Efficient Fully-Leakage Resilient One-More Signature Schemes

Antonio Faonio

IMDEA Software Institute, Madrid, Spain

In a recent paper Faonio, Nielsen and Venturi (ICALP 2015) gave new constructions of leakage-resilient signature schemes. The signature schemes proposed remain unforgeable against an adversary leaking arbitrary information on the entire state of the signer, including the random coins of the signing algorithm. The main feature of their signature schemes is that they offer a graceful degradation of security in situations where standard existential unforgeability is impossible. The notion, put forward by Nielsen, Venturi, and Zottarel (PKC 2014), define a *slack parameter* γ which, roughly speaking, describes how gracefully the security degrades. Unfortunately, the standard model signature scheme of Faonio, Nielsen and Venturi has slack parameter that depends on number of signature queried by the adversary.

In this paper we show two new constructions in the standard model where the above limitation is avoided. Specifically, the first scheme we propose achieves slack parameter $O(1/\lambda)$ where λ is the security parameter and it is based on standard number theoretic assumptions, the second scheme achieves optimal slack parameter (i.e. $\gamma = 1$) and it is based on knowledge of the exponent assumptions. Our constructions are efficient and have leakage rate $1 - o(1)$, most notably our second construction has signature size of only 8 group elements which makes it the leakage-resilient signature scheme with the shortest signature size known to the best of our knowledge.

1 Introduction

In the last years a lot of effort has been put into constructing cryptographic primitives that remain secure even in case the adversary obtains partial information of the secrets used within the system. This effort is motivated by the existence of the so-called side-channel attacks (see, e.g. [29,30,20]) which can break provably secure cryptosystem exploiting physical characteristics of the crypto-devices where such scheme are implemented.

A common way to model leakage attacks, is to give to the adversary a leakage oracle, taking as input adaptively chosen functions f_i and returning $f_i(\alpha)$ where α is the current secret state of the cryptosystem under attack. Some restriction on the functions f_i must be impose, as otherwise there is no hope for security. In this paper we consider the *bounded leakage model* where we just assume that the total bit-length of the leakage obtained via the functions f_i is smaller than some a priori determined leakage bound ℓ. Leakage-resilient schemes in this model include public-key and identity-based encryption [32,3,2,9,11,25], signature schemes [27,3,9,11,8,31], and more (see, e.g., [21,7,4,14]).

Graceful degradation. For a signature scheme to remain existentially unforgeable in the bounded leakage model, it is necessary that the length of a signature is larger than the length of the secret key, as otherwise an adversary could simply leak a forgery. The first consequence is that signatures are very long, as the goal is to tolerate large leakage, which is impractical, the 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.

Recently Nielsen, Venturi and Zottarel [33] addressed this issue introducing a "graceful degradation" property, which essentially requires that an adversary should not be able to produce more forgeries than what he could have leaked via leakage queries. In particular, to break unforgeability, an adversary has to produce n forgeries where $n \approx \ell/(\gamma \cdot s) + 1$ for signatures of size s, a total of ℓ bits of leakage, and a "slack parameter" $\gamma \in (0, 1]$ measuring how close to optimal security a scheme is. This enables to design signature schemes where the size of the secret key do not depend on the signature size, leading to shorter signatures, which still allows for interesting applications (e.g., to leaky identification [33]). Subsequently, Faonio, Nielsen and Venturi [15] (journal version in [16]), extended the model to the fully-leakage resilient setting, where the adversary can leak arbitrary information of the entire secret state, including the random coins of the signing algorithm.

Interestingly, while in the (not-fully) leakage-resilient regime the authors of [33] showed a signature scheme with almost-optimal graceful degradation (i.e. γ is a constant in $(0, 1]$), in the fully-leakage-resilient regime the best signature scheme in the standard model known has graceful degradation $O(1/q)$ where q is the number of signature oracle queries performed by the adversary. While the latter result still allows for some meaningful applications, in practice, the leakage security of the scheme is hard to estimate as it degrades as function of the number of signatures which in principle could be really big.

Our contributions. In this paper we solve the above problem by constructing fully leakage-resilient signatures in the bounded leakage model with graceful degradation that does not depend on the number of signatures issued. In particular, we construct two signature schemes, the former has slack parameter $O(1/\lambda)$ and it is based on standard number theoretic assumptions, the latter, most notably, has optimal graceful degradation (i.e. $\gamma = 1$) and it is based on knowledge of the exponent assumptions (see, e.g. [10,5,1,23]). The first signature scheme is based on a recent paper of Fujisaki [18], while the second construction is based on a quasi-adaptive NIZK for linear space with adaptive *weak* knowledge soundness. In particular, as minor contribution of independent interest, we show how to modify the elegant construction of Kiltz and Wee [28] to get an efficient quasi-adaptive argument system with weak knowledge soundness for linear-space relationship.

A Technical Overview. We start recalling the scheme of [33], for future reference we call the scheme NVZ14. Roughly speaking, the secret key material for NVZ14 is a polynomial $\boldsymbol{\delta}$ in $\mathbb{Z}_p[X]$

of degree d and a signature for $m \in \mathbb{Z}_p$ is composed by a commitment C^* to the evaluation of the polynomial $\boldsymbol{\delta}$ on the point m together with a simulation-extractable NIZK that the commitment, indeed, commit to such evaluation. More in details, the polynomial $\boldsymbol{\delta}$ is published in the verification key using an homomorphic commitment scheme (for example, the classical Pedersen's commitment scheme [35]) so that from the verification key we can derive the homomorphic evaluation of the polynomial in a committed form C_m and the NIZK needs only to prove that the commitments C_m and C^* contain the same value. The key idea is that from an adversary we can extract n evaluations of the polynomial $\boldsymbol{\delta}$ however, because of the bound on the leakage performed, at most $\ell/\log p \approx n-1$ could be possibly be uniquely defined. The authors show that this property is sufficient to find two different openings for a commitment and therefore to break the binding property of the commitment scheme.

The main insight of [15] is that, in the fully-leakage resilient setting, neither simulation-extractable NIZK nor a standard commitment scheme for the generation of signatures are sufficient. In fact two contrasting requirements are necessary: on one hand, to guarantee that the queried signatures information-theoretically hide the polynomial $\boldsymbol{\delta}$, even in presence of leakage from the randomness, the commitment scheme needs to be statistically hiding and the NIZK proof needs to be statistically zero-knowledge. On the other hand, to extract the n evaluations of $\boldsymbol{\delta}$ we need that either the commitment scheme or the NIZK proof are perfectly binding. Unfortunately, neither statistically NIZK nor perfectly hiding commitment can be perfectly binding (under non-falsiable assumptions). To solve this problem in [15] is it shown a construction of a commitment scheme that with probability $1/q$ is perfectly binding and with probability $1 - 1/q$ is perfectly hiding. In this way, almost all the signatures queried by the adversary will be perfectly hiding while over the $n \approx O(q \cdot \ell)$ forged signature (so that $\gamma = O(1/q)$) with overwhelming probablity strictly more than $(\ell/\log p) + 1$ signatures are perfectly binding. The security follows by a noticing that a winning adversary gets in input ℓ bit of information about $\boldsymbol{\delta}$ and outputs strictly more than ℓ bits of information about $\boldsymbol{\delta}$, this adversary cannot exists as otherwise a basic information theoretic principle would be violated.

We describe our two new signature schemes. In the first construction we substitute the commitment scheme of [15] with an All-But-Many Encryption (ABM-Enc) scheme. Roughly speaking, an ABM-Enc is an encryption scheme where all the ciphertext created by the adversary can be successfully decrypted (knowing the secret key) while, with the knowledge of a special trapdoor, it is possible to create an unbounded number of fake ciphertexts that are equivocable. The proof of security is quite straight-forward (actually even easier then in [15]): with the knowledge of the trapdoor all the signatures are equivocated and with the knowledge of the secret key of the ABM-Enc all the forged signature are extracted. Fujisaki [18], building over a paper of Hofheinz [26], showed two constructions of ABM-Enc. The first construction achieves constant overhead (the ratio between ciphertext size and message size) and it is based on the decision Composite Residuosity (DCR) assumption while the latter is based on DDH and achieves $\lambda/\log \lambda$ overhead. At first sight, by plugging the constant overhead ABM-Enc of Fujisaki in our signature scheme we would get a fully-leakage resilient signature with almost-optimal slack paramater, the problem is that efficient NIZK [24] and the the construction over DCR groups do not quite match. In particular, a Groth-Sahai proof for the needed statement would commit the witness bit-by-bit so that the total size of the signature is $O(\lambda^2)$ groups elements. Since each forged signature carries only $\log p$ bits of information this, unfortunately, implies that the slack parameter is $1/\mathtt{poly}(\lambda)$. On the other hand, the second construction of ABM-Enc based on DDH fits better with the Groth-Sahai NIWI mechanism, as to prove the necessary statement we only a costant number pairing-product equations in the size of the ciphertext.

The second construction is inspired by the following observation: if we used a zk-SNARK [23,22,34] instead of Groth-Sahai then the construction sketched above would have signature size $O(\lambda)$ and therefore almost-optimal slack parameter. However, at second thought, employing zk-SNARKs is definitely an over-killing, as what we need is only the ability of simultaneously information-theoretically hide and extract the commitments. Instead, we consider the commitment scheme of

Scheme	Fully	No Erasure	KGen	G. D.	Assumption	Efficiency	
						leak	signature size
NVZ14	✗	-	-	$O(1)$	DLIN	$\frac{1}{2} - o(1)$	$O(1)$
$\mathsf{FNV15}_1$	✓	✗	✓	$O(1)$	DLIN	$1 - \epsilon$	$O(\epsilon^{-1})$
$\mathsf{FNV15}_2$	✓	✓	✓	$O(1/q)$	DLIN	$1 - \epsilon$	$O(\epsilon^{-1} \cdot \log \lambda)$
$\mathsf{FNV15}_3$	✓	✓	✓	$O(1)$	BDH*	$1 - \epsilon$	$O(\epsilon^{-1} \cdot \log \lambda)$
$\mathcal{SS}_1$	✓	✓	✓	$O(1/\lambda)$	SXDH	$1 - \epsilon$	$O(\epsilon^{-1} \cdot \lambda)$
$\mathcal{SS}_2$	✓	✓	✗	1	KerLin_2+q-KE*	$1 - \epsilon$	8λ

Table 1: Comparison of known efficient leakage-resilient one-more signature schemes in the bounded leakage model. The * symbol means the scheme is in the random oracle model; G.D. stands for graceful degradation. The signature size is computed in number of group elements. The value ϵ is parameter set at initialization phase and it can be any inverse polynomial of the security parameter. DLIN stands for the decision linear assumption, BDH stands for the bilinear Diffie-Hellman assumption, SXDH stands for the external decision diffie-hellman assumption.

Abe and Fehr [1] based on the knowledge of the exponent assumption (KEA3) of Bellare and Palacio [5] (see also Damgaard [10]). Then we notice that, for this kind of commitments, proving that two commitments commit to the same message reduces to prove that a vector in $\mathbb{G}^2$ lies in a specific subspace of $\mathbb{G}^2$. The second tuning is, therefore, to consider the quasi-adaptive zero-knowledge (QA-NIZK) proof for linear subspace of Kiltz and Wee [28]. The last technical point is that, to assure that the extracted value is indeed an evaluation of the polynomial $\boldsymbol{\delta}$ we additionally need to extract the witness of the QA-NIZK proof. To do so, we resort again to the knowledge of the exponent assumption constructing a knowledge sound variation of the scheme of [28].

Comparison. We compare our signature schemes with the signatures of [33] and [16] (see Table 1). Four different signature schemes are presented in [16], we select the three most interesting[1] and we denote them with $\mathsf{FNV15}_1$,$\mathsf{FNV15}_2$ and $\mathsf{FNV15}_3$.

The third column in Table 1 (namely, "No Erasure") reefers to a weak model of fully leakage resilient signature considered in [15]. Specifically, the scheme $\mathsf{FNV15}_1$ is proved secure under the assumption that the cryptographic device can perfectly erase the random coins used in previous invocations. We call $\mathcal{SS}_1$ the signature scheme based on ABM-Enc schemes and $\mathcal{SS}_2$ the scheme based on knowledge of the exponent assumption. From an efficiency point of view we notice that $\mathcal{SS}_1$ is less efficient than $\mathsf{FNV15}_2$ but achieves asymptotically better graceful degradation. On the other hand, $\mathcal{SS}_1$ is both less efficient and with worse graceful degradation respect to $\mathsf{FNV15}_1$ and $\mathsf{FNV15}_3$, however, $\mathsf{FNV15}_1$ needs perfect erasure of the randomness and $\mathsf{FNV15}_3$ is only proved secure in the random oracle model. The signature scheme $\mathcal{SS}_2$ is proved secure in a fully-leakage model where the key generation phase is leak free. We consider this a reasonable assumption, in fact, in almost all practical scenarios we could safely assume that the cryptographic devices are initialized in a safe environment before being used *in the wild*. The technical reason behind this limitation is that commitment schemes based on the knowledge of the exponent assumptions do not admits oblivious sampling of the parameters. The scheme $\mathcal{SS}_2$ achieves optimal graceful degradation, moreover the signature size is independent of the parameter ϵ and, notably, more compact (both asymptotically and practically) even than the signature scheme $\mathsf{FNV15}_3$ in the random oracle model.

2 Notations and Preliminaries

Throughout the paper we let λ denote the security parameter. We say that a function f is negligible in the security parameter λ, and we write $f \in \mathtt{negl}(\lambda)$, if it vanishes asymptotically faster than

[1] As the forth scheme is a variation of $\mathsf{FNV15}_1$ and it achieves worse efficiency parameters

Experiment $\mathbf{Exp}^{\mathtt{q-KE^*}}_{\mathsf{A},\mathcal{S},\mathsf{Ext},\mathsf{Setup}_{BG}}(1^\lambda)$:

1. Let $\mathtt{prm} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, G_1, G_2, G_T, e) \leftarrow \mathsf{Setup}_{BG}(1^\lambda)$;
2. Sample $[\boldsymbol{g}]_1 \leftarrow\$\ \mathbb{G}_1^q$ and $\alpha \leftarrow\$\ \mathbb{Z}_p$ and set $[\mathcal{M}]_1 \leftarrow (1, \alpha)^T \cdot [\boldsymbol{g}]_1^T \in \mathbb{G}_1^{2,q}$, sample $r \leftarrow\$\ \{0,1\}^\lambda$;
3. Let $\boldsymbol{Y} \leftarrow \mathsf{A}([\mathcal{M}]_1, [\alpha]_2, \mathcal{S}([\mathcal{M}]_1, [\alpha]_2; r))$ and $\boldsymbol{z} \leftarrow \mathsf{Ext}([\mathcal{M}]_1, [\alpha]_2, \mathcal{S}([\mathcal{M}], [\alpha]_2; r))$;
4. Output 1 iff $\boldsymbol{Y} \in \mathit{Span}([1, \alpha]_1)$ and $\boldsymbol{Y} \neq [\mathcal{M}]_1 \cdot \boldsymbol{z}$.

Fig. 1: The q-KE* assumption.

the inverse of any polynomial. We use the classic notion of probabilistic polynomial time (PPT) algorithms. We write $x \leftarrow\$\ \mathcal{D}$ (resp. $x \leftarrow\$\ \mathsf{A}(y)$) to denote that x is chosen at random from the distribution $\mathcal{D}$ (resp. an PPT algorithm A run on input y), and we write $x \leftarrow \mathsf{A}(y; r)$ to denote that we assign to x the output of A. For two ensembles $\mathcal{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$, we write $\mathcal{X} \equiv \mathcal{Y}$ to denote that $\mathcal{X}$ and $\mathcal{Y}$ are identically distributed, and $\mathcal{X} \approx_s \mathcal{Y}$ (resp., $\mathcal{X} \approx_c \mathcal{Y}$) to denote that $\mathcal{X}$ and $\mathcal{Y}$ are statistically (resp., computationally) indistinguishable. Vectors and matrices are typeset in boldface.

Given an element $m \in \mathbb{Z}$ and a vector $\boldsymbol{v}$ of length d, we denote $\boldsymbol{v}(m) := \boldsymbol{v}^T \cdot (1, m^1, \ldots, m^{d-1})^T$, meaning the evaluation of the polynomial with coefficients $\boldsymbol{v}$ at point m. We consider also the natural extension of the notion to matrix, $\boldsymbol{V}(m) := \boldsymbol{V} \cdot (1, m^1, \ldots, m^{d-1})^T$. All the algorithms take as input (group) parameters $\mathtt{prm}$, for readability, whenever it is clear from the context we consider it implicit. Given two random variables X and Y, we define the average conditional min-entropy $\widetilde{\mathbb{H}}_\infty(X|Y) = \max_{\mathsf{P}} (-\log \Pr[\mathsf{P}(Y) = X])$, similarly we can generalize (see Alwen, Dodis and Wichs [3]) to *interactive* predictors that participate in some randomized experiment Y with the goal of guess X. In Appendix we state the basic properties about that we make us of. A (bilinear) group generator Setup_{BG} is an algorithm that upon input the security parameter 1^λ outputs the description $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, G_1, G_2, G_T, e)$ of three groups equipped with a (non-degenerate) bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$. We use additive notation for the group operation, and we denote group elements using the bracket notation introduced by Escala *et al.* in [13]. Namely, for a $y \in \mathbb{Z}_p$ we let $[y]_X$ be the element $y \cdot G_X \in \mathbb{G}_X$ for $X \in \{1, 2, T\}$. Given $[x]_1$ and $[y]_2$ we can write $[x \cdot y]_T$ as shorthand for $e([x]_1, [y]_2)$.

We recall the standard notion of of collision resistant hash function (CRH). A tuple of PPT $(\mathsf{Gen}_{CRH}, \mathsf{H})$ is a CRH where Gen_{CRH} upon security parameter 1^λ produce an hash key hk and the algorithm H upon input the hash key hk and a message in $\{0,1\}^*$ outputs a string in $\{0,1\}^\lambda$. Collision resistance states that for a randomly sampled hk, it is hard to find two different messages which are mapped to the same output, where the messages are computed as function of hk.

2.1 Number-Theoretic Assumptions

Knowledge of the Exponent Assumption. Consider the experiment in Fig 1 between an adversary A, a randomness sampler $\mathcal{S}$ and an extractor Ext and a bilinear group generator Setup_{BG}. The q-KE* assumption over bilinear groups states the following:

Definition 1. *Given a bilinear group generator Setup_{BG} and a value $q \in \mathbb{N}$, we say that the q-KE* assumption holds for Setup_{BG} if for any PPT A and any sampler $\mathcal{S}$ there exists a PPT Ext such that:*

$$\mathtt{Adv}^{q-KE^*}_{\mathsf{A},\mathsf{Ext},\mathsf{Setup}_{BG}}(\lambda) := \Pr\left[\mathbf{Exp}^{\mathtt{q-KE^*}}_{\mathsf{A},\mathsf{Ext},\mathsf{Setup}_{BG}}(1^\lambda)\right] \in \mathtt{negl}(\lambda).$$

Some remarks follows. The reason why we add the sampler $\mathcal{S}$ in the way we state the knowledge of the exponent assumption is because we deal with adversaries with oracle access (for example to the signature oracle or the leakage oracle). In this setting, as shown by Fiore and Nitulescu [17], we need to take particular care on how the adversary can interact with its oracles. In particular, as we

will show in the proof of security in Sec. 5, with the help of the sampler, we can reduce the queries of the adversary to be non-adaptive.

Second, notice that in bilinear groups the test $\boldsymbol{Y} \in Span([1, \alpha]_1)$ can be efficiently performed using the bilinear map $e(\boldsymbol{Y}_0, [\alpha]_2) = e(\boldsymbol{Y}_1, [1]_2)$. Also, we can naturally scale down the assumption to non bilinear groups, in this case, the adversary does not get $[\alpha]_2$. Given a (no-bilinear) group generator Setup_G the assumption for $q = 1$ is not stronger than the KEA [10] for non-uniform PT adversaries, while for $q = 3$ is not stronger than the KEA3 assumption [5] for non-uniform PT adversaries. For a bilinear group Setup_{BG}, and any polynomial q, the q-KE* assumption is not stronger than the q-PKE assumption of [23], indeed it is easy to show that if q-PKE holds than also q-KE* holds, however, the reverse implication is not known. The extractability assumptions for non-uniform adversaries consider an extractor that works for any auxiliary inputs. As shown in [6] this sometimes can be dangerous. Notice that in our assumption the only "auxiliary input" is generated by the random sampler $\mathcal{S}$ which does not take the secret material $\boldsymbol{g}, \alpha \in \mathbb{Z}_p$ on clear[2].

Kernel Diffie-Hellman Assumptions. Given parameter $\mathtt{prm}$, we call $\mathcal{D}_k$ a matrix distribution if it outputs a matrix in $\mathbb{Z}_p^{k+1,k}$ of full rank k in polynomial time.

Definition 2 (Escala *et al.* [13]). *Given a bilinear group generator* Setup_{BG}*, we say that the* $\mathcal{D}_k$*-Kernel Diffie-Hellman assumption (*$\mathcal{D}_k$-KerMDH*) holds for* Setup_{BG} *if for any PPT* A*:*

$$\mathtt{Adv}^{\mathcal{D}_k-\mathsf{KerMDH}}_{\mathsf{A},\mathsf{Setup}_{BG}}(\lambda) := \Pr\left[\boldsymbol{c}^T \cdot \boldsymbol{A} = 0 \wedge \boldsymbol{c} \neq \boldsymbol{0} \;|[\boldsymbol{c}]_1 \leftarrow \mathsf{A}(\mathtt{prm}, [\boldsymbol{A}]_2)\right]$$

where $\boldsymbol{A} \leftarrow\!\$\ \mathcal{D}_k$ *and* $\mathtt{prm} \leftarrow\!\$\ \mathsf{Setup}_{BG}(1^\lambda)$.

More specifically, in this paper we consider the KerLin_2 assumption which is equivalent to the $\mathcal{D}'_k$-KerMDH where $\mathcal{D}'_\lambda$ outputs matrix of the form: $\mathcal{D}'_k = \left\{ \begin{pmatrix} 1 & 1 \\ a_1 & 0 \\ 0 & a_2 \end{pmatrix} : a_0, a_1 \in \mathbb{Z}_p \right\}$.

2.2 All-but-Many Encryption

An all-but-many encryption scheme (ABM-Enc) is a tuple $\mathcal{ABME} = (\mathsf{Gen}, \mathsf{Sample}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{EquivEnc}, \mathsf{FakeEnc})$ such that: (1) Gen upon input the security parameter 1^λ outputs $(\mathsf{pk}, (\mathsf{sk}^s, \mathsf{sk}^e))$. The public key pk defines an *tag space* that we denote with $\mathcal{U}$ and a message space $\mathcal{M}$. (2) Sample upon input $(\mathsf{pk}, \mathsf{sk}^e)$ and $t \in \{0,1\}^\lambda$ outputs $u \in \mathcal{U}_{\mathsf{pk}}$. (3) Enc upon input $\mathsf{pk}, (t, u)$ and a message $\mu \in \mathcal{M}$ outputs a ciphertext C. (4) Dec upon input $\mathsf{sk}^e, (t, u)$ and a ciphertext C outputs a message μ. (5) $\mathsf{FakeEnc}$ upon input $\mathsf{pk}, (t, u), \mathsf{sk}^s$ outputs a ciphertext C and auxiliary information aux. (6) $\mathsf{EquivEnc}$ upon input (t, u) and aux and a message μ outputs random coins r; Let $\mathcal{L}^s_{\mathsf{pk}} = \{(t, u) : t \in \{0,1\}^\lambda, u \leftarrow \mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^e, t)\}$ and let $\mathcal{L}^e = \{0,1\}^\lambda \times \mathcal{U}_{\mathsf{pk}} \setminus \mathcal{L}^s$. (For simplicity we will omit the subscript pk when it is clear from the context.) We require that an ABM-Enc satisfy the following properties:

Pseudorandomness. For every PPT adversary A the following advantage is negligible:

$$\mathtt{Adv}^{pprf}_{\mathcal{ABME}}(\lambda) := \left| \Pr\left[\begin{matrix} (\mathsf{pk}, \mathsf{sk}^e, \mathsf{sk}^s) \leftarrow\!\$\ \mathsf{Gen}(1^\lambda) \\ \mathsf{A}(\mathsf{pk})^{\mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, \cdot)} = 1 \end{matrix} \right] - \Pr\left[\begin{matrix} (\mathsf{pk}, \mathsf{sk}^e, \mathsf{sk}^s) \leftarrow\!\$\ \mathsf{Gen}(1^\lambda) \\ \mathsf{A}(\mathsf{pk})^{\mathcal{O}_{\mathsf{pk}}(\cdot)} = 1 \end{matrix} \right] \right|$$

Where the oracle $\mathcal{O}_{\mathsf{pk}}(\cdot)$ samples at random from the distribution $\mathcal{U}_{\mathsf{pk}}$

Unforgeability. For every PPT adversary A the following advantage is negligible:

$$\mathtt{Adv}^{unf}_{\mathcal{ABME}}(\lambda) := \Pr\left[(t^*, u^*) \in \mathcal{L}^e, t^* \notin \mathcal{Q} \;:\; \begin{matrix} (\mathsf{pk}, \mathsf{sk}^e, \mathsf{sk}^s) \leftarrow\!\$\ \mathsf{Gen}(1^\lambda) \\ (t^*, u^*) \leftarrow \mathsf{A}(\mathsf{pk})^{\mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, \cdot)} \end{matrix} \right]$$

Where $\mathcal{Q}$ is the set of query made by A to the oracle Sample.

[2] Also notice that we quantify the extractor after the sampler, so to avoid pathological situation where the adversary A simply forwards the output of the sampler $\mathcal{S}$.

Dual Mode. The scheme can work in two different modes:

- **Decryption Mode:** For all $k\lambda \in \mathbb{N}$ For a hybrid linearly homomorphic commitment scheme all $\mathsf{pk}, \mathsf{sk}^e, \mathsf{sk}^s \in \mathsf{Gen}(1^\lambda)$ and all $\tau = (t, u) \in \mathcal{L}^e$ and all $\mu \in \mathcal{M}$ it holds that $\mathsf{Dec}(\mathsf{sk}^e, \tau, \mathsf{Enc}(\mathsf{pk}, \tau, \mu)) = \mu$.
- **Trapdoor Mode:** For all $k\lambda \in \mathbb{N}$ all $\mathsf{pk}, \mathsf{sk}^e, \mathsf{sk}^s \in \mathsf{Gen}(1^\lambda)$ all $\tau = (t, u) \in \mathcal{L}^s$ and all $\mu \in \mathcal{M}$ it holds that let $C, aux \leftarrow \mathsf{FakeEnc}(\mathsf{pk}, \tau, \mathsf{sk}^s)$ and $r \leftarrow \mathsf{EquivEnc}(\tau, aux, \mu)$ then $C = \mathsf{Enc}(\mathsf{pk}, \tau, \mu; r)$.

Moreover for all $(\mathsf{pk}, \mathsf{sk}^s, \mathsf{sk}^e) \in \mathsf{Gen}(1^\lambda)$ all $t \in \{0,1\}^\lambda$ and $\mu \in \mathcal{M}$ the following ensembles are statistically indistinghuishable:

$$\left\{ (u, C, r) : \begin{array}{c} u \leftarrow\!\!\$\ \mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^e, t), \\ r \leftarrow\!\!\$\ \{0,1\}^\lambda, \\ C \leftarrow \mathsf{Enc}(\mathsf{pk}, \tau, \mu; r) \end{array} \right\}_{\lambda \in \mathbb{N}} \quad \text{and} \quad \left\{ (u, c, r) : \begin{array}{c} u \leftarrow\!\!\$\ \mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^e, t), \\ C, aux \leftarrow \mathsf{FakeEnc}(\mathsf{pk}, \tau, \mathsf{sk}^s), \\ r \leftarrow \mathsf{EquivEnc}(\tau, aux, \mu) \end{array} \right\}_{\lambda \in \mathbb{N}}$$

Theorem 1 (Fujisaki, [19]). *If DDH assumption holds then there exist an ABM-Enc scheme. Moreover, the scheme admits an algorithm* $\widetilde{\mathsf{Gen}}$ *that obliviously samples the public parameter.*

2.3 Homomorphic Trapdoor Commitment Schemes

A trapdoor commitment scheme $\mathcal{COM} = (\mathsf{Setup}, \mathsf{Com}, \mathsf{ECom}, \mathsf{EOpen})$ is a tuple of algorithms specified as follows: (1) Algorithm Setup takes as input the security parameter and outputs a verification key ϑ and a trapdoor ψ; (2) Algorithm Com 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 $Com = \mathsf{Com}(\vartheta, m; r)$. (3) Algorithm ECom takes as input ψ and outputs a pair (Com, aux); (4) Algorithm EOpen takes as input (ψ, m, aux) and outputs $r \in \mathcal{R}$. We recall the standard security notions of *trapdoor hiding* and *computationally binding*. Roughly speaking, the former says that given a trapdoor is possible to create *fake commitments* using ECom which later on can be equivocated to open any message, and that the distribution of the tuple fake commitment equivocated opening is indistinguishable from the the distribution of a real commitment and its honest opening. The latter instead says that for any PPT adversary it is unfeasible to find two different messages/opening for the same commitment. We state the properties formally in Appendix B. For simplicity in the exposition we set $\mathcal{M}$ and $\mathcal{R}$ to be $\mathbb{Z}_p$ for a prime p. We say that $\mathcal{COM}$ is *linearly homomorphic* in the following sense: Given commitments Com and Com' (that commit to m and m') and $a \in \mathbb{Z}_p$, one can compute commitments $Com^* := a \cdot Com + Com'$ that opens to $a \cdot m + m'$. We will write the We will write the mappings as $Com^* = \mathsf{Com}(\vartheta, a \cdot m + m'; a \cdot r + r')$.

Moreover, we will require the following additional property. Let $(\vartheta, \psi) \leftarrow \mathsf{Setup}(1^\lambda)$, $(Com_1, aux_1) \leftarrow \mathsf{ECom}(\vartheta, \psi)$ and $(Com_2, aux_2) \leftarrow \mathsf{ECom}(\vartheta, \psi)$. Then we can use auxiliary information $a \cdot aux_1 + aux_2$ to equivocate $a \cdot Com_1 + Com_2$.

Finally, we consider commitment schemes where there is an oblivious way to sample the verification key which we denote with $\widetilde{\mathsf{Setup}}(1^\lambda)$.

2.4 (Quasi-Adaptive) NIZK and NIWI argument systems

Let $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$ be an NP-relation; the language associated with $\mathcal{R}$ is $\mathcal{L}_\mathcal{R} := \{x : \exists w \text{ s.t. } (x, w) \in \mathcal{R}\}$. We assume that $(x, w) \in \mathcal{R}$ is efficiently verifiable. An argument system $\mathcal{NIZK} := (\mathsf{Init}, \mathsf{P}, \mathsf{V})$ for $\mathcal{R}$ is a tuple of PPT algorithms specified as follows: (1) The initialization algorithm Init takes as input the security parameter 1^λ, and creates a common reference string (CRS) $\mathsf{crs} \in \{0,1\}^*$; (2) The prover algorithm P takes as input the CRS crs, a pair (x, w) such that $(x, w) \in \mathcal{R}$, and produces a proof $\pi \leftarrow\!\!\$\ \mathsf{P}(\mathsf{crs}, x, w)$; (3) The verifier algorithm V takes as input the CRS crs, a pair (x, π), and outputs a decision bit $\mathsf{V}(\mathsf{crs}, x, \pi)$. Additionally, we say that argument system is quasi-adaptive if the CRS generator algorithm Init takes as additional input the NP-relation $\mathcal{R}$ (or more formally a description of it).

$\mathbf{Exp}^{\texttt{one-more}}_{\mathcal{SS},\mathsf{A}}(\lambda,\ell,\gamma)$:

1. $(vk,\mathsf{sk}) \leftarrow\!\!\$\ \mathsf{Gen}(1^\lambda; r_0)$, return vk to A; let $\alpha = r_0$.
2. Run $\mathsf{A}(vk)$ with oracle access to $\mathsf{Sign}(\mathsf{sk},\cdot)$ and the leakage oracle.
   - Upon query $m \in \mathcal{M}$ to the signature oracle, let $\sigma := \mathsf{Sign}(\mathsf{sk}, m; r)$, $r \leftarrow\!\!\$\ \{0,1\}^\lambda$ and udpate the state $\alpha := \alpha \cup \{r\}$.
   - Upon query f to the leakage oracle, return $f(\alpha)$ where α is the current state.
3. Let $\mathcal{Q}$ be the set of signing queries issued by A, and let $\Lambda \in \{0,1\}^*$ be the total amount of information leaked by the adversary. A outputs n pairs $(m_1^*,\sigma_1^*),\ldots,(m_n^*,\sigma_n^*)$.
4. The experiment outputs 1 if and only if the following conditions are satisfied:
   (a) $\mathsf{Verify}(vk,(m_i^*,\sigma_i^*)) = 1$ and $m_i^* \notin \mathcal{Q}$, for all $i \in [n]$.
   (b) The messages $m_1^*,\ldots,m_n^*$ are pairwise distinct.
   (c) $n \geq \lfloor \ell/(\gamma \cdot s) \rfloor + 1$, where $s := |\sigma|$ and $|\Lambda| \leq \ell$.

Fig. 2: The fully-leakage one-more unforgeability experiment.

We consider distribution $\mathcal{D}_R$ over NP-relation. As for all the algorithms in this paper, the distribution can depends on the parameters $\mathtt{prm}$ (for example, $\mathtt{prm}$ could be the description of a bilinear group). Completeness means that for all CRSs crs output by $\mathsf{Init}(1^\lambda)$ (or for $\mathcal{R}$ and crs output by $\mathsf{Init}(1^\lambda,\mathcal{R})$ in the quasi-adaptive case), and for all pairs $(x,w) \in \mathcal{R}$, we have that $\mathsf{V}(\mathsf{crs},x,\mathsf{P}(\mathsf{crs},x,w)) = 1$ with all but a negligible probability. We require the following security properties (cf. Appendix C).

- **Perfect zero-knowledge:** Honestly computed proofs do not reveal anything beyond the validity of the statement, meaning that they can be perfectly simulated given only the statement itself (and a trapdoor information).
- **Adaptive knowledge soundness:** For any adversary that with input the CRS produces a valid NIZK proof for a statement x there exists an extractor (which might take as input a trapdoor information) that outputs the witness w such that $(x,w) \in \mathcal{R}$.
- **Statistical witness-indistinghuishability:** Given two different witness w, w' for the same instance x, a proof generated with the witness w is statistically indistinguishable to a proof generated with the witness w'.
- **Adaptive soundness:** No PPT adversary can forge a verifying proof for an adaptively chosen invalid statement.

We consider arguments that admit oblivious sampling of the CRS and we denote it with $\tilde{\mathsf{Init}}$.

3 Fully-Leakage One-More Unforgeability

A signature scheme is a triple of algorithms $\mathcal{SS} = (\mathsf{Gen},\mathsf{Sign},\mathsf{Verify})$ defined as follows: (1) The key generation algorithm takes as input the security parameter λ and outputs a verification key/signing key pair $(vk,\mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$; (2) The signing algorithm takes as input a message $m \in \mathcal{M}$ and the signing key sk and outputs a signature $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk},m)$; (3) The verification algorithm takes as input the verification key vk and a pair (m,σ) and outputs a bit $\mathsf{Verify}(vk,(m,\sigma)) \in \{0,1\}$. 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 Gen, we have that $\mathsf{Verify}(vk,(m,\mathsf{Sign}(\mathsf{sk},m)))$ returns 1. Given a signature scheme $\mathcal{SS}$, consider the experiment in Fig. 2 running with a PPT adversary A and parametrized by the security parameter $\lambda \in \mathbb{N}$, the leakage parameter $\ell \in \mathbb{N}$, and the slack parameter $\gamma := \gamma(\lambda)$:

Definition 3 (Fully-leakage one-more unforgeability). *We say that* $\mathcal{SS} = (\mathsf{Gen},\mathsf{Sign},\mathsf{Verify})$ *is* (ℓ,γ)-fully-leakage one-more unforgeable *if for every PPT adversary* A *we have that:*

$$\mathtt{Adv}^{\texttt{one-more}}_{\mathcal{SS},\mathsf{A}}(\lambda,\ell,\gamma) := \Pr\left[\mathbf{Exp}^{\texttt{one-more}}_{\mathcal{SS},\mathsf{A}}(\lambda,\ell,\gamma) = 1\right] \in \mathtt{negl}(\lambda).$$

Key Generation. Let $d, \mu \in \mathbb{N}$ be parameters. Let $\mathcal{NIWI} = (\mathsf{Init}, \mathsf{P}, \mathsf{V})$ be a NIWI argument system for the following polynomial-time relation:

$$\mathcal{R} := \left\{ (\vartheta, \mathsf{pk}, \tau, Com, C); (m^*, r^*, s) \;\middle|\; \begin{array}{c} Com = \mathsf{Com}(\vartheta, m^*; r^*) \\ C = \mathsf{Enc}(\mathsf{pk}, \tau, m^*; s) \end{array} \right\} .$$

Run $hk \leftarrow\!\$\ \mathsf{Gen}_{CRH}(1^\lambda)$ $\mathsf{crs} \leftarrow \tilde{\mathsf{Init}}(1^\lambda)$, $\vartheta \leftarrow \tilde{\mathsf{Setup}}(1^\lambda)$ and $\mathsf{pk} \leftarrow\!\$\ \tilde{\mathsf{Gen}}(1^\lambda)$.
Sample $\boldsymbol{\Delta} \leftarrow\!\$\ \mathbb{Z}_p^{\mu,d+1}$ and $\mathbf{r} = (r_0, \ldots, r_d) \leftarrow\!\$\ \mathcal{R}^{d+1}$, and compute $Com_i \leftarrow \mathsf{Com}(\vartheta, \boldsymbol{\delta}_i; r_i)$ for $i \in [0, d]$, where $\boldsymbol{\delta}_i \in \mathbb{Z}_p^\mu$ is the j-th column of $\boldsymbol{\Delta}$. Let $\boldsymbol{Com} = (Com_0, \ldots, Com_d)$
Output

$$\mathsf{sk} = (\boldsymbol{\Delta}, \mathbf{r}) \qquad vk = (\mathsf{crs}, \vartheta, \mathsf{pk}, \boldsymbol{Com}).$$

Signature. To sign a message $m \in \mathbb{Z}_p$ compute $m^* \leftarrow \boldsymbol{\Delta}(m)$ and $r^* \leftarrow \mathbf{r}(m)$. Pick $u \leftarrow \mathcal{U}_{\mathsf{pk}}$ and set $\tau = (\mathsf{H}(hk, m), u)$ and compute $C \leftarrow \mathsf{Enc}(\mathsf{pk}, \tau, m^*; s)$ where $s \leftarrow\!\$\ \mathcal{R}$.
Generate a NIWI argument π for $(\vartheta, \mathsf{pk}, \tau, \boldsymbol{Com}(m), C)$, using the witness (m^*, r^*, s). Output $\sigma = (C, \tau, \pi)$.

Verification. Given a pair (m, σ) and the verification key vk, parse σ as $(C, \tau = (t, u), \pi)$ and parse vk as $(\mathsf{crs}, \vartheta, \mathsf{pk}, \boldsymbol{Com})$ Output 1 if and only if:

$$\mathsf{H}(hk, m) = t \quad \wedge \quad \mathsf{V}(\mathsf{crs}, \pi, (\vartheta, \mathsf{pk}, \tau, \boldsymbol{Com}(m), C)).$$

Fig. 3: The signature scheme $\mathcal{SS}_1$.

Moreover, let $\mathcal{A}$ be the class of PPT adversaries which leakage functions do not depends on r_0 (see step 1 of the security experiment). We say that $\mathcal{SS}$ is (ℓ, γ)-fully-leakage one-more unforgeable with leak-free keygen *if for every adversary $\mathsf{A} \in \mathcal{A}$ the above equation holds.*

The number of signatures the adversary has to forge depends on the length of the leakage. In particular (ℓ, γ)-fully-leakage one-more unforgeability implies standard unforgeability for any adversary asking no leakage. The slack parameter γ specifies how close to optimal security $\mathcal{SS}$ is. In particular, in case $\gamma = 1$ one-more unforgeability requires that 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).

4 Signature scheme based on ABM-Enc Schemes

Our scheme $\mathcal{SS} = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ has message space equal to $\mathbb{Z}_p$ and is described in Fig. 3. The scheme is based on a homomorphic commitment scheme $\mathcal{COM}$, an ABM-Enc scheme $\mathcal{ABME}$, a NIWI argument system $\mathcal{NIWI}$ and a CRH function $(\mathsf{Gen}_{CRH}, \mathsf{H})$. The scheme follows the basic template described in Sec. 1, however instead of using just one single polynomial $\boldsymbol{\delta} \in \mathbb{Z}_p[X]$ of degree d, we use $\mu \in \mathbb{N}$ different polynomials arranged in the matrix $\boldsymbol{\Delta}$.

The correctness follows from the correctness of the NIWI argument system, and from linearly homomorphic property.

Theorem 2. *Let $\mu \in \mathbb{N}$ and let p be a prime larger than 2^λ. Assume that: (i) the commitment scheme $\mathcal{COM}$ is a trapdoor hiding, linearly homomorphic with message space $\mathbb{Z}_p^\mu$; (ii) the $\mathcal{ABME}$ is a secure ABME-Enc scheme with message space $\mathbb{Z}_p^\mu$ and ciphertext length s_1; (iii) $\mathcal{NIWI}$ is a statistical non-interactive witness indistinguishable argument system for the relation $\mathcal{R}$ described in Fig. 3 with proof length s_2. Then, let $s = s_1 + s_2$ and let $\gamma = \mu \log p / s$, for any $0 \leq \ell \leq (d \log \lambda) - \lambda$, the above signature scheme $\mathcal{SS}_1$ is (ℓ, γ)-fully-leakage one-more unforgeable.*

The security follows similarly to the proof in [15], thanks to the property of the ABM Encryption scheme the analysis is actually even easier.For space reason, we defer the proof in Appendix D.

Proof Sketch. The proof follows the standard games-hops paradigm. We define a series of hybrids and show that the hybrids are indistinguishable. Let A be an adversary that wins the fully-leakage one-more unforgeability game with probability ε. We denote with $((m_1^*, \sigma_1^* = (C_1^*, \tau_1^*, \pi_1^*)), \ldots,$

$(m_n^*, \sigma_n^* = (C_n^*, \tau_n^*, \pi_n^*)))$ the list of forgeries of A and with $(r_0, \boldsymbol{\Delta}, \mathbf{r}, (s_j, t_j)_{j\in[q]})$ the secret state. Notice that, because of the oblivious sampling of the parameters, the randomness r_0 such that $vk, \mathsf{sk} = \mathsf{KGen}(1^\lambda; r_0)$ can be computed efficiently as function of the verification key vk, we therefore omit r_0 from the state α. For each hybrid $\mathbf{H}_i$ we compute the winning probability ε_i of A:

Hybrid $\mathbf{H}_0$: the hybrid is the fully-leakage one more unforgeability game but where we condition on the soundness of the forged NIWI proofs.
Hybrid $\mathbf{H}_1$: We switch the way the parameters are sampled, specifically we use do not use the oblivious sampling, so that we gets the secret keys $\mathsf{sk}^s, \mathsf{sk}^e$ of the ABM-Enc and the equivocation trapdoor ψ of the commitments scheme.
Hybrid $\mathbf{H}_2$: We equivocate the commitments $\boldsymbol{Com}$ and, upon leakage oracle query, we recompute the equivocated randomness as function of the secret key $\boldsymbol{\Delta}$, so that the full secret state α can be written as $((\boldsymbol{\Delta}, \boldsymbol{r}(\boldsymbol{\Delta})),\ (s_j, z_j)_{i\in[q]})$.
Hybrid $\mathbf{H}_3$: For each j, at the j-th signature oracle query, we sample the tag $\tau_j = (t_j, u_j)$ differently, specifically, the element u_j is computed as $\mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, t_j)$. The indistinguishability comes from the pseudorandomness property of the ABM-Enc scheme.
Hybrid $\mathbf{H}_4$: For each j, at the j-th signature oracle query, we compute the encryption C_j differently. Specifically, we compute C_j using the trapdoor mode $C_j, aux_j \leftarrow\$ \mathsf{FakeEnc}(\mathsf{pk}, \tau_j, \mathsf{sk}^s)$, successively we compute the randomness $s_j(\boldsymbol{\Delta})$ as function of the secret key using the algorithm $\mathsf{EquivEnc}$. The indistinguishability comes from the dual mode property of the ABM-Enc scheme.
Hybrid $\mathbf{H}_5$: For each j, at the j-th signature oracle query with message m_j, we compute the NIWI proof π_j using a different witness. Specifically, instead of using the real witness $(\boldsymbol{\Delta}(m_j), \mathbf{r}(m_j), s_j)$ we use the witness $(0, r'_j, s'_j)$ where r'_j is an opening of the equivocated commitment $\boldsymbol{Com}(m_j)$ to 0 and s'_j is an opening of the fake encryption to 0. Moreover, we compute randomness $z_j(\boldsymbol{\Delta})$ (as function of $\boldsymbol{\Delta}$) by sampling z_j from the set $\{z : \pi_j = \mathsf{P}(\mathsf{crs},\ (\vartheta, \mathsf{pk}, \tau_j, \boldsymbol{Com}(m), C_j),\ (0, r'_j, s'_j)\}$. The key point which was originally proved by Boyle, Segev and Wichs [**?**], and then abstracted by Faonio, Nielsen, Venturi [15], is that, if the argument system is statistical NIWI then the sampled randomness z_j is distributed as it should have if we first sampled z_j and then computed the proof for the honest statement. Notice that this hybrid cannot be computed efficiently, however we are using a statistical security property of the underlying building block.

Notice that in $\mathbf{H}_5$ the secret state α can be computed as a function of the secret key $\boldsymbol{\Delta}$, so any function $f(\alpha)$ could be rephrased as a function $f'(\boldsymbol{\Delta})$. Moreover, by indistinguishability, the value ε_5, namely, the winning probability of A in $\mathbf{H}_5$ is approximately ε. We can now describe a predictor P with oracle access to $\boldsymbol{\Delta}$ which guesses $\boldsymbol{\Delta}$. The predictor runs the experiment $\mathbf{H}_4$ and leaks at most ℓ bits of information from $\boldsymbol{\Delta}$, then when the adversary outputs its forgery, the predictor extracts $n+1$ evaluation points of the polynomials $\boldsymbol{\delta}_i$ (for each i) and guesses the remaining $d-n$. The predictor guesses $\boldsymbol{\Delta}$ with probability approximately $1/p^{\mu(n-d)} \cdot \varepsilon$, meaning that the average conditional min-entropy of $\boldsymbol{\Delta}$ given the leakage provided by $\mathbf{H}_5$ is upper-bounded by roughly $\mu(n-d)\log p + \log \varepsilon$. On the other hand, we can prove using a standard argument that the average conditional min-entropy of $\boldsymbol{\Delta}$ in $\mathbf{H}_5$ is lower-bounded by $d\mu \log p - \ell$. So, since we set $\gamma = \log p/s$, and therefore $n \approx \ell/\log p + 1$, it must be that the value ε is negligible to have both equations simultaneously true.

Concrete Instantiation. As mentioned in Sec. 1 we instantiate the ABM-Scheme with the construction $\mathcal{ABME}_{DDH}$ of [19] based on DDH assumption, the NIWI argument system with Groth-Sahai [24] and the trapdoor commitment with the classical Pedersen commitment scheme. A ciphertext C of $\mathcal{ABME}_{DDH}$ is composed by $5\lambda/\log(\lambda)$ groups elements and the encryption procedure can be described by $5\lambda\log(\lambda)$ pairing-product equations. The message space can be parsed as $\mathbb{Z}_n^{\lambda/\log\lambda}$ where $n = \mathtt{poly}(\lambda)$ and its "encoded in the exponent". We additionally need $O(\lambda/\log\lambda)$ equations to describe that the plaintext and the opening of the commitment match. Summing up, the value s in the theorem is equal to $O(\lambda/\log\lambda)$. Finally, we notice that since we use the same groups for NIWI and $\mathcal{ABME}_{DDH}$ we need to use the external Diffie-Hellman (SXDH) assumption.

Let $\mathcal{COM} := (\mathsf{Setup}, \mathsf{Com})$ be the following commitment scheme:

Setup. The algorithm Setup upon input group parameters $\mathtt{prm}$ parse it as $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, G_1, G_2, G_T)$, pick at a random $[\boldsymbol{g}]_1 \leftarrow\!\!\$\ \mathbb{G}_1^{\mu}$, $\alpha \leftarrow\!\!\$\ \mathbb{Z}_p$ and $[h]_1 \leftarrow\!\!\$\ \mathbb{G}_1$, set $[\mathcal{M}]_1 \leftarrow (1, \alpha)^T \cdot [\boldsymbol{g}^T, h]_1$, set $[\boldsymbol{h}]_1 = [h, \alpha \cdot h]_1^T$ be the last column of $[\mathcal{M}]_1$, and set $[\alpha]_2$. Output the verification key $\vartheta = ([\mathcal{M}]_1, [\alpha]_2)$.

Commit. The algorithm Com upon input the verification key $[\mathcal{M}]_1, [\alpha]_2$ a message $\boldsymbol{m} \in \mathbb{Z}_p^{\mu}$, pick a uniformly random $r \leftarrow\!\!\$\ \mathbb{Z}_p$ and set $Com = [\mathcal{M}] \cdot (\boldsymbol{m}^T, r)^T \in \mathbb{G}_1^2$. The opening of the commitment is the randomizer r.

Let $\mathcal{NIZK}_{ext} = (\mathsf{Init}, \mathsf{P}, \mathsf{V})$ be the following QA-NIZK argument system:

Init. Let $\mathtt{prm}$ be the parameters defining a bilinear group, the algorithm Init upon input a matrix $[\boldsymbol{H}]_1 \in \mathbb{G}_1^{n,t}$ (and the parameter $\mathtt{prm}$) samples $\beta \leftarrow\!\!\$\ \mathbb{Z}_p$, $\boldsymbol{A} \leftarrow\!\!\$\ \mathcal{D}_k$ and $\boldsymbol{K} \leftarrow\!\!\$\ \mathbb{Z}_q^{n,k}$, it computes $\boldsymbol{P} \leftarrow \boldsymbol{H}^T \cdot \boldsymbol{K}$, $\boldsymbol{C} \leftarrow \boldsymbol{K} \cdot \boldsymbol{A}$ and $\boldsymbol{P}' \leftarrow \beta \cdot \boldsymbol{P}$, $\boldsymbol{C}' \leftarrow \beta \cdot \boldsymbol{C}$ and it outputs $\mathsf{crs} = ([\boldsymbol{P}]_1, [\boldsymbol{P}']_1, [\boldsymbol{C}]_2, [\boldsymbol{C}']_2, [\boldsymbol{A}]_2)$.

Prove. The algorithm P upon input crs and a tuple $[\boldsymbol{y}]_1, \boldsymbol{x}$ such that $[\boldsymbol{y}]_1 = [\boldsymbol{H}]_1 \cdot \boldsymbol{x}$ outputs (π, π') such that:

$$\pi = \boldsymbol{x}^T \cdot [\boldsymbol{P}]_1 \text{ and } \pi = \boldsymbol{x}^T \cdot [\boldsymbol{P}']_1.$$

Verify. The algorithm V upon input crs and a tuple $[\boldsymbol{y}]_1, \pi$ output 1 iff:

$$e(\pi, [\boldsymbol{A}]_2) = e([\boldsymbol{y}^T]_1, [\boldsymbol{C}]_2) \text{ and } e(\pi', [\boldsymbol{A}]_2) = e([\boldsymbol{y}^T]_1, [\boldsymbol{C}']_2).$$

Fig. 4: The commitment scheme $\mathcal{COM}$ and the QA-NIZK argument of knowledge $\mathcal{NIZK}_{ext}$.

5 A Signature Scheme based on Knowledge of the Exponent Assumption

Before describing the signature scheme we give more details on the building blocks used. Consider the commitment scheme $\mathcal{COM} := (\mathsf{Setup}, \mathsf{Com})$ (with implicit parameters an integer μ and a group generator Setup_{BG}) described in Fig 4. Notice that for any two messages $\boldsymbol{m}_0, \boldsymbol{m}_1$ and randomness r_0 there exists an unique assignment for r_1 such that $[\mathcal{M}]_1 \cdot (\boldsymbol{m}_0^T, r_0)^T = [\mathcal{M}]_1 \cdot (\boldsymbol{m}_0^T, r_1)^T$ holds, therefore $\mathcal{COM}$ is perfectly hiding.

The quasi-adaptive NIZK argument system in Fig 4 is a variation of [28]. As opposed to adaptive soundness of the scheme of Kiltz and Wee, the argument system $\mathcal{NIZK}_{ext}$ has a (weak flavor of) adaptive knowledge soundness[3]. Roughly speaking, the scheme is a two-fold version of the scheme of Kiltz and Wee. For technical reason, our scheme is secure only for distribution $\mathcal{D}_R$ that are witness sampleable. Given a distribution $\mathcal{D}_R$ (with parameter the description of a blinear group) over matrices $[\boldsymbol{H}]_1 \in \mathbb{G}^{n,t}$ we say that the distribution $\mathcal{D}_R$ is *witness sampleable* if there exists another efficiently sampleable distribution $\mathcal{D}'_R$ over matrices $\boldsymbol{H}' \in \mathbb{Z}_p^{n,t}$ such that $[\boldsymbol{H}]_1 \equiv [\boldsymbol{H}']_1$. For space reason, we defer the proof of the following theorem in Appendix C.

Theorem 3. *The quasi-adaptive argument system $\mathcal{NIZK}_{ext}$ in Fig. 4 is perfect zero-knowledge and if the $\mathcal{D}_k$-KerMDH assumption and the 1-KE* assumption hold for Setup_{BG} then the argument system is adaptive weak knowledge sound.*

The Signature Scheme. The signature scheme $\mathcal{SS}_2$ is described in Fig. 5. We first show that the scheme is correct. Specifically, for any tuple m, σ where σ is a valid signature for m with verification key $vk = (\mathsf{crs}, \vartheta, \boldsymbol{Com})$, let parse σ as C, π, we have:

$$\begin{aligned}\boldsymbol{Com}(m) - C &= \textstyle\sum_i Com_i \cdot m^i - C = \sum_i [\mathcal{M}]_1 \cdot (\boldsymbol{\delta}_i^T, r_i)^T \cdot m^i - [\mathcal{M}]_1 \cdot (\boldsymbol{\Delta}(m)^T, s)^T = \\ &= [\mathcal{M}]_1 \cdot \textstyle\sum_i (\boldsymbol{\delta}_i^T, r_i)^T \cdot m^i - [\mathcal{M}]_1 \cdot (\boldsymbol{\Delta}(m)^T, s)^T = \\ &= [\mathcal{M}]_1 \cdot ((\boldsymbol{\Delta}(m)^T, \mathbf{r}(m))^T - (\boldsymbol{\Delta}(m)^T, s)^T) = [\boldsymbol{h}]_1 \cdot (\mathbf{r}(m) - s).\end{aligned}$$

The last equation follows by noticing that $[\boldsymbol{h}]_1$ is the last column on $[\mathcal{M}]_1$. The correctness of the signature scheme follows by the equation above and the correctness of the quasi-adaptive NIZK scheme.

[3] In particular, we reverse the order of the quantifiers of the usual definition of knowledge soundness. Namely, for each adversary A there exists an extractor Ext. See Appendix C for more details.

Let $\mathcal{SS}_2 = (\mathsf{KGen}, \mathsf{Sign}, \mathsf{Verify})$ with message space $\mathbb{Z}_p$ be defined as follow:

Key Generation. Let $d, \mu \in \mathbb{N}$ be parameters. Let $\mathtt{prm} \leftarrow\!\$\ \mathsf{Setup}_{BG}(1^\lambda)$ be parameter describing an asymmetric bilinear group , let $\vartheta = [\mathcal{M}]_1, [\alpha]_2 \leftarrow\!\$\ \mathsf{Setup}(\mathtt{prm})$ and let $[\boldsymbol{h}]_1$ be the last column of $[\mathcal{M}]_1$. Consider the NP relation $\mathcal{R}$ defined as follow:

$$\mathcal{R} = \{(\boldsymbol{Y}, r) : \boldsymbol{Y} = r \cdot [\boldsymbol{h}]_1\}$$

Run $\mathsf{crs}, tp \leftarrow \mathsf{Init}(1^\lambda, \mathcal{R})$, sample $\boldsymbol{\Delta} \leftarrow\!\$\ \mathbb{Z}_p^{\mu,d+1}$ and $\mathbf{r} = (r_0, \ldots, r_d) \leftarrow\!\$\ \mathbb{Z}_p^{d+1}$, and compute commitments $Com_i \leftarrow \mathsf{Com}(\vartheta, \boldsymbol{\delta}_i; r_i)$ for $i \in [0, d]$, where $\boldsymbol{\delta}_i \in \mathbb{Z}_p^\mu$ is the j-th column of $\boldsymbol{\Delta}$. Let $\boldsymbol{Com} = (Com_i)_{i=0}^d$ and output

$$\mathsf{sk} = (\boldsymbol{\Delta}, \mathbf{r}) \qquad vk = (\mathsf{crs}, \vartheta, \boldsymbol{Com}).$$

Signature. To sign a message $m \in \mathbb{Z}_p$ compute $m^* = \boldsymbol{\Delta}(m)$ and let $C \leftarrow \mathsf{Com}(\vartheta, m^*, s)$ for $s \leftarrow\!\$\ \mathbb{Z}_p$, and compute $\pi \leftarrow \mathsf{P}(\mathsf{crs}, (\mathbf{r}(m) - s) \cdot [\boldsymbol{h}], \mathbf{r}(m) - s)$.
Output $\sigma = (C, \pi)$.

Verification. Given a pair (m, σ) and the verification key vk, parse σ as (C, π) and parse vk as $(\mathsf{crs}, \vartheta, \boldsymbol{Com})$. Output 1 if and only if $\mathsf{V}(\mathsf{crs}, \boldsymbol{Com}(m) - C, \pi)$ and $e(C_0, [1]_2) = e(C_1, [\alpha]_2)$.

Fig. 5: The signature scheme $\mathcal{SS}_2$.

Theorem 4. *Let $\mu, d \in \mathbb{N}$ and $\mu > 8$. If the $(\mu+1)$-KE* assumption and the KerLin$_2$ assumption hold over Setup_{BG} then, for any $0 \leq \ell \leq (d \log \lambda) - \lambda$, the signature scheme $\mathcal{SS}_2$ described Fig. 5 is $(\ell, 1)$-fully-leakage one-more unforgeable with leak-free key generation.*

Proof. Let A be an adversary such that $\mathtt{Adv}^{\mathtt{one\text{-}more}}_{\mathsf{A},\mathcal{SS}_2}(\lambda, \ell, 1) = \varepsilon$ for parameter ℓ as described in the statement of the theorem. Let $\mathbf{H}_0(\lambda)$ be the experiment $\mathbf{Exp}^{\mathtt{one-more}}_{\mathcal{SS},\mathsf{A}}(\lambda)$. Denote with $((m_1^*, (C_1^*, \pi_1^*)), \ldots, (m_n^*, (C_n^*, \pi_n^*)))$ the list of forgeries of A. During the experiment the adversary has oracle access to $\alpha = (\boldsymbol{\Delta}, \mathbf{r},\ (s_j, z_j)_{j \in [q]})$ where s_j is the randomness used by Com and z_j is the randomness used by P (the prover of the NIZK proof system). We now define a series of hybrid experiments. Let $\mathtt{Forge}_i$ (resp. $\mathtt{Forge}_{i,j}$) be the event that $\mathbf{H}_i$ (resp. $\mathbf{H}_{i,j}$) returns 1, so that $\mathsf{P}[\mathtt{Forge}_0] = \varepsilon$.

Hybrid 1. The hybrid $\mathbf{H}_1$ runs the same as the hybrid $\mathbf{H}_0$ but with a slightly different syntax. More in details, consider the following sampler $\mathcal{S}$:

Sampler $\mathcal{S}([\mathcal{M}]_1, [\alpha]_2)$:

1. Sample $r_\mathsf{A} \leftarrow \{0,1\}^\lambda$ and $\boldsymbol{\Delta} \leftarrow\!\$\ \mathbb{Z}_p^{\mu,d+1}, \mathbf{r} \leftarrow \mathbb{Z}_p^{d+1}$, set $\mathsf{sk} = (\boldsymbol{\Delta}, \mathbf{r})$ and compute the verification key vk as described in KGen using $[\mathcal{M}]_1$;
2. Run $\mathsf{A}(vk; r_\mathsf{A})$ and answer all the signature oracle queries using $\mathsf{Sign}(\mathsf{sk}, \cdot)$ and the leakage oracle queries with α. Let $\mathsf{View} = (\sigma_1, \ldots, \sigma_q, \mathsf{Leak})$ be the full transcript of the interactions between A and the oracles;
3. Output $(vk, \mathsf{View}, r_\mathsf{A})$.

The hybrid $\mathbf{H}_1$ (1) creates the parameters $(\mathtt{prm}_{BG}, [\mathcal{M}]_1, [\alpha]_2)$, (2) executes the sampler $(vk, \mathsf{View}, r_\mathsf{A}) \leftarrow\!\$\ \mathcal{S}([\mathcal{M}]_1, [\alpha]_2)$, (3) runs $\mathsf{A}(vk; r_\mathsf{A})$ and answers all the oracle queries using the information in View. (4) Eventually it runs the extractors $\mathsf{Ext}'_1, \ldots, \mathsf{Ext}'_n$ and outputs 1 if and only if all the forged signatures verify correctly for vk and all the messages are different. The change between the two hybrids is only syntactical, therefore $\varepsilon_0 = \varepsilon_1$.

Hybrid 2.i. The hybrid $\mathbf{H}_{2.i}$ takes as parameters i different extractors $\mathsf{Ext}_1, \ldots, \mathsf{Ext}_i$ and runs the same as the hybrid $\mathbf{H}_1$ but, also, it runs the extractors and output 1 if and only if the extracted values match the commitments $C_1^*, \ldots, C_i^*$. More in details, the hybrid $\mathbf{H}_{2.i}$ first creates the parameters $(\mathtt{prm}_{BG}, [\mathcal{M}]_1, [\alpha]_2)$, then it executes the sampler $(vk, \mathsf{View}, r_\mathsf{A}) \leftarrow \mathcal{S}([\mathcal{M}]_1, [\alpha]_2)$, then it runs $\mathsf{A}(vk; r_\mathsf{A})$ and answers all the oracle queries using the information in View. Eventually, A outputs its forgeries $(m_1^*, \sigma_1^*), \ldots, (m_n^*, \sigma_n^*)$ where $\sigma_i^* = (C_i^*, \pi_i^*)$, and for $j = 1, \ldots, i$ the hybrid $\mathbf{H}_{1.i}$ computes $\boldsymbol{x}_i \leftarrow \mathsf{Ext}_i(([\mathcal{M}]_1, [\alpha]_2), (vk, \mathsf{View}, r_\mathsf{A}))$ and outputs 1 if and only if:

(a) all the forged signatures verify correctly for vk and all the messages are different and,

(b) for any $j = 1, \ldots, i$ we have $C_j^* = [\mathcal{M}]_1 \cdot \boldsymbol{x}_j$.

Claim. There exist PPT extractors $\mathsf{Ext}_1, \ldots, \mathsf{Ext}_n$ such that for any $i > 1$, $|\varepsilon_{1.i-1} - \varepsilon_{1.i}| \in \mathtt{negl}(\lambda)$. Moreover, $\varepsilon_1 = \varepsilon_{2.0}$.

Proof. First we prove second sentence of the claim. The change between $\mathbf{H}_1$ and $\mathbf{H}_{2.0}$ is only syntactical. In fact, the winning condition is the same in both hybrids, as $\mathbf{H}_{2.0}$ does not check the condition (b). Now we prove the first sentence. We define an adversary A_i' for the $(\mu+1)$-KE* assumption:

Adversary $\mathsf{A}_i'([\mathcal{M}]_1; r')$:
1. Parse r' as $(vk, \mathsf{View}, r_{\mathsf{A}})$;
2. Run $\mathsf{A}(vk; r_{\mathsf{A}})$ and answers all the oracle queries using the information in View;
3. Eventually, A outputs its forgeries $(m_1^*, \sigma_1^*), \ldots, (m_n^*, \sigma_n^*)$;
4. If all the forged signatures verify correctly for vk and all the messages are different parse σ_i^* as C_i^*, π_i^* and output $\boldsymbol{Y} := C_i^*$.

For any PPT Ext_i the two hybrids diverge when $([\mathcal{M}]_1 \cdot \boldsymbol{x}_i \neq \boldsymbol{Y})$, where $\boldsymbol{x}_i$ is the output of the extractor, but the signature σ_i^* verifies correctly. Notice that the verification algorithm check that $e(\boldsymbol{Y}_0, [\alpha]_2) = e(\boldsymbol{Y}_1, [1]_2)$, and so $\boldsymbol{Y} \in \mathit{Span}([1, \alpha]_1)$. Therefore:

$$|\varepsilon_{1.i-1} - \varepsilon_{1.i}| \leq \Pr\left[[\mathcal{M}]_1 \cdot \boldsymbol{x}_i \neq \boldsymbol{Y} \ \wedge \ \boldsymbol{Y} \in \mathit{Span}([1, \alpha]_1)\right]$$

We can apply the security of the $\mu+1$-KE* assumption. In particular, there must exist an extractor Ext_i such that the difference above is negligible.

Hybrid 3.*i*. The hybrid $\mathbf{H}_{3.i}$ takes as parameters i different extractors $\mathsf{Ext}_1', \ldots, \mathsf{Ext}_i'$ and runs the same as the hybrid $\mathbf{H}_{2.n}$ but also for any $j = 1, \ldots, i$ it computes $w_i \leftarrow \mathsf{Ext}_i'(\mathsf{crs}, tp, r')$ where $r' = (\boldsymbol{\Delta}, \mathbf{r}, [\boldsymbol{g}, h], \alpha)$ and the winning conditions are changed as follow:

(a) All the forged signatures verify correctly for vk and all the messages are different,
(b) for any $j = 1, \ldots, n$ we have $C_j^* = [\mathcal{M}]_1 \cdot \boldsymbol{x}_j$ and,
(c) for any $j = 1, \ldots, i$ check $\boldsymbol{Com}(m_i^*) - C_i^* = w_i \cdot [h, \alpha h]$.

Claim. There exist PPT extractors $\mathsf{Ext}_1', \ldots, \mathsf{Ext}_n'$ such that for any $i > 1$, $|\varepsilon_{1.i-1} = \varepsilon_{2.i}| \in \mathtt{negl}(\lambda)$. Moreover, $\varepsilon_{3.0} = \varepsilon_{2.n}$.

Proof. Clearly $\varepsilon_{3.0} = \varepsilon_{2.n}$, as the point (c) is not checked in $\mathbf{H}_{3.0}$. We define an adversary A_i' for the adaptive weak knowledge soundness of the QANIZK $\mathcal{NIZK}_{ext}$:

Adversary $\mathsf{A}_i'(\mathsf{crs}; r')$:
1. Parse r' as $(\boldsymbol{\Delta}, \mathbf{r}, [\boldsymbol{g}^T, h], \alpha)$, define ϑ as described in Setup, compute $\boldsymbol{Com}$ using $\boldsymbol{\Delta}$ and $\mathbf{r}$, and set the verification key $vk = (\mathsf{crs}, \vartheta, \boldsymbol{Com})$ and $\mathsf{sk} = (\boldsymbol{\Delta}, \mathbf{r})$;
2. Run $\mathsf{A}(vk; r_{\mathsf{A}})$ and answer all the signature oracle queries using $\mathsf{Sign}(\mathsf{sk}, \cdot)$ and the leakage oracle queries with α;
3. Eventually, A outputs its forgeries $(m_1^*, \sigma_1^*), \ldots, (m_n^*, \sigma_n^*)$;
4. If all the forged signatures verify correctly for vk and all the messages are different parse σ_i^* as C_i^*, π_i^* and output the statement $(\boldsymbol{Com}(m_i^*) - C_i^*)$ and the proof π_i^*.

For any PPT Ext_i' the two hybrids diverge when conditions (a) and (b) holds but $\boldsymbol{Com}(m_i^*) - C_i^* \neq w_i \cdot [h, \alpha h]$ happens. Clearly, condition (a) implies that $\boldsymbol{Com}(m_i^*) - C_i^* \in \mathit{Span}([h, \alpha h])$ as the signature verification check it explicitly. Let Ext_i' be the extractor prescribed by the QA-NIZK adaptive weak knowledge soundness property, if the above event happens with noticeable probability, then the adversary A_i' breaks adaptive weak knowledge soundness of the $\mathcal{NIZK}_{ext}$.

Hybrid 4. The hybrid $\mathbf{H}_4$ is the same as $\mathbf{H}_{3.n}$ but the winning condition are changed as follow:

(a) All the forged signatures verify correctly for vk and all the messages are different,
(b) for any $j = 1, \ldots, n$ we have $C_j^* = [\mathcal{M}]_1 \cdot \boldsymbol{x}_j$,
(c) for any $j = 1, \ldots, n$ check $\boldsymbol{Com}(m_i^*) - C_i^* = w_i \cdot [h, \alpha h]$ and,
(d) for any $j = 1, \ldots, n$, let $\boldsymbol{x}'_j$ be the projection of $\boldsymbol{x}_j$ to the first μ coordinates, check $\boldsymbol{x}'_j = \boldsymbol{\Delta}(m_j^*)$.

Claim. $|\varepsilon_{3.n} = \varepsilon_4| \in \mathtt{negl}(spar)$.

Proof. The two hybrids diverge conditions (a),(b),(c) hold but one of the extracted values $\boldsymbol{x}_i$ is such that $\boldsymbol{x}'_i \neq \boldsymbol{\Delta}(m_i^*)$. We show that the probability of this event is negligible. To do so we reduce to the representation problem over $(\mathbb{G}_1, p, [\boldsymbol{g}]_1)$. The representation problem asks to find two vector $\boldsymbol{x}, \boldsymbol{y}$ such that $\boldsymbol{x} \neq \boldsymbol{y}$ but $[\boldsymbol{g}]_1^T \cdot \boldsymbol{x} = [\boldsymbol{g}]_1^T \cdot \boldsymbol{y}$. It is well known that if DLOG problem over $(\mathbb{G}_1, p)$ is hard then representation problem over $(\mathbb{G}_1, p, [\boldsymbol{g}]_1)$ for random $[\boldsymbol{g}]_1$ is hard too. Consider the following adversary for the representation problem:

Adversary $\mathsf{B}([\boldsymbol{g}]_1)$:

1. Run the hybrid $\mathbf{H}_4$ with the parameter set to $[\boldsymbol{g}]_1$, in particular sample $[h]_1 \leftarrow \beta \cdot [g_{j^*}]_1$ where $\beta \leftarrow\$\ \mathbb{Z}_p$, the index $j \leftarrow\$\ [\mu]$, the secret key $\boldsymbol{\Delta} \leftarrow\$\ \mathbb{Z}_p^{\mu, d+1}$ and $\mathbf{r} \leftarrow \mathbb{Z}_p^{d+1}$.
2. If the winning conditions (a),(b),(c) are met but not condition (d), then let i be the index such that $\boldsymbol{x}'_i \neq \boldsymbol{\Delta}(m_i^*)$.
3. Parse $\boldsymbol{x}_i$ as $(x_{i,1}, \ldots, x_{i,\mu+1})$ and $\boldsymbol{\Delta}(m_i^*)$ as $(y_1, \ldots, y_\mu)$, output the vectors

$$\bar{\boldsymbol{x}} = (x_{i,1}, \ldots, x_{i,j^*} + \beta \cdot (x_{i,\mu+1} + w - y_{\mu+1}), \ldots, x_{i,\mu}) \text{ and } \bar{\boldsymbol{y}} = (y_1, \ldots, y_\mu).$$

Let k be an index such that $\boldsymbol{x}'_{i,k} \neq \boldsymbol{\Delta}(m_i^*)_k$ then with probability $1 - 1/\mu$ the index $k \neq j^*$ (because j^* is information theoretically hidden), and when this happens then $\bar{\boldsymbol{x}}$ and $\bar{\boldsymbol{y}}$ are different.

Moreover, notice that, because $\mathsf{Com}(m_i^*) = C_i^* + w_i \cdot [h, \alpha h]$ and $h = \beta \cdot g_{j^*}$ we have that $[\boldsymbol{g}]_1^T \cdot \bar{\boldsymbol{x}} = [\boldsymbol{g}]_1^T \cdot \bar{\boldsymbol{y}}$. So the adversary B breaks the representation problem for $(\mathbb{G}_1, p, [\boldsymbol{g}]_1)$.

Hybrid 5. The hybrid $\mathbf{H}_5$ is the same as $\mathbf{H}_4$ but we revert the changes introduced in the hybrids $\mathbf{H}_{2,i}$ for all $i \in [n]$. The winning condition are changed and in particular they are less stringent as do not consider the condition (c). As the condition is not checked then the hybrid does not need to execute the extractors Ext'_i for $i \in [n]$. Notice that the set of conditions are relaxed, so the probability of the event cannot decrease, namely $\varepsilon_5 \geq \varepsilon_4$.

The Predictor P. The predictor runs the same as the hybrid $\mathbf{H}_5$ but the sampler $\mathcal{S}$ is run *externally*. In particular, the parameters for $\mathcal{S}$ are sampled, then first the sampler is executed and then the predictor P is executed with input the output produced by $\mathcal{S}$. Eventually, the predictors (which runs internally A) receives n forgeries $(m_1^*, \sigma_1^*), \ldots, (m_n^*, \sigma_n^*)$. The predictor checks the winning conditions (a),(b),(d) of the hybrid $\mathbf{H}_4$ and if they hold, for $j \in [\mu]$ it samples a polynomial $\boldsymbol{\delta}_j^*$ in $\mathbb{Z}_p[X]$ of degree d such that $\boldsymbol{\delta}_j(m_i^*) = \boldsymbol{x}'_{i,j}$ for $i \in [n]$, and it outputs $\boldsymbol{\Delta}^* = (\boldsymbol{\delta}_1^*, \ldots, \boldsymbol{\delta}_\mu^*)$.
Recall that the advantage of A in the one-more unforgeability game is ε:

Lemma 1. $\Pr[\mathsf{P}(\mathcal{S}([\mathcal{M}]_1, [\alpha]_2)) = \boldsymbol{\Delta}] \geq \exp(((n-d) \cdot \mu) \log p) \cdot (\varepsilon - \mathtt{negl}(\lambda))$.

Proof. By the triangular inequality and the claims above we have that $\varepsilon_5 \geq \varepsilon - \mathtt{negl}(\lambda)$. When the event $\mathtt{Forge}_5$ happens then $\boldsymbol{\Delta}(m_i^*) = \boldsymbol{x}'_i$ for $i \in [n]$ so the event that $\boldsymbol{\Delta}^* = \boldsymbol{\Delta}$ is equivalent to the event that the predictor P correctly guesses the remaining $d - n$ zeros of the polynomials $\boldsymbol{\delta}_i$ for $i \in [\mu]$ which is equal to $1/p^{\mu(d-n)} = \exp(((n-d) \cdot \mu) \log p)$.

Lemma 2. *For any parameter* $\mathtt{prm} \leftarrow\$\ \mathsf{Setup}_{BG}(1^\lambda)$ *and any* $([\mathcal{M}]_1, [\alpha]_2) \in \mathbb{G}_1^{2,\mu+1} \times \mathbb{G}_2$ *we have* $\widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \mathcal{S}([\mathcal{M}]_1, [\alpha]_2)) \geq |\boldsymbol{\Delta}| - \ell$.

Proof. We define two samplers $\mathcal{S}_1$ and $\mathcal{S}_2$, we prove that their output distribution is equivalent to the distribution of $\mathcal{S}$, and moreover, their outputs are independent of the secret material $\boldsymbol{\Delta}$ as sampled by $\mathcal{S}$.

The sampler $\mathcal{S}_1$ executes the same of $\mathcal{S}$ but the elements $\boldsymbol{Com}$, the signature queries, and the leakage oracle queries are computed in the following way:

- The elements $\boldsymbol{Com}$ are sampled as uniformly element from $Span([g, \alpha g])$.
- At the j-th signature oracle query with message m the element C_j is sampled as uniformly element from $Span([g, \alpha g])$.
- Define the function $\mathbf{r}(\boldsymbol{\Delta})$ that outputs the vector $(r_0, \ldots, r_d)$ computing r_i such that $Com_i = [\mathcal{M}]_1 \cdot (\boldsymbol{\delta}_i^T, r_i)^T$. Similarly, define the functions $s_j(\boldsymbol{\Delta})$ that output the vector s_j such that $C_j = [\mathcal{M}]_1 \cdot (\boldsymbol{\delta}_i^T, s_j)^T$. For each leakage oracle query f the answer of f is computed as $f(\boldsymbol{\Delta}, \mathbf{r}(\boldsymbol{\Delta}), (s_i(\boldsymbol{\Delta}), z_i)_{i \leq q})$.

Claim. For any parameter $\mathtt{prm} \leftarrow\!\!\$\ \mathsf{Setup}_{BG}(1^\lambda)$ and any $([\mathcal{M}]_1, [\alpha]_2) \in \mathbb{G}_1^{2,\mu+1} \times \mathbb{G}_2$ the outputs of the samplers $\mathcal{S}$ and $\mathcal{S}_1$ are identically distributed, $\mathcal{S}([\mathcal{M}]_1, [\alpha]_2) \equiv \mathcal{S}_1([\mathcal{M}]_1, [\alpha]_2)$.

Proof. We notice that for any $\boldsymbol{m}$ the commitment to $\boldsymbol{m}$ is uniformly distributed over $Span([1, \alpha])$. Therefore, for any Com_i (resp. C_i), it always exists such r_i (resp. s_i), and moreover, once $\boldsymbol{\Delta}$ and $\boldsymbol{Com}$ (resp. C_i) are fixed its value is uniquely defined.

The sampler $\mathcal{S}_2$ executes the same of $\mathcal{S}_1$ but, for all the signatures, the NIZK proofs π_i are computed using the simulator S of NIZK and, moreover, the randomness $z_i \leftarrow\!\!\$\ \{z_i : \pi_i = \mathsf{P}(\mathsf{crs}, (\boldsymbol{r}(\boldsymbol{\Delta})(m_i) - \boldsymbol{s}(\boldsymbol{\Delta})) \cdot [\boldsymbol{h}], (\boldsymbol{r}(\boldsymbol{\Delta})(m_i) - \boldsymbol{s}(\boldsymbol{\Delta})))\}$ where $\boldsymbol{r}(\boldsymbol{\Delta})$ is the vector of the randomness as computed by $\mathcal{S}_1$.

Claim. For any parameter $\mathtt{prm} \leftarrow\!\!\$\ \mathsf{Setup}_{BG}(1^\lambda)$ and any $([\mathcal{M}]_1, [\alpha]_2) \in \mathbb{G}_1^{2,\mu+1} \times \mathbb{G}_2$ the outputs of the samplers $\mathcal{S}_1$ and $\mathcal{S}_2$ are identically distributed, $\mathcal{S}_1([\mathcal{M}]_1, [\alpha]_2) \equiv \mathcal{S}_2([\mathcal{M}]_1, [\alpha]_2)$.

Proof. By the perfect zero-knowledge property of the quasi-adaptive NIZK, the proofs π_i are distributed equivalently to the real proofs. Notice that perfect zero-knowledge implies that the set of the simulated proofs and the set of real proofs (for any instance and witness) is exactly the same. Moreover, for all i, we sample s_i' uniformly at random from the set of possible randomness that match with the proof π_i, therefore s_i' is equivalently distributed to s_i, the randomness used to compute the proofs in $\mathcal{S}_1$. We write $z_i(\boldsymbol{\Delta})$ to stress that z_i is computed as function of $\boldsymbol{\Delta}$, for each leakage oracle query f the answer of f is computed as $f(\boldsymbol{\Delta}, \mathbf{r}(\boldsymbol{\Delta}), (s_i(\boldsymbol{\Delta}), z_i(\boldsymbol{\Delta}))_{i \leq q})$.

Claim. $\widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \mathcal{S}_3([\mathcal{M}]_1, [\alpha]_2)) \geq \mu \cdot (d+1) \log p - \ell$.

Proof. Let q be the number of signature queries made by A and let Leak the concatenation of all the leakage performed by A:

$$\begin{aligned}\widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \mathcal{S}_2([\mathcal{M}]_1, [\alpha]_2)) &= \widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid vk, \mathsf{View}, r_\mathsf{A}) = \widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \mathsf{View}) \\ &= \widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \sigma_1, \ldots, \sigma_q, \mathsf{Leak}) = \widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \mathsf{Leak}) \geq \mu \cdot (d+1) \log p - \ell.\end{aligned}$$

Where the second equality holds because vk, r_A are sampled independent from $\boldsymbol{\Delta}$ (Lemma 4), the forth because the signatures $\sigma_1, \ldots, \sigma_q$ are sampled independent from $\boldsymbol{\Delta}$, and last inequality we simply applied the chain rule (Lemma 3) and notice that $|\boldsymbol{\Delta}| = \mu(d+1) \log p$ and $|\mathsf{Leak}| = \ell$.

By putting together the claims above and by Lemma 5 the lemma follows.

Applying the definition of the average conditional min-entropy to Lemma 1 and combining it with Lemma 2:

$$d \log p - \ell \leq (d - n) \log p - \log(\varepsilon - \mathtt{negl}(\lambda))$$

By easy calculation we can derive that $\ell \geq n \log p + \log(\varepsilon - \mathtt{negl}(\lambda))$, and by the fact that $n \geq \frac{\ell}{s \cdot \gamma} + 1$ and $\gamma = 1$ we can derive that: $-\log(\varepsilon - \mathtt{negl}(\lambda)) \geq s \geq \lambda$. Notice that, for the equation above to hold, we necessarily need that ε is at least negligible in λ.

References

1. M. Abe and S. Fehr. Perfect NIZK with adaptive soundness. In S. P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 118–136. Springer, Heidelberg, Feb. 2007.
2. J. Alwen, Y. Dodis, M. Naor, G. Segev, S. Walfish, and D. Wichs. Public-key encryption in the bounded-retrieval model. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 113–134. Springer, Heidelberg, May 2010.
3. J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 36–54. Springer, Heidelberg, Aug. 2009.
4. G. Ateniese, A. Faonio, and S. Kamara. Leakage-resilient identification schemes from zero-knowledge proofs of storage. In J. Groth, editor, *15th IMA International Conference on Cryptography and Coding*, volume 9496 of *LNCS*, pages 311–328. Springer, Heidelberg, Dec. 2015.
5. M. Bellare and A. Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In M. Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 273–289. Springer, Heidelberg, Aug. 2004.
6. N. Bitansky, R. Canetti, O. Paneth, and A. Rosen. On the existence of extractable one-way functions. In D. B. Shmoys, editor, *46th ACM STOC*, pages 505–514. ACM Press, May / June 2014.
7. E. Boyle, S. Goldwasser, A. Jain, and Y. T. Kalai. Multiparty computation secure against continual memory leakage. In H. J. Karloff and T. Pitassi, editors, *44th ACM STOC*, pages 1235–1254. ACM Press, May 2012.
8. E. Boyle, G. Segev, and D. Wichs. Fully leakage-resilient signatures. In K. G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 89–108. Springer, Heidelberg, May 2011.
9. Z. Brakerski, Y. T. Kalai, J. Katz, and V. Vaikuntanathan. Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In *51st FOCS*, pages 501–510. IEEE Computer Society Press, Oct. 2010.
10. I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In J. Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 445–456. Springer, Heidelberg, Aug. 1992.
11. Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs. Cryptography against continuous memory attacks. In *51st FOCS*, pages 511–520. IEEE Computer Society Press, Oct. 2010.
12. Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 523–540. Springer, Heidelberg, May 2004.
13. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. Villar. An algebraic framework for Diffie-Hellman assumptions. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, Aug. 2013.
14. A. Faonio and J. B. Nielsen. Fully leakage-resilient codes. In S. Fehr, editor, *PKC 2017, Part I*, volume 10174 of *LNCS*, pages 333–358. Springer, Heidelberg, Mar. 2017.
15. A. Faonio, J. B. Nielsen, and D. Venturi. Mind your coins: Fully leakage-resilient signatures with graceful degradation. In M. M. Halldórsson, K. Iwama, N. Kobayashi, and B. Speckmann, editors, *ICALP 2015, Part I*, volume 9134 of *LNCS*, pages 456–468. Springer, Heidelberg, July 2015.
16. A. Faonio, J. B. Nielsen, and D. Venturi. Fully leakage-resilient signatures revisited: Graceful degradation, noisy leakage, and construction in the bounded-retrieval model. *Theor. Comput. Sci.*, 660:23–56, 2017.
17. D. Fiore and A. Nitulescu. On the (in)security of SNARKs in the presence of oracles. In M. Hirt and A. D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 108–138. Springer, Heidelberg, Oct. / Nov. 2016.
18. E. Fujisaki. All-but-many encryption - A new framework for fully-equipped UC commitments. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 426–447. Springer, Heidelberg, Dec. 2014.
19. E. Fujisaki. All-but-many encryption. *J. Cryptology*, 31(1):226–275, 2018.
20. K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic analysis: Concrete results. In Çetin Kaya. Koç, D. Naccache, and C. Paar, editors, *CHES 2001*, volume 2162 of *LNCS*, pages 251–261. Springer, Heidelberg, May 2001.
21. S. Garg, A. Jain, and A. Sahai. Leakage-resilient zero knowledge. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 297–315. Springer, Heidelberg, Aug. 2011.
22. R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic span programs and succinct NIZKs without PCPs. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
23. J. Groth. Short pairing-based non-interactive zero-knowledge arguments. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, Dec. 2010.
24. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In N. P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 415–432. Springer, Heidelberg, Apr. 2008.
25. S. Halevi and H. Lin. After-the-fact leakage in public-key encryption. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 107–124. Springer, Heidelberg, Mar. 2011.
26. D. Hofheinz. All-but-many lossy trapdoor functions. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 209–227. Springer, Heidelberg, Apr. 2012.
27. J. Katz and V. Vaikuntanathan. Smooth projective hashing and password-based authenticated key exchange from lattices. In M. Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 636–652. Springer, Heidelberg, Dec. 2009.

28. E. Kiltz and H. Wee. Quasi-adaptive NIZK for linear subspaces revisited. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 101–128. Springer, Heidelberg, Apr. 2015.
29. P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In N. Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 104–113. Springer, Heidelberg, Aug. 1996.
30. P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In M. J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 388–397. Springer, Heidelberg, Aug. 1999.
31. T. Malkin, I. Teranishi, Y. Vahlis, and M. Yung. Signatures resilient to continual leakage on memory and computation. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 89–106. Springer, Heidelberg, Mar. 2011.
32. M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 18–35. Springer, Heidelberg, Aug. 2009.
33. J. B. Nielsen, D. Venturi, and A. Zottarel. Leakage-resilient signatures with graceful degradation. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 362–379. Springer, Heidelberg, Mar. 2014.
34. B. Parno, J. Howell, C. Gentry, and M. Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.
35. T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In J. Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 129–140. Springer, Heidelberg, Aug. 1992.

A Information-Theoretic Lemmas

We state the following lemmas from Dodis *et al.* [12]

Lemma 3. *Let X, Y and Z be random variables. If Y has at most 2^ℓ possible values, then $\widetilde{\mathbb{H}}_\infty(X|Y,Z) \geq \widetilde{\mathbb{H}}_\infty(X,Y|Z) - \ell \geq \widetilde{\mathbb{H}}_\infty(X|Z) - \ell$.*

Lemma 4. *Let X, Y be random variables. If Y is independent of X then $\widetilde{\mathbb{H}}_\infty(X|Y) = \mathbb{H}_\infty(X)$.*

Lemma 5. *Let X, Y and Y' be random variables. If $(X,Y) \equiv (X,Y')$ then $\widetilde{\mathbb{H}}_\infty(X|Y) = \widetilde{\mathbb{H}}_\infty(X|Y)$.*

Lemma 6. *If X is uniformly random over a set of 2^ℓ elements, then $\mathbb{H}_\infty(X) = \ell$.*

B Commitment Schemes

B.1 Security properties

A trapdoor commitment $\mathcal{COM} = (\mathsf{Setup}, \mathsf{Com}, \mathsf{ECom}, \mathsf{EOpen})$ scheme has three properties, known a binding, extractability and trapdoor hiding.

Binding Property. Consider the following probability:

$$\mathsf{P}[\mathsf{Com}(\vartheta, m_0; r_0) = \mathsf{Com}(\vartheta, m_1; r_1) : \ \vartheta \leftarrow \mathsf{Setup}(1^\lambda); ((m_0, r_0), (m_1, r_1)) \leftarrow \mathsf{A}(\vartheta)].$$

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

Trapdoor Hiding Property. For all $(\vartheta, \psi) \leftarrow \mathsf{Setup}(1^\lambda)$ and for all $m \in \mathcal{M}$ the following probability distributions are indistinguishable:

$$\left\{ (Com, r) \;\middle|\; \begin{array}{c} r \leftarrow\!\$\ \mathcal{R}, \\ Com := \mathsf{Com}(\vartheta, m; r) \end{array} \right\} \text{ and } \left\{ (Com, r) \;\middle|\; \begin{array}{c} (Com, aux) \leftarrow\!\$\ \mathsf{ECom}(\vartheta, \psi), \\ r := \mathsf{EOpen}(\psi, m, aux) \end{array} \right\}.$$

Trapdoor hiding implies the less stringent notion of *perfect hiding* where for any two messages m_0, m_1 adaptively chosen as function of the verification key the distribution $\mathsf{Com}(\vartheta, m_0)$ and $\mathsf{Com}(\vartheta, m_1)$ are indistinghuishable.

C Quasi-Adaptive NIZK and NIWI argument systems

C.1 NIWI argument systems

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

Definition 4 (Adaptive soundness). *Let $\mathcal{NIWI}$ be a non-interactive argument system for a language $\mathcal{L}$. We say that $\mathcal{NIWI}$ satisfies adaptive soundness, if for all PPT adversaries A we have*

$$\Pr\left[\mathsf{V}(\mathsf{crs}, x, \pi) = 1 \wedge x \notin \mathcal{L} : \ \mathsf{crs} \leftarrow \mathsf{Init}(1^\lambda), (x, \pi) \leftarrow \mathsf{A}(1^\lambda, \mathsf{crs})\right] \in \mathtt{negl}(\lambda).$$

Definition 5 (Statistical witness indistinguishability). *Let $\mathcal{NIWI}$ be a non-interactive argument system for a relation $\mathcal{R}$. We say that $\mathcal{NIWI}$ satisfies statistical witness indistinguishability if for any triplet (x, w, w') such that $(x, w) \in \mathcal{R}$ and $(x, w') \in \mathcal{R}$, the distributions $\{(\mathsf{crs}, \pi) | \mathsf{crs} \leftarrow \mathsf{Init}(1^\lambda), \pi \leftarrow \mathsf{P}(\mathsf{crs}, x, w)\}$ and $\{(\mathsf{crs}, \pi) |\ \mathsf{crs} \leftarrow \mathsf{Init}(1^\lambda), \pi \leftarrow \mathsf{P}(\mathsf{crs}, x, w')\}$ are statistically indistinguishable.*

Experiment $\mathbf{Exp}^{\mathtt{snd}-\mathcal{QANIZK}}_{\mathsf{A},\mathcal{NIZK},\mathcal{D}_R}(\lambda)$:

1. $\mathtt{prm} \leftarrow\!\$\ \mathsf{Setup}_{BG}(\lambda), \rho \leftarrow\!\$\ \mathcal{D}_R(\mathtt{prm})$
   $\mathsf{crs}, tp \leftarrow\!\$\ \mathsf{Init}(\mathtt{prm}, \rho)$;
2. $(x, \pi) \leftarrow \mathsf{A}(\mathsf{crs})$;
3. Output 1 iff $x \notin \mathcal{L}_\rho$ and $\mathsf{V}(\mathsf{crs}, x, \pi) = 1$.

Experiment $\mathbf{Exp}^{\mathtt{knw}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{NIZK},\mathcal{D}_R}(\lambda)$:

1. $\mathtt{prm} \leftarrow\!\$\ \mathsf{Setup}_{BG}(\lambda)$, $\rho \leftarrow\!\$\ \mathcal{D}_R(\mathtt{prm})$,
   $\mathsf{crs}, tp \leftarrow\!\$\ \mathsf{Init}(\mathtt{prm}, \rho)$;
2. $(x, \pi) \leftarrow \mathsf{A}(\mathsf{crs}; r)$, $w \leftarrow \mathsf{Ext}(\mathsf{crs}; r)$ and $r \leftarrow\!\$\ \{0,1\}^\lambda$;
3. Output 1 iff $(x, w) \notin \mathcal{R}_\rho$ and $\mathsf{V}(\mathsf{crs}, x, \pi) = 1$.

Let $\mathcal{NIZK}_{KW} := (\mathsf{Init}, \mathsf{P}, \mathsf{V})$ be the defined as follow:

Init. Let $\mathtt{prm}$ be the parameters defining a bilinear group, the algorithm Init upon input a matrix $[\boldsymbol{H}]_1 \in \mathbb{G}_1^{n,t}$ where $n > t$ (and the parameter $\mathtt{prm}$) samples $\boldsymbol{A} \leftarrow\!\$\ \mathcal{D}_k$ and $\boldsymbol{K} \leftarrow\!\$\ \mathbb{Z}_q^{n,k}$, it computes $\boldsymbol{P} \leftarrow \boldsymbol{H}^T \cdot \boldsymbol{K}$, $\boldsymbol{C} \leftarrow \boldsymbol{K} \cdot \boldsymbol{A}$ and it outputs $\mathsf{crs} = ([\boldsymbol{P}]_1, [\boldsymbol{C}]_2, [\boldsymbol{A}]_2)$.

Prove. The algorithm P upon input crs and a tuple $[\boldsymbol{y}]_1, \boldsymbol{x}$ such that $[\boldsymbol{y}]_1 = [\boldsymbol{H}]_1 \cdot \boldsymbol{x}$ outputs $\pi = \boldsymbol{x}^T \cdot [\boldsymbol{P}]_1$.

Verify. The algorithm V upon input crs and a tuple $[\boldsymbol{y}]_1, \pi$ output 1 iff $e(\pi, [\boldsymbol{A}]_2) = e([\boldsymbol{y}^T]_1, [\boldsymbol{C}]_2)$.

Fig. 6: The QA-NIZK with adaptive soundness of Kiltz and Wee.

C.2 Quasi-Adaptive NIZK argument systems

We define both standard soundness and weak knowledge soundness for QA-NIZK argument systems.

Definition 6. *For any $\mathsf{A}, \mathcal{NIZK}$ and $\mathcal{D}_R$ define the following advantage:*

$$\mathtt{Adv}^{\mathtt{snd}-\mathcal{QANIZK}}_{\mathsf{A},\mathcal{NIZK},\mathcal{D}_R}(\lambda) := \Pr\left[\mathbf{Exp}^{\mathtt{snd}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{D}_R}(\lambda) = 1\right].$$

We say that $\mathcal{NIZK}$ is adaptive sound if for every PPT adversary A and for any distribution $\mathcal{D}_R$ $\mathtt{Adv}^{\mathtt{snd}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{NIZK},\mathcal{D}_R}(\lambda) \in \mathtt{negl}(\lambda)$.

Definition 7. *For any $\mathsf{A}, \mathsf{Ext}, \mathcal{NIZK}$ and $\mathcal{D}_R$ define the following advantage:*

$$\mathtt{Adv}^{\mathtt{knw}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{NIZK},\mathcal{D}_R}(\lambda) := \Pr\left[\mathbf{Exp}^{\mathtt{knw}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{D}_R}(\lambda) = 1\right].$$

We say that $\mathcal{NIZK}$ is adaptive weak knowledge soundness if for every PPT adversary A there exist a PPT extractor Ext such that for any distribution $\mathcal{D}_R$: $\mathtt{Adv}^{\mathtt{knw}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{NIZK},\mathcal{D}_R}(\lambda) \in \mathtt{negl}(\lambda)$.

We describe below the scheme of Kiltz and Wee in Fig. 6.

Theorem 5 (Kiltz and Wee, [28]). *The argument system $\mathcal{NIZK}_{KW}$ is a QA-NIZK argument. Furthermore, under $\mathcal{D}_k$−KerMDH Assumption for Setup_{BG}, it has adaptive soundness.*

In the following we prove that our scheme is adaptive weak knowledge sound. We restate the theorem from Sec. 5.

Theorem 3. *The quasi-adaptive argument system $\mathcal{NIZK}_{ext}$ in Fig. 4 is perfect zero-knowledge and if the $\mathcal{D}_k$-KerMDH assumption and the 1-KE* assumption hold for Setup_{BG} then the argument system is adaptive weak knowledge sound.*

Proof. We first prove that the argument system is adaptive soundness. This easily come from Theorem 5. In fact, given a PPT adversary A for $NIZK_{ext}$ we can create an adversary A' for $\mathcal{NIZK}_{KW}$. The adversary A' upon input a CRS of $\mathcal{NIZK}_{KW}$ and a matrix $[\boldsymbol{H}]$ samples $\beta \leftarrow\!\$\ \mathbb{Z}_q$ and computes $[\boldsymbol{P}']_1 \leftarrow \beta \cdot [\boldsymbol{P}]_1$ and $[\boldsymbol{C}']_1 \leftarrow \beta \cdot [\boldsymbol{C}]_1$ to create a CRS for $NIZK_{ext}$. Eventually, A outputs a tuple statement $[\boldsymbol{y}]$ and a proof π, π' and A' outputs $[\boldsymbol{y}], \pi$. It is easy to check that if A breaks soundness then A' does too.

Let A be an adversary for the adaptive (standard) soundness experiment, let $\mathtt{Win} := \mathbf{Exp}^{\mathtt{snd}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{D}_R}$ and let $\varepsilon := \Pr[\mathtt{Win}]$. Let Sound be the event that the the element $[\boldsymbol{y}]_1$ output by A is indeed in the column span of $[\boldsymbol{H}]$ and let $\mathtt{Forge}$ be the event that $\mathtt{Win} \wedge \mathsf{Sound}$, meaning that the proof verify and $[\boldsymbol{y}]_1$ is a valid instance. Let $\mathbf{H}_0$ be the hybrid experiment that it is equivalent to $\mathbf{Exp}^{\mathtt{snd}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{D}_R}$. We define $\varepsilon_i := \Pr[\mathtt{Forge}]$ where the probability is taken over the experiment $\mathbf{H}_i$. By the argument given above we know that $|\varepsilon - \varepsilon_1| \leq \Pr[\neg\mathsf{Sound}] \leq \mathtt{negl}(\lambda)$.

Let $\mathbf{H}_1$ be the same as $\mathbf{H}_0$ but the verification of the argument system additionally check that $e(\pi', [1]_2) = e(\pi, [\beta]_2)$.

Claim. $|\varepsilon_1 - \varepsilon_0| \leq \mathtt{negl}(\lambda)$.

Proof. The two hybrids diverge when the proof (π, π') verify, the instance $[\boldsymbol{y}]_1$ is in the language but $e(\pi, [1]_2) \neq e(\pi', [\beta]_2)$. We prove that, if the event happens with noticeable probability then we can break the $\mathcal{D}_k-\mathsf{KerMDH}$ Assumption in $\mathbb{G}_2$. Assuming that $e(\pi, [\beta]_2) \neq e(\pi', [1]_2)$ then it means that $\pi \cdot \beta - \pi' \neq 0$. On the other hand, the two verification equations tell us that $e(\pi, [\boldsymbol{A}]_2) = e([\boldsymbol{y}]_1, [\boldsymbol{C}]_2)$ and $e(\pi', [\boldsymbol{A}]_2) = e([\boldsymbol{y}]_1, [\beta \cdot \boldsymbol{C}]_2)$, and therefore $e(\pi \cdot \beta, [\boldsymbol{A}]_2) = e(\pi', [\boldsymbol{A}]_2)$. Now, let $[\boldsymbol{z}]_1 \leftarrow \pi \cdot \beta - \pi'$ we have that $e([\boldsymbol{z}]_1, [\boldsymbol{A}]_2) = 0$ buth $[\boldsymbol{z}]_1 \neq 0$, so we clearly break the $\mathcal{D}_k-\mathsf{KerMDH}$ assumption. (Notice that in the reduction we can sample $\beta \leftarrow\!\!\$\ \mathbb{Z}_p$).

Let $\mathbf{H}_{2,i}$ be an hybrid that takes as parameter and extractor $\mathsf{Ext}_1, \ldots, \mathsf{Ext}_i$, which runs the same as $\mathbf{H}_1$ but where at the end if all the conditions of $\mathbf{H}_1$ are met additionally runs $\tilde{x}_i \leftarrow \mathsf{Ext}([1, \beta]_1, [\beta]_2, r)$ where r is all the randomness of the experiment excluded the sampling of $[1, \beta]_1, [\beta]_2$ and the winning condition modified to be valid only for all $j \leq i$ we have $\pi_j = [\tilde{x}_i]$.

Claim. There exists PPT $\mathsf{Ext}_1, \ldots, \mathsf{Ext}_{k+1}$ such that $|\varepsilon_{2,k+1} - \varepsilon_{2,0}| \leq \mathtt{negl}(\lambda)$.

Proof. For any index i the hybrids $\mathbf{H}_{2,i-1}$ and $\mathbf{H}_{2,i}$ diverge when the extractor Ext_i does not output $\tilde{x}_i$ such that $\pi_j = [\tilde{x}_i]$. We define an adversary for the KE assumption on bilinear group.

Adversary $\mathsf{A}'_i([\beta]_1, [\beta]_2)$:

1. Create all the parameters of the crs, in particular, since $\mathcal{D}_R$ is witness sampleable, first sample $\boldsymbol{H} \leftarrow \mathbb{Z}_q^{n,t}$, $\boldsymbol{A} \leftarrow\!\!\$\ \mathcal{D}_k$ and $\boldsymbol{K} \leftarrow \mathbb{Z}_q^{n,k}$ and set $[\boldsymbol{P}]_1 = [\boldsymbol{M}^T \cdot \boldsymbol{K}]$, $[\boldsymbol{P}']_1 = [\beta]_1 \cdot (\boldsymbol{M}^T \cdot \boldsymbol{K})$ and similarly $[\boldsymbol{C}]_2$ and $[\boldsymbol{C}']_2 = [\beta]_2 \cdot (\boldsymbol{K} \cdot \boldsymbol{A})$.
2. Run the adversary A on the common reference string created and receive π, π'.
3. Outputs π_i, π'_i (the i-element of the vector π, resp. π').

Notice that the distribution of the CRS created by A'_i is exactly the same as the real CRS. Moreover, for this adversary there exists an extractor Ext_i such that the outputs $\tilde{x}_i$ of Ext_i is $\pi_i = [\tilde{x}_i]_1$, as otherwise we would break the KE assumption over bilinear groups.

Lastly we define an extractor for the knowledge soundness of $NIZK_{ext}$.

Extractor $\mathsf{Ext}([\beta]_1, [\beta]_2, \boldsymbol{K}; r)$:

- If $\boldsymbol{K}$ has rank strictly less than t then abort, else find $\boldsymbol{T}$ such let $(\boldsymbol{H}^T \cdot \boldsymbol{K}) \cdot \boldsymbol{T}$ is equal to $\boldsymbol{I}_t$ (the identity matrix).
- For $i = 1, \ldots k+1$ executes $\tilde{x}_i \leftarrow \mathsf{Ext}_i([\beta]_1, [\beta]_2; (r, \boldsymbol{K}))$;
- let $\tilde{\boldsymbol{x}} = (\tilde{x}_1, \ldots, \tilde{x}_{k+1})$ then output $\tilde{\boldsymbol{x}} \cdot \boldsymbol{K}^{-1}$.

Claim. $|\mathtt{Adv}^{\mathtt{knw}-\mathcal{QANIZK}}_{\mathsf{A},\mathsf{Ext},\mathcal{D}_R}(\lambda) - \varepsilon_{2,n}| \leq \mathtt{negl}(\lambda)$.

Proof. The only difference between the two experiment is that additionally Ext check that the matrix $\boldsymbol{K}$ has rank n. However, if we assume $k+1 > t$ then with overwhelming probability $\boldsymbol{K}$ has at least rank t. Moreover, since $\boldsymbol{H}$ has rank t then we can always find the matrix $\boldsymbol{T}$. Finally we need to check that what the extractor outputs is a valid witness, but notice that $[\tilde{\boldsymbol{x}}] = [\boldsymbol{x}^T \cdot \boldsymbol{H}^T \cdot \boldsymbol{K}]$, and so the output $\tilde{\boldsymbol{x}} \cdot \boldsymbol{T} = \boldsymbol{x}^T \cdot (\boldsymbol{H}^T \cdot \boldsymbol{K} \cdot \boldsymbol{T}) = \boldsymbol{x}^T$.

We can conclude the proof by noticing that the adaptive knowledge soundness of the argument system $NIZK_{ext}$ and its adaptive (standard) soundness are negligibly close.

D Proof of Theorem 4

Theorem 2. *Let $\mu \in \mathbb{N}$ and let p be a prime larger than 2^λ. Assume that: (i) the commitment scheme $\mathcal{COM}$ is a trapdoor hiding, linearly homomorphic with message space $\mathbb{Z}_p^\mu$; (ii) the $\mathcal{ABME}$ is a secure ABME-Enc scheme with message space $\mathbb{Z}_p^\mu$ and ciphertext length s_1; (iii) $\mathcal{NIWI}$ is a statistical non-interactive witness indistinguishable argument system for the relation $\mathcal{R}$ described in Fig. 3 with proof length s_2. Then, let $s = s_1 + s_2$ and let $\gamma = \mu \log p/s$, for any $0 \leq \ell \leq (d \log \lambda) - \lambda$, the above signature scheme $\mathcal{SS}_1$ is (ℓ, γ)-fully-leakage one-more unforgeable.*

Proof. Let A be an adversary such that $\mathtt{Adv}^{\mathtt{one\text{-}more}}_{\mathsf{A},\mathcal{SS}_1}(\lambda, \ell, \gamma) = \varepsilon$ for parameter ℓ, γ as described in the statement of the theorem. Let $\mathbf{H}_0(\lambda)$ be the experiment $\mathbf{Exp}^{\mathtt{one-more}}_{\mathcal{SS},\mathsf{A}}(\lambda)$. Denote with $((m_1^*, \sigma_1^* = (C_1^*, \tau_1^*, \pi_1^*)), \ldots, (m_n^*, \sigma_n^* = (C_n^*, \tau_n^*, \pi_n^*)))$ the list of forgeries of A. Let $\mathtt{Forge}_0$ be the event that $\mathbf{H}_0$ returns 1, so that $\mathsf{P}[\mathtt{Forge}_0] = \varepsilon$. Define $\mathtt{False}_0$ 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., $\mathtt{False}_0$ is verified if in $\mathbf{H}_0$ there exists $i \in [n]$ for which $\mathsf{Dec}(\mathsf{sk}^e, \tau_i^*, C_i^*) = m_i'$ and $m_i' \neq \boldsymbol{\Delta}(m_i^*)$. Define $\mathtt{Collision}_0$ to be the event that there exists $i \in [n]$ and $j \in [q]$, for which $\mathsf{H}(hk, m_i^*) = \mathsf{H}(hk, m_j)$ where m_j is the j-th signature oracle's query made by the adversary.

Let $\varepsilon_0 := \mathsf{P}[\mathtt{Forge}_0 \wedge \neg\mathtt{False}_0 \wedge \neg\mathtt{Collision}_0]$.

Claim. $\varepsilon - \varepsilon_0 \in \mathtt{negl}(\lambda)$.

Proof. The claim is proved in two steps, first we prove $|\varepsilon - \mathsf{P}[\mathtt{Forge}_0 \wedge \neg\mathtt{False}_0]| \in \mathtt{negl}(\lambda)$ and then we prove $|\mathsf{P}[\mathtt{Forge}_0 \wedge \neg\mathtt{False}_0] - \mathsf{P}[\mathtt{Forge}_0 \wedge \neg\mathtt{False}_0 \wedge \neg\mathtt{Collision}_0]| \in \mathtt{negl}(\lambda)$.

By adaptive computational soundness of the NIWI argument system, we must have that $\mathsf{P}[\mathtt{False}_0] \in \mathtt{negl}(\lambda)$. In fact, from an adversary A provoking $\mathtt{False}_0$, we can easily construct an adversary breaking adaptive soundness by simply emulating the entire experiment for A and outputing the proof for which the statement of the event holds (we can do this efficiently as we know sk^e).

The second part of the proof follows easily by collision resistance of H.

We now define a series of hybrid experiments. For each hybrid $\mathbf{H}_i$, we write ε_i for the probability of the event $\mathtt{Forge}_i \wedge \neg\mathtt{False}_i \wedge \neg\mathtt{Collision}_i$. During the experiment the adversary has leakage oracle access to $\alpha = (\boldsymbol{\Delta}, \mathbf{r}, (s_j, z_j)_{j\in[q]})$ where s_j is the randomness used by Enc and z_j is the randomness used by P (the proving algorithm of the NIWI). Notice that, because of the oblivious sampling of the parameters, the randomness r_0 such that $vk = \mathsf{KGen}(1^\lambda; r_0)$ can be computed efficiently as function of the verification key vk. To keep the exposition lighter, we therefore omit r_0 from the secret state α.

Hybrid 1. The experiment $\mathbf{H}_1$ is the same as $\mathbf{H}_0$, expect that the parameters are not sampled by the oblivious algorithms but instead we sample $(\mathsf{pk}, (\mathsf{sk}^s, \mathsf{sk}^e)) \leftarrow\!\$\ \mathsf{Gen}(1^\lambda)$ and $\vartheta, \psi \leftarrow\!\$\ \mathsf{Setup}(1^\lambda)$ and $\mathsf{crs} \leftarrow\!\$\ \mathsf{Init}(1^\lambda)$.

Claim. $\varepsilon_1 - \varepsilon_0 \in \mathtt{negl}(\lambda)$.

The claim follows by the oblivious sampling properties of $\tilde{\mathsf{Setup}}, \tilde{\mathsf{Gen}}$ and $\tilde{\mathsf{Init}}$. Details omitted.

Hybrid 2. The experiment $\mathbf{H}_2$ is the same as $\mathbf{H}_1$, except that now the commitments $\{Com_i\}_{i=0}^d$ to the columns $\boldsymbol{\delta}_i$ are replaced by equivocable commitments, i.e. $(Com_i, r_i') \leftarrow\!\$\ \mathsf{ECom}_1(\vartheta, \psi)$ for all $i \in [0, d]$. Notice that the actual randomness r_i, used to produce Com_i in $\mathbf{H}_0$, can be recovered efficiently as a function of the coefficients $\boldsymbol{\delta}_i$ and the fake randomness r_i', as $r_i(\boldsymbol{\Delta}) := \mathsf{EOpen}(\psi, \boldsymbol{\delta}_i, r_i')$. Given r_i the signature computation is identical. We write $\mathbf{r}(\boldsymbol{\Delta}) = (r_0(\boldsymbol{\Delta}), \ldots, r_d(\boldsymbol{\Delta}))$ the vector of randomness $\mathbf{r}$ computed as function of $\boldsymbol{\Delta}$ (with the help of ψ and $(r_0', \ldots, r_d')$).

Claim. $\varepsilon_2 - \varepsilon_1 \in \mathtt{negl}(\lambda)$.

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.

Hybrid 3. The experiment $\mathbf{H}_2$ is the same as $\mathbf{H}_1$, except that now the signature are computed differently. Specifically, for all j upon the j-th query with message m the signature oracle compute the signature by computing $u_j \leftarrow\$ \mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, t_j)$ where $t_j = \mathsf{H}(hk, m)$.

Claim. $\varepsilon_3 - \varepsilon_2 \in \mathtt{negl}(\lambda)$.

Proof. The trapdoor pseudorandomness property of the ABM-Enc scheme implies that, for any t_j the distribution $\mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, t_j)$ and $u_j \leftarrow\$ \mathcal{U}_{\mathsf{pk}}$ are computationally close. The claim follows by a standard hybrid argument over all the signature queries.

Let $\mathtt{ForgeTag}_i$ be the event that there exists a forgery (m^*, σ^*) of the adversary such that $\sigma^* = (C^*, (t^*, u^*), \pi^*)$ and $(t^*, u^*) \in \mathcal{L}^e$.

Let ε'_i be the probability of the event $\mathtt{Good}_i$ defined as follow:

$$\mathtt{Good}_i := \mathtt{Forge}_i \wedge \neg\mathtt{False}_i \wedge \neg\mathtt{Collision}_i \wedge \neg\mathtt{ForgeTag}_i.$$

Claim. $\varepsilon'_3 - \varepsilon_3 \in \mathtt{negl}(\lambda)$.

Proof. Notice that $\varepsilon'_3 - \varepsilon_3 \leq \Pr[\mathtt{ForgeTag}_i]$, so we bound the probability of this event. Consider the adversary B of the unforgeability experiment of ABM-Enc scheme that runs the experiment $\mathbf{H}_2$ but instead of computing u_j by its own it uses its oracle access to $\mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, \cdot)$. Eventually, the adversary A outputs n forgeries and if all of them are valid and for different messages then the adversary B picks an index $i \leftarrow\$ [n]$ and outputs t^*_i, u^*_i as its own forgery. It is easy to see that the adversary B wins with probability $\Pr[\mathtt{ForgeTag}_i]/n$ the unforgeability experiment of the AMB-Enc scheme.

Hybrid 4. The experiment $\mathbf{H}_4$ is the same as $\mathbf{H}_3$, except that now the signature are computed differently. Specifically, for all j upon the j-th query with message m the signature oracle compute the signature by computing $C_j, aux \leftarrow \mathsf{FakeEnc}(\mathsf{pk}, (t_j, u_j), \mathsf{sk}^s)$ and compute the randomness for the leakage oracle as $s_j \leftarrow \mathsf{EquivEnc}(\tau, aux, \boldsymbol{\Delta}(m))$. To stress that the randomness s_j can be computed as function of $\boldsymbol{\Delta}$, we write $s_j(\boldsymbol{\Delta}) := \mathsf{EquivEnc}(\tau, aux_j, \boldsymbol{\Delta}(m_j))$ where m_j is the message queried at the j-th signature oracle call.

Claim. $\varepsilon'_4 - \varepsilon'_3 \in \mathtt{negl}(\lambda)$.

Proof. The dual mode property of the ABM-Enc scheme implies that, for any $t_j \in \{0,1\}^\lambda$ and $u_j \leftarrow\$ \mathsf{Sample}(\mathsf{pk}, \mathsf{sk}^s, t_j)$ the distribution that computes the ciphertext C_j using Enc and the distribution that compute it with $\mathsf{FakeEnc}$ even given the equivocated randomness are statistically indistinguishable. The claim follows by a standard hybrid argument.

Hybrid 5. This experiment is identical to the previous hybrid, except that it uses a different witness w' to compute the NIWI arguments. In particular, given the j-th query m, the experiment generates the argument π by running

$$\mathsf{P}(\mathsf{crs}, \underbrace{(\vartheta, \mathsf{pk}, \tau_j, \boldsymbol{Com}(m), C_j)}_{x}, \underbrace{(0, \mathsf{EOpen}(\psi, 0^\mu, \boldsymbol{r}'(m)), \mathsf{EquivEnc}(\tau, aux_j, \mathbf{0})}_{w'}; z'_j),$$

where $\boldsymbol{r}'(m) = \sum_{i=0}^{d} r'_i \cdot m^i$ is computed using the randomness $\{r'_i\}_{i=0}^d$. Notice that the randomness z used to generate the NIWI argument in the previous experiment can be sampled (inefficiently) as a

function of the (real) witness $w := (\boldsymbol{\Delta}(m), \mathsf{EOpen}(\psi, \boldsymbol{\Delta}(m), r'(m)), \mathsf{EquivEnc}(\tau, aux, \boldsymbol{\Delta}(m)))$ and z'_j. In particular, $z_j(\boldsymbol{\Delta}) := z'_j$ where z'_j is sampled from the distribution $\{z : \pi_j = \mathsf{P}(\mathsf{crs},\ (\vartheta, \mathsf{pk}, \tau_j, \boldsymbol{Com}(m), C_j),\ (0, r'_j)\}$. The state used to answer the leakage query f is set to:

$$\alpha(\boldsymbol{\Delta}) = ((\boldsymbol{\Delta}, \mathbf{r}(\boldsymbol{\Delta})), (s_j(\boldsymbol{\Delta}), z_j(\boldsymbol{\Delta}))_{i\in[q]}).$$

Notice that the function $\alpha(\boldsymbol{\Delta})$ needs the values $(r'_i)_{i\in[d]}, (m_j, \tau_j, aux_j, z'_j)_{j\in[n]}$ to be computed. We hardwire such values in the definition of the function $\alpha(\cdot)$.

Claim. $\varepsilon'_4 - \varepsilon'_3 \in \mathtt{negl}(\lambda)$.

Proof. The linear homomorphic property of the hybrid commitment scheme ensures that the value $r'(m)$ is the right randomness to equivocate $\boldsymbol{Com}(m)$. The claim follows by a simply hybrid argument over all the signature queries. For a specific query j, by statistical witness indistinguishability, the distribution of the proof π_j is statistically close in the two hybrids.

Moreover, for any $(x, w) \in \mathcal{R}$, and for any $\mathsf{crs} \leftarrow\!\$\ \mathsf{Init}(1^\lambda)$, let $\Pi_w := \{\pi|\ \pi = \mathsf{P}(\mathsf{crs}, x, w; z)\}$. The statistical witness indistinguishability property implies:

$$\Pr_{\pi\leftarrow\!\$\ \mathsf{P}(\mathsf{crs},x,w)}[\pi \notin \Pi_{w'}] \in \mathtt{negl}(\lambda).$$

This is because 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, we can sample z' from the distribution $\{z|\ \pi = \mathsf{P}(\mathsf{crs}, x, w'; z)\}$ as the distribution is not empty.

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

The Predictor $\mathsf{P}^{\mathcal{O}_{\boldsymbol{\Delta}}(\cdot)}$. The predictor runs the same as the previous hybrid, with the difference that $\boldsymbol{\Delta}$ is not sampled by the predictor as part of the signing key, but can instead be accessed via $\mathcal{O}_{\boldsymbol{\Delta}}(\cdot)$. In particular, all signature queries are handled as in $\mathbf{H}_4$. Moreover, whenever the adversary A queries with a leakage oracle query f the predictor define $f'(\cdot) := f(\alpha(\cdot))$ and forwards it to its own leakage oracle. Finally the predictors receives from A the n forgeries (m^*_i, σ^*_i) and does as follow:

1. Check that all the forgeries are valid and that the messages are different, otherwise return $\perp$;
2. Parse σ^*_i as C^*_i, τ^*_i, π^*_i and compute $\boldsymbol{y}^*_i = \mathsf{Dec}(\mathsf{sk}^e, \tau^*_i, C^*_i)$;
3. For $j \in [\mu]$ sample a polynomial $\boldsymbol{\delta}^*_j$ in $\mathbb{Z}_p[X]$ of degree d such that $\boldsymbol{\delta}_j(m^*_i) = y^*_{i,j}$ for all $i \in [n]$;
4. Outputs $\boldsymbol{\Delta}^* = (\boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}^*_\mu)$.

Let View be the full view of the predictor P after have interacted with the leakage oracle and fully executed the experiment. In particular, View contains all the public parameters $\mathsf{crs}, \mathsf{pk}, \vartheta$, the trapdoors $\mathsf{sk}^e, \mathsf{sk}^s, \psi$, the randomness $(r'_j)_{i\in[d]}$, the auxiliary information and all the signatures $(aux_j, \sigma_j)_{j\in[q]}$, moreover it contains the leakage Leak performed by B.

Lemma 7. $\widetilde{\mathbb{H}}_\infty(\boldsymbol{\Delta} \mid \mathsf{View}) \geq (d \cdot \mu) \log p - \ell$.

Proof. The proof of the lemma is rather easy. Notice that $|\boldsymbol{\Delta}| \geq (d \cdot \mu) \log p$ moreover $\mathsf{crs}, \mathsf{pk}, \vartheta$, $\mathsf{sk}^e, \mathsf{sk}^s, \psi$ and $(r'_j)_{i\in[d]}, (aux_j, \sigma_j)_{i\in[q]}$ are independent of $\boldsymbol{\Delta}$. Lastly, the size of the leakage is ℓ so, by Lemma 3, it can decrease the average conditional min entropy only of ℓ bits.

Lemma 8. *If there exists a PPT adversary* A *such that* $\mathtt{Adv}^{\mathcal{SS},\mathsf{A}}_{one-more}(\lambda) = \varepsilon$ *then* $\Pr\left[\mathsf{P}^{\mathcal{O}_{\boldsymbol{\Delta}}(\cdot)} = \boldsymbol{\Delta}\right] \geq \exp(-((d-n)\cdot\mu)\log p) \cdot \varepsilon'_4$.

Proof. Conditioning on the event $\mathtt{Good}_4$ and by the correctness of the ABM-Enc scheme we have that for all $i \in [0, q]$ and $j \in [\mu]$ the equation $y^*_{i,j} = \boldsymbol{\delta}_j(m^*_i)$ holds. The predictor P in this case sample uniformly at random μ polynomials that evaluates as $\boldsymbol{\Delta}$ in those positions. Therefore, for any $j \in [\mu]$, the predictor guesses the right polynomial with probability $\exp(-((d-n))|\mathbb{Z}_p|)$ (because it has to guess only $d-n$ coefficients) in this conditional space.

We finish the proof by noticing that ε'_4 is equal to $\varepsilon - \texttt{negl}(\lambda)$ moreover, by Lemma 7:

$$-\log\left(\exp((n-d)\cdot\mu)\log p\right)\cdot(\varepsilon - \texttt{negl}(\lambda))) \geq (d\cdot\mu)\log p - \ell.$$

In fact, the lemma gives an upperbound on the guessing probability of $\boldsymbol{\Delta}$ given the view. By easy calculation we can derive that the following equation holds:

$$n\cdot\mu\log p - \log(\varepsilon - \texttt{negl}(\lambda)) \leq \ell$$

By setting the slack parameter $\gamma = s/(2\mu\lambda)$ and noticing that $n \geq \lfloor\frac{\ell}{\gamma\cdot s}\rfloor + 1$ and $\log p = \lambda$ then it must be $\varepsilon \in \texttt{negl}(\lambda)$ for the equation above to hold.