In *INDOCRYPT*, pages 60–79, 2012.
- [31] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
- [32] Marc Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In *CRYPTO*, pages 152–168, 2005.
- [33] J. Longo Galea, D. Martin, E. Oswald, D. Page, M. Stam, , and M. Tunstall. Simulatable leakage: Analysis, pitfalls, and new constructions. Cryptology ePrint Archive, Report 2014/357, 2014. <http://eprint.iacr.org/>.
- [34] Sanjam Garg, Abhishek Jain, and Amit Sahai. Leakage-resilient zero knowledge. In *CRYPTO*, pages 297–315, 2011.

- [35] Jens Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *ASIACRYPT*, pages 444–459, 2006.
- [36] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In *EUROCRYPT*, pages 339–358, 2006.
- [37] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In *EUROCRYPT*, pages 415–432, 2008.
- [38] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest we remember: Cold boot attacks on encryption keys. In *USENIX Security Symposium*, pages 45–60, 2008.
- [39] Shai Halevi and Huijia Lin. After-the-fact leakage in public-key encryption. In *TCC*, pages 107–124, 2011.
- [40] Dennis Hofheinz and Eike Kiltz. Programmable hash functions and their applications. In *CRYPTO*, pages 21–38, 2008.
- [41] Yael Tauman Kalai, Bhavana Kanukurthi, and Amit Sahai. Cryptography with tamperable and leaky memory. In *CRYPTO*, pages 373–390, 2011.
- [42] Jonathan Katz and Vinod Vaikuntanathan. Signature schemes with bounded leakage resilience. In *ASIACRYPT*, pages 703–720, 2009.
- [43] Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *CRYPTO*, pages 104–113, 1996.
- [44] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *CRYPTO*, pages 388–397, 1999.
- [45] Benoît Libert, Thomas Peters, Marc Joye, and Moti Yung. Non-malleability from malleability: Simulation-sound quasi-adaptive NIZK proofs and cca2-secure encryption from homomorphic signatures. In *EUROCRYPT*, pages 514–532, 2014.
- [46] Vadim Lyubashevsky, Adriana Palacio, and Gil Segev. Public-key cryptographic primitives provably as secure as subset sum. In *TCC*, pages 382–400, 2010.
- [47] Tal Malkin, Isamu Teranishi, Yevgeniy Vahlis, and Moti Yung. Signatures resilient to continual leakage on memory and computation. In *TCC*, pages 89–106, 2011.
- [48] Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In *TCC*, pages 278–296, 2004.
- [49] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. In *CRYPTO*, pages 18–35, 2009.
- [50] Jesper Buus Nielsen, Daniele Venturi, and Angela Zottarel. On the connection between leakage tolerance and adaptive security. In *Public Key Cryptography*, pages 497–515, 2013.
- [51] Jesper Buus Nielsen, Daniele Venturi, and Angela Zottarel. Leakage-resilient signatures with graceful degradation. In *Public-Key Cryptography*, pages 362–379, 2014.
- [52] Rafael Pass. On deniability in the common reference string and random oracle model. In *CRYPTO*, pages 316–337, 2003.
- [53] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *CRYPTO*, pages 129–140, 1991.
- [54] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *J. Cryptology*, 13(3):361–396, 2000.
- [55] Jean-Jacques Quisquater and David Samyde. Electromagnetic analysis (EMA): Measures and counter-measures for smart cards. In *E-smart*, pages 200–210, 2001.
- [56] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *CRYPTO*, pages 566–598, 2001.
- [57] Hovav Shacham and Brent Waters. Compact proofs of retrievability. In *ASIACRYPT*, pages 90–107, 2008.

- [58] Francois-Xavier Standaert, Tal Malkin, and Moti Yung. A unified framework for the analysis of side-channel key recovery attacks. In *EUROCRYPT*, pages 443–461, 2009.
- [59] Francois-Xavier Standaert, Olivier Pereira, and Yu Yu. Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In *CRYPTO*, pages 335–352, 2013.
- [60] Brent Waters. Efficient identity-based encryption without random oracles. In *EUROCRYPT*, pages 114–127, 2005.