

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), defines 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 a slack parameter that depends on the number of signatures 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 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.

Keywords: signature scheme; leakage resilience; efficient scheme; knowledge assumptions.

Table of Contents

Efficient Fully-Leakage Resilient One-More Signature Schemes	1
<i>Antonio Faonio</i>	
1 Introduction.....	2
2 Notations and Preliminaries	6
2.1 All-but-Many Encryption.....	10
3 Fully-Leakage One-More Unforgeability	11
4 Signature scheme based on ABM-Encryption	12
5 A Signature Scheme based on KEA	18
6 Acknowledgements	26
A Information-Theoretic Lemmas	28
B Commitment Schemes	28
B.1 Security properties	28
C Quasi-Adaptive NIZK and NIWI argument systems	29
C.1 NIWI argument systems.....	29
C.2 Quasi-Adaptive NIZK argument systems	29
C.3 Constructions	30

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. [26,27,19]) which can break provably secure cryptosystems exploiting physical characteristics of the crypto-devices where such schemes are implemented.

A common way to model leakage attacks is to give to the adversary a leakage oracle. Such oracle stores the current secret state of the cryptosystem under attack (let it be α), takes as input leakage functions f_i and returns $f_i(\alpha)$.

The leakage functions need to belong to a restricted set of functions, as otherwise there is no hope for security. In this paper we consider the *bounded leakage model* where we assume that the total bit-length of the leakage obtained via the leakage functions is smaller than some a priori determined leakage bound ℓ . Leakage-resilient schemes in this model include public-key, identity-based encryption, signature schemes and identification schemes [29,3,2,8,10,7,28,4,13].

Graceful degradation. For any existentially unforgeable signature scheme in the bounded leakage model, necessarily, the length of a signature is larger than the leakage bound, as otherwise an adversary could simply leak a forgery. The main consequence is that, if the goal is to tolerate large amount of leakage then, the signature size needs to be very large but the latter makes the schemes unpractical. Recently Nielsen, Venturi and Zottarel [30] addressed this issue

introducing a new notion of security for signature schemes which requires that an adversary should not be able to produce more forgeries than what he could have leaked via leakage queries.

In particular, if s is the length in bits of a signature of size and ℓ is the leakage bound, to break unforgeability, an adversary must produce n forgeries where $n \approx \ell/(\gamma \cdot s) + 1$, where $\gamma \in (0, 1]$ is a value that we call the “slack parameter”. Roughly speaking, the slack parameter measures how close to optimal security the scheme is. When $\gamma = 1$ we say that the scheme has *optimal* graceful degradation of security, as the number of forged signatures requested is exactly one more than what an adversary could possibly leak. When γ is a constant smaller than 1 we say that the scheme has *almost-optimal* graceful degradation, as in this case, the number of forged signature requested is a constant factor more than what an adversary could leak¹. Notably, this new security notion enables to design signature schemes where the size of the secret key (and the leakage bound) does not depend on the signature size, leading to short signatures.

Subsequently, Faonio, Nielsen and Venturi [14] (journal version in [15]), extended the model to the fully-leakage resilient setting, where the adversary can leak arbitrary information of the entire secret state, including all the random coins of the signing algorithm.

Interestingly, while in the (not-fully) leakage-resilient regime the authors of [30] showed a signature scheme with almost-optimal graceful degradation, in the fully-leakage-resilient regime the best signature scheme known (in the standard model) has slack parameter $\gamma = O(1/q)$ 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 two new fully leakage-resilient signature schemes in the bounded leakage model where the slack parameter does not depend on the number of signatures issued.

The first signature scheme has slack parameter $O(1/\lambda)$. The construction makes use of an All-but-Many Encryption scheme (Fujisaki [17]) and a Non-Interactive Witness-Indistinguishable system and is instantiated under standard number theoretic assumptions.

The second signature scheme has optimal graceful degradation. The construction is based on a specific extractable and perfectly hiding commitment scheme (Abe and Fehr [1]) and on a quasi-adaptive NIZK for linear space (Jutla and Roy [24]). For technical reason, we need a NIZK system with a weak form of knowledge soundness. As minor contribution of independent interest, we show how to modify the elegant construction of Kiltz and Wee [25] to get an efficient quasi-adaptive NIZK system for linear-space relationship with (weak) knowledge soundness. Both the components of the second schemes are instantiated under the knowledge of the exponent assumption (see, e.g. [9,5,1,21]).

¹ In [30], the authors show that the notion, even for small value of the slack parameter, allows for interesting applications such as leakage-resilient identification schemes.

A Technical Overview. We recall the scheme of [30], for future reference we call it NVZ14. The secret key of NVZ14 is a polynomial δ in $\mathbb{Z}_p[X]$ of degree d and a signature for a message $m \in \mathbb{Z}_p$ is composed by a commitment C^* to the evaluation of the polynomial δ on the point m together with a sim-extractable NIZK that the commitment, indeed, commits to such evaluation. The polynomial δ is published in the verification key using an homomorphic commitment scheme (for example, the classical Pedersen’s commitment scheme [32]). The verification of a signature works in two stages: first from the verification key it derives (using the homomorphic property of the commitment scheme) a commitment C_m to the evaluation of the polynomial δ on point m , second it verifies the NIZK for the statement (C^*, C_m) which proves that the commitments C_m and C^* open to the same value, therefore proving that the commitment C^* commits to an evaluation of δ on the point m . The leakage bound of the scheme is roughly $\ell \approx d \log p$ and the slack parameter is a constant. The key idea for the unforgeability is that from $n \approx d + 1$ signatures we can extract $d + 1$ evaluations of the polynomial δ , however, because of the bound on the leakage performed, at most d evaluation points could be possibly be uniquely defined. The latter implies that one of the commitment produced by the adversary can be opened in two different way therefore breaking the binding property of the commitment scheme.

The construction proposed by [14] follows the same blue print. Their main idea is to convert leakage functions over the full state (namely, the secret key and the randomness) to leakage functions of the secret key only. In this way, they reduce the task of proving fully leakage resilient to the easier task of proving (standard) leakage resilience.

We give a glimpse of their technique with a toy example. As in the scheme NVZ14, in the construction of [14], a signature $\sigma = (C^*, \pi)$ is composed by a commitment C^* and a proof of consistency for the commitment π . So the randomness of a signature is equal to (r, t) where r is the randomness for the commitment and t is the randomness for the NIZK. Their first idea is to use an equivocable commitment scheme. Recall that a commitment scheme is equivocable if, roughly speaking, we can sample a fake commitment C such that, given a trapdoor, for any message m we can produce randomness r' such that $C = \text{Com}(\text{vk}, m, r')$, namely, the fake commitment C opens to the message m . For the sake of this toy example, let us consider a leakage function $f(\delta, r)$ that does not depend on the randomness t of the NIZK. In [14], the authors show that we can construct a new leakage function $\hat{f}(\delta)$ that first computes r' equivocating the commitment C^* to $\delta(m)$ and then it computes $f(\delta, r')$. The function \hat{f} converts the leakage on the randomness as leakage of the secret key only.

The main technical problem that [14] had to solve is that standard equivocable commitments scheme were not sufficient. In fact two contrasting requirements are necessary: on one hand, both the commitment scheme and the NIZK need to be equivocable (so that we can reduce fully-leakage resilience to standard leakage resilience as shown in the toy example above), on the other hand, to extract the n evaluations of the polynomial δ we need that either the commitment scheme or the NIZK system is perfectly binding. To solve this problem

the authors of [14] showed a construction of a commitment scheme where any commitment created is perfectly binding with probability $1/q$ and equivocal with probability $1 - 1/q$. 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 signatures (so that $\gamma = O(1/q)$) strictly more than $(\ell/\log p) + 1$ signatures are perfectly binding (with overwhelming probability). The unforgeability of the scheme follows because a winning adversary gets in input exactly ℓ bits of information about δ and outputs strictly more than ℓ bits of information about δ : this adversary cannot exist as otherwise a basic information-theoretic principle would be violated.

New Ideas. We describe our two new signature schemes. For the first construction we substitute the commitment scheme of [14] with an All-But-Many Encryption (ABM-Enc) scheme. Roughly speaking, an ABM-Enc is an encryption scheme where all the ciphertexts created by the adversary can be successfully decrypted (knowing the secret key) while, with the knowledge of a special trapdoor, we can create an unbounded number of fake ciphertexts that are equivocal. The proof of security is quite straight-forward (actually even easier than in [14]): 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 [17], building over a paper of Hofheinz [23], 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 (where λ is the security parameter). 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 parameter, the problem is that efficient NIZK [22] and the Fujisaki’s 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 $1/\text{poly}(\lambda)$. Luckily, the ABM-Enc based on DDH of Fujisaki fits better with the NIZK of Groth-Sahai, as to prove the necessary statement we need only a constant number, in the size of the ciphertext, of pairing-product equations.

The second construction is inspired by the following observation: if we used a zk-SNARK [21,20,31] 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 a zk-SNARK is an over killing, as what we need is the ability of simultaneously equivocate and extract the commitments, and in particular, we do not need succinctness. Therefore, instead of naively use zk-SNARKs, we “open the box” of zk-SNARKs. In particular, we consider the commitment scheme of Abe and Fehr [1] based on the knowledge of the exponent assumption (KEA3) of Bellare and Palacio [5] (see also Damgaard [9]). Nicely, for this kind of commitments, we can reduce the relation that two commitments open to the same message to the fact that a certain vector in \mathbb{G}^2

lies in a specific subspace. The latter allows us to get faster and shorter signatures, thanks to recent advances in efficiency of quasi-adaptive NIZK systems for linear realations (see for example, [24,25]).

More in details, the proof technique for the second construction diverges significantly to the proof technique of [14]. The main reason is that the commitment scheme of Abe and Fehr is simultaneously extractable and perfect hiding but it is not efficiently equivocable². Our strategy is to first apply all the computational steps and then use the fact that the commitment scheme is perfectly hiding. Therefore we can “equivocate” a commitment by brute force it and open it to the desired value.

Comparison. We compare our signature schemes with the signature schemes of [30] and [14,15] (see Table 1). Four different signature schemes are presented in [15], we select the three most interesting³ and we denote them with FNV15₁, FNV15₂ and FNV15₃. The third column in the Table 1 (namely, “No Erasure”) refers to a weaker model of fully leakage resilient signature considered in [14]. Specifically, the scheme FNV15₁ is proved secure under the assumption that the cryptographic device can perfectly erase the random coins used in the previous invocations. We call \mathcal{SS}_1 the signature scheme based on ABM-Enc scheme 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 FNV15₂ but achieves asymptotically better graceful degradation. On the other hand, \mathcal{SS}_1 is both less efficient and with worse graceful degradation respect to FNV15₁ and FNV15₃, however, FNV15₁ needs perfect erasure of the randomness and FNV15₃ 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 the commitment scheme based on the knowledge of the exponent assumption does not admit oblivious sampling of the parameters. The scheme \mathcal{SS}_2 achieves optimal graceful degradation, the signature size is independent of the ϵ and, notably, more compact (both asymptotically and practically) even than the signature scheme FNV15₃ 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 $f \in \text{negl}(\lambda)$, if it vanishes asymptotically faster than the inverse of any polynomial. We use the classic notion of probabilistic polynomial time (PPT) algorithms. We write $x \leftarrow \$ \mathcal{D}$ (resp. $x \leftarrow \$ A(y)$) to denote that x is chosen at random from the

² Intuitively, any trapdoor for equivocation would break the knowledge of the exponent assumption.

³ As the forth scheme is a variation of FNV15₁ and it achieves worse efficiency parameters

Scheme	Fully	No Erasure	KGen	G. D.	Assumption	Efficiency	
						leak	signature size
NVZ14	✗	-	-	$O(1)$	DLIN	$\frac{1}{2} - o(1)$	$O(1)$
FNV15 ₁	✓	✗	✓	$O(1)$	DLIN	$1 - \epsilon$	$O(\epsilon^{-1})$
FNV15 ₂	✓	✓	✓	$O(1/q)$	DLIN	$1 - \epsilon$	$O(\epsilon^{-1} \cdot \log \lambda)$
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 decisional diffie-hellman assumption.

distribution \mathcal{D} (resp. an PPT algorithm A run on input y), and we write $x \leftarrow A(y; r)$ to denote that we assign to x the output of A run with randomness r . 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 \mathbf{v} of length d , we denote $\mathbf{v}(m) := \mathbf{v}^T \cdot (1, m^1, \dots, m^{d-1})^T$, meaning the evaluation of the polynomial with coefficients \mathbf{v} at point m . We consider also the natural extension of the notion to matrix, $\mathbf{V}(m) := \mathbf{V} \cdot (1, m^1, \dots, m^{d-1})^T$. All the algorithms take as input (group) parameters prm , for readability, whenever it is clear from the context we consider them implicit. 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 $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. We use additive notation for the group operation, and we denote group elements using the bracket notation introduced by Escala *et al.* in [12]. 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 write $[x \cdot y]_T$ as shorthand for $e([x]_1, [y]_2)$. We recall the standard notion of collision-resistant hash function (CRH). A tuple of PPT (Gen_{CRH}, H) is a CRH where Gen_{CRH} upon security parameter 1^λ produces a 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 m_0, m_1 chosen as function of hk , such that $H(hk, m_0) = H(hk, m_1)$.

Knowledge of the Exponent Assumption. Consider the experiment in Fig. 1 between an adversary A , a randomness sampler \mathcal{S} , an extractor Ext and a bilinear group generator Setup_{BG} .

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 deterministic PT A*

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

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

Fig. 1: The experiment of the q -KE* assumption.

and any PPT sampler \mathcal{S} there exists a PT Ext such that:

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

In contrast with the standard definition of the knowledge of the exponent assumption, in our definition we additionally have a sampler \mathcal{S} . The technical reason is that we deal with adversaries with oracle access (for example, to the signature oracle or the leakage oracle). In fact, in this setting, as shown by Fiore and Nitulescu [16], 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. Notice, in bilinear groups the test $[\mathbf{y}]_1 \in \text{Span}([1, \alpha]_1^T)$ can be efficiently performed using the bilinear map $e([\mathbf{y}_0]_1, [\alpha]_2) = e([\mathbf{y}_1]_1, [1]_2)$. Also, we can naturally scale down the assumption to non bilinear groups, in this case, the adversary does not get $[\alpha]_2$. Given a (non-bilinear) group generator Setup_G the assumption for $q = 1$ is not stronger than the KEA [9] 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 [21], indeed it is easy to show that if q -PKE holds then 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 $\mathbf{g}, \alpha \in \mathbb{Z}_p$ on clear⁴.

Kernel Diffie-Hellman Assumptions. Given parameter 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. [12]). Given a bilinear group generator Setup_{BG} , we say that the \mathcal{D}_k -Kernel Diffie-Hellman assumption (\mathcal{D}_k -KerMDH) holds for

⁴ Also notice that we quantify the extractor after the sampler, so to avoid pathological situation where the adversary \mathbf{A} simply forwards the output of the sampler \mathcal{S} .

Setup_{BG} if for any PPT A :

$$\text{Adv}_{A, \text{Setup}_{BG}}^{\mathcal{D}'_k\text{-KerMDH}}(\lambda) := \Pr[\mathbf{c}^T \cdot \mathbf{A} = 0 \wedge \mathbf{c} \neq \mathbf{0} \mid [\mathbf{c}]_1 \leftarrow A(\text{prm}, [\mathbf{A}]_2)] \in \text{negl}(\lambda).$$

where $\mathbf{A} \leftarrow \mathcal{D}_k$ and $\text{prm} \leftarrow \text{Setup}_{BG}(1^\lambda)$.

More specifically, in this paper we consider the KerLin_2 assumption which is equivalent to the $\mathcal{D}'_k\text{-KerMDH}$ where \mathcal{D}'_k outputs a matrix which columns are $(1, a_1, 0)^T$ and $(1, 0, a_2)^T$ for a_1, a_2 uniformly chosen in \mathbb{Z}_p .

Homomorphic Trapdoor Commitment Schemes. A trapdoor commitment scheme $\mathcal{COM} = (\text{Setup}, \text{Com}, \text{ECom}, \text{EOpen})$ is a tuple of algorithms where: (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 $Com = \text{Com}(\vartheta, m; r)$. (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 to any message in a indistinguishable way. The latter instead says that no PPT adversary can open the same commitment to two different messages without the knowledge of the trapdoor ψ . 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* if given commitments Com and Com' (that commit to m and m') and $a \in \mathbb{Z}_p$, one can compute the commitment $Com^* := a \cdot Com + Com'$ that opens to $a \cdot m + m'$. We write the mappings as $Com^* = \text{Com}(\vartheta, a \cdot m + m'; a \cdot r + r')$.

Moreover, we require the following additional property. Let $(\vartheta, \psi) \leftarrow \text{Setup}(1^\lambda)$, $(Com_1, aux_1) \leftarrow \text{ECom}(\vartheta, \psi)$ and $(Com_2, aux_2) \leftarrow \text{ECom}(\vartheta, \psi)$. We can use the auxiliary information $a \cdot aux_1 + aux_2$ to equivocate the commitment $a \cdot Com_1 + Com_2$. Finally, we consider commitment schemes with an additional algorithm Setup which samples the verification key *obliviously*.

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 non-interactive argument system $\mathcal{NIZK} := (\text{Init}, \text{P}, \text{V})$ for \mathcal{R} is a tuple of PPT algorithms where: (1) The initialization algorithm Init takes as input the security parameter 1^λ , and creates a common reference string (CRS) $\text{crs} \in \{0, 1\}^*$; (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 $\pi \leftarrow \text{P}(\text{crs}, x, w)$; (3) The verifier algorithm V takes as input the CRS crs , a pair (x, π) , and outputs a decision bit $\text{V}(\text{crs}, x, \pi)$. Additionally, we say that an 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). We consider distribution $\mathcal{D}_{\mathcal{R}}$ over NP-relation. As for all the algorithms in this paper, the distribution can depends on the parameters prm (for example,

prm could be the description of a bilinear group). We require the standard notion of completeness, meaning that for any CRS crs output by $\text{Init}(1^\lambda)$ (or for any \mathcal{R} and any crs output by $\text{Init}(1^\lambda, \mathcal{R})$ in the quasi-adaptive case), and for any pair $(x, w) \in \mathcal{R}$, we have that $V(\text{crs}, x, P(\text{crs}, x, w)) = 1$ with all but a negligible probability. We consider argument systems that admit oblivious sampling of the CRS and we denote it with Init . We require the following security properties (cf. Appendix C).

- **Perfect zero-knowledge:** 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.
- **Perfect witness-indistinguishability:** Given two different witnesses valid for the same instance, a proof generated with the first witness is equivalently distributed to a proof generated with the second witness.
- **Adaptive weak knowledge soundness:** For any PPT adversary that on input the CRS produces a valid NIZK proof for a statement x there exists a PPT extractor that outputs a witness w such that $(x, w) \in \mathcal{R}$.
- **Adaptive soundness:** No PPT adversary can forge a verifying proof for an adaptively chosen invalid statement.

2.1 All-but-Many Encryption

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

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

$$\text{Adv}_{\mathcal{ABM}}^{\text{pprf}}(\lambda) := \left| \Pr \left[(\text{pk}, \text{sk}^e, \text{sk}^s) \leftarrow \text{Gen}(1^\lambda) \right] \Pr \left[A(\text{pk})_{\text{Sample}(\text{pk}, \text{sk}^e, \cdot)} = 1 \right] - \Pr \left[(\text{pk}, \text{sk}^e, \text{sk}^s) \leftarrow \text{Gen}(1^\lambda) \right] \Pr \left[A(\text{pk})_{\mathcal{O}_{\text{pk}}(\cdot)} = 1 \right] \right|$$

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

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

$$\text{Adv}_{\mathcal{ABM}}^{\text{unf}}(\lambda) := \Pr \left[(t^*, u^*) \in \mathcal{L}^e, t^* \notin \mathcal{Q} : \begin{array}{l} (\text{pk}, \text{sk}^e, \text{sk}^s) \leftarrow \text{Gen}(1^\lambda) \\ (t^*, u^*) \leftarrow A(\text{pk})_{\text{Sample}(\text{pk}, \text{sk}^e, \cdot)} \end{array} \right]$$

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

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

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

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

$$\left\{ (u, C, r) : \begin{array}{l} u \leftarrow \$ \text{Sample}(\text{pk}, \text{sk}^e, t), \\ r \leftarrow \$ \{0, 1\}^\lambda, \\ C \leftarrow \text{Enc}(\text{pk}, \tau, \mu; r) \end{array} \right\}_{\lambda \in \mathbb{N}},$$

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

Theorem 1 (Fujisaki, [18]). *If DDH assumption holds in Setup_G then there exists an ABM-Enc scheme. Moreover, the scheme admits an algorithm $\tilde{\text{Gen}}$ that obviously samples the public parameter.*

3 Fully-Leakage One-More Unforgeability

A signature scheme is a triple of algorithms $\mathcal{SS} = (\text{Gen}, \text{Sign}, \text{Verify})$ where: (1) The key generation algorithm takes as input the security parameter λ and outputs a verification key/signing key pair $(\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$; (2) The signing algorithm takes as input a message $m \in \mathcal{M}$ and the signing key sk and outputs a signature $\sigma \leftarrow \text{Sign}(\text{sk}, m)$; (3) The verification algorithm takes as input the verification key vk and a pair (m, σ) and outputs a bit $\text{Verify}(\text{vk}, (m, \sigma)) \in \{0, 1\}$.

Given a signature scheme \mathcal{SS} , consider the experiments 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 $\mathcal{SS} = (\text{Gen}, \text{Sign}, \text{Verify})$ is (ℓ, γ) -fully-leakage one-more unforgeable if for every PPT adversary A we have that:*

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

Moreover, We say that \mathcal{SS} is (ℓ, γ) -fully-leakage one-more unforgeable with leak-free keygen if for every PPT adversary A we have that:

$$\text{Adv}_{\mathcal{SS}, A}^{\text{one-more}^*}(\lambda, \ell, \gamma) := \Pr [\text{Exp}_{\mathcal{SS}, A}^{\text{one-more}^*}(\lambda, \ell, \gamma) = 1] \in \text{negl}(\lambda).$$

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

1. $(\mathbf{vk}, \mathbf{sk}) \leftarrow \mathfrak{s} \text{Gen}(1^\lambda; r_0)$, return \mathbf{vk} to \mathbf{A} ; let $\alpha = r_0$, $\text{let } \alpha = \mathbf{sk}$.
2. Run $\mathbf{A}(\mathbf{vk})$ with oracle access to $\text{Sign}(\mathbf{sk}, \cdot)$ and the leakage oracle.
 - Upon query $m \in \mathcal{M}$ to the signature oracle, let $\sigma := \text{Sign}(\mathbf{sk}, m; r)$, $r \leftarrow \mathfrak{s} \{0, 1\}^\lambda$ and update 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 \mathbf{A} , and let $A \in \{0, 1\}^*$ be the concatenation of all the leakage. \mathbf{A} outputs n pairs $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$.
4. The experiment outputs 1 if and only if the following conditions are satisfied:
 - (a) $\text{Verify}(\mathbf{vk}, (m_i^*, \sigma_i^*)) = 1$ and $m_i^* \notin \mathcal{Q}$, for all $i \in [n]$.
 - (b) The messages m_1^*, \dots, m_n^* are pairwise distinct.
 - (c) $n \geq \lceil \ell / (\gamma \cdot s) \rceil + 1$, where $s := |\sigma|$ and $|A| \leq \ell$.

Fig. 2: The fully-leakage one-more unforgeability experiment and the fully-leakage one-more unforgeability experiment with leak-free key gen. The second experiment is equal to the first but it additionally executes the operations described the box.

The number of signatures the adversary musts 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 the signature scheme \mathcal{SS} is to the optimal security \mathcal{SS} . In particular, in the case $\gamma = 1$ one-more unforgeability requires that the adversary \mathbf{A} cannot forge even a single signature more than what it could have (partially) 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-Encryption

Our scheme $\mathcal{SS} = (\text{Gen}, \text{Sign}, \text{Verify})$ 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{ABM} , a NIWI argument system \mathcal{NIWI} and a CRH function $(\text{Gen}_{\text{CRH}}, \text{H})$. The scheme follows the basic template described in Sec. 1, however instead of using just one single polynomial $\delta \in \mathbb{Z}_p[X]$ of degree d , we use $\mu \in \mathbb{N}$ different polynomials arranged in the matrix Δ . The correctness follows from the completeness of the NIWI argument system, and from the linearly homomorphic property.

Theorem 2. *Let $\mu \in \mathbb{N}$. Assume that: (i) the commitment scheme \mathcal{COM} is trapdoor hiding and linearly homomorphic with message space \mathbb{Z}_p^μ ; (ii) the ABM is a secure ABME-Enc scheme with message space \mathbb{Z}_p^μ and ciphertexts of length s_1 bits; (iii) \mathcal{NIWI} is a perfect NIWI argument system for the relation \mathcal{R} described in Fig. 3 with proofs of length s_2 bits. Then, let $s = s_1 + s_2$ and let*

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

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

Run $hk \leftarrow \text{Gen}_{CRH}(1^\lambda)$, $\text{crs} \leftarrow \tilde{\text{Init}}(1^\lambda)$, $\vartheta \leftarrow \tilde{\text{Setup}}(1^\lambda)$ and $\text{pk} \leftarrow \text{Gen}(1^\lambda)$.
 Sample $\Delta \leftarrow \mathbb{Z}_p^{\mu, d+1}$ and $\mathbf{r} = (r_0, \dots, r_d) \leftarrow \mathbb{R}^{d+1}$, and compute $\text{Com}_i \leftarrow \text{Com}(\vartheta, \delta_i; r_i)$ for $i \in [0, d]$, where $\delta_i \in \mathbb{Z}_p^\mu$ is the i -th column of Δ . Let $\mathbf{Com} = (\text{Com}_0, \dots, \text{Com}_d)$
 Output

$$\text{sk} = (\Delta, \mathbf{r}) \quad \text{vk} = (\text{crs}, \vartheta, \text{pk}, \mathbf{Com}).$$

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

Verification. Given a pair (m, σ) and vk , parse σ as $(C, \tau = (t, u), \pi)$ and parse vk as $(\text{crs}, \vartheta, \text{pk}, \mathbf{Com})$. Output 1 if and only if $\text{H}(hk, m) = t$ and $\text{V}(\text{crs}, \pi, (\vartheta, \text{pk}, \tau, \mathbf{Com}(m), C))$.

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

$\gamma = \mu \log p/s$, for any $0 \leq \ell \leq ((d+1)\mu \log p) - \lambda$, the signature scheme \mathcal{SS}_1 is (ℓ, γ) -fully-leakage one-more unforgeable.

We first provide a sketch, the formal proof follows below. The proof is similar to the proof of [14], the following proof sketch highlights the main differences.

Proof Sketch. We denote with $(r_0, \Delta, \mathbf{r}, (s_j, t_j)_{j \in [q]})$ the full secret state. Notice that, because of the oblivious sampling of the parameters, the randomness r_0 such that $\text{vk}, \text{sk} = \text{KGen}(1^\lambda; r_0)$ can be computed efficiently as function of both vk and sk , we therefore omit r_0 from the state α . The first hybrid \mathbf{H}_0 is the fully-leakage one-more unforgeability game but we additionally condition on the validity of the forged proofs. By the adaptive soundness of the NIWI the real experiment and \mathbf{H}_0 are indistinguishable. In the next hybrid \mathbf{H}_1 we switch the way the parameters are sampled, so that we get the secret keys sk^s, sk^e of the ABM-Enc and the equivocation trapdoor ψ of the commitment scheme. The hybrids \mathbf{H}_0 and \mathbf{H}_1 are indistinguishable because of the dual mode property of the ABM-Enc and the equivocability of the commitment scheme.

In the hybrid \mathbf{H}_2 we equivocate the commitments \mathbf{Com} in the public key. Notice that the full secret state α can be written as $((\Delta, \mathbf{r}(\Delta)), (s_j, z_j)_{j \in [q]})$ where $\mathbf{r}(\Delta)$ is a function of the secret key computed by EOpen.

In the hybrid \mathbf{H}_3 for each signature oracle query we sample the tag $\tau = (t, u)$ such that $u = \text{Sample}(\text{pk}, \text{sk}^s, t)$. The indistinguishability comes from the pseudorandomness property of the ABM-Enc scheme.

Thanks to the last change, in the hybrid \mathbf{H}_4 , for each signature oracle query we can sample the encryption C using the trapdoor mode FakeEnc. Notice that

the full secret state α can be written as $((\Delta, \mathbf{r}(\Delta)), (s_j(\Delta), z_j)_{i \in [q]})$ where for any j , the value $s_j(\Delta)$ is a function of the secret key Δ computed using the algorithm **EquivEnc**. The dual mode property of the ABM-Enc scheme assures that the two hybrids are indistinguishable.

In the hybrid \mathbf{H}_5 we compute the NIWI proof using the witness $(0, r', s')$ where r' is an opening of the equivocated commitment $\mathbf{Com}(m)$ to 0 and s' is an opening of the fake encryption to 0. This step follows exactly as in the proof of security in [14].

In this last hybrid the full secret state α can be written as $((\Delta, \mathbf{r}(\Delta)), (s_j(\Delta), z_j(\Delta))_{i \in [q]})$, namely, all the state can be written as a deterministic function of the secret polynomials Δ . In particular, any function $f(\alpha)$ could be rephrased as a function $f'(\Delta)$.

The last part of the proof proceeds similarly as in [14] so here we give just an intuition. Informally, an adversary A that wins the fully-leakage one-more unforgeability game with probability ϵ will win with probability negligibly close to ϵ in the hybrid \mathbf{H}_5 . Recall that a winning adversary returns $n := \lceil \ell / \mu \log p \rceil + 1$ valid signatures. By the unforgeability of the ABM Encryption and the change introduced in \mathbf{H}_0 , from the forged signatures $(m_i^*, \sigma_i^* = (C_i^*, \tau_i^*, \pi_i^*))_{i \in [n]}$, by decrypting the ciphertext C_i^* , we can extract the values $\Delta(m_i^*)$. Notice that each $\Delta(m_i^*)$ gives us $\mu \log p$ bits of information about Δ . Putting all together, with probability negligibly close to ϵ from the adversary we can extract $n \cdot (\mu \log p) > \ell$ bits of information about Δ . On the other hand, in \mathbf{H}_5 , the adversary gets at most ℓ bits of information about Δ , the latter implies that ϵ must be negligible.

Formal Proof.

Proof. Let A be an adversary such that $\text{Adv}_{\mathbf{A}, \mathbf{SS}_1}^{\text{one-more}}(\lambda, \ell, \gamma) = \epsilon$ for parameter ℓ, γ as described in the statement of the theorem. Let $\mathbf{H}_0(\lambda)$ be the experiment $\text{Exp}_{\mathbf{SS}, \mathbf{A}}^{\text{one-more}}(\lambda)$. Denote with $((m_1^*, \sigma_1^* = (C_1^*, \tau_1^*, \pi_1^*)), \dots, (m_n^*, \sigma_n^* = (C_n^*, \tau_n^*, \pi_n^*)))$ the list of forgeries of A . Let Forge_0 be the event that \mathbf{H}_0 returns 1, so that $\mathbb{P}[\text{Forge}_0] = \epsilon$. Define False_0 to be the event that at least one of the proofs contained in the adversary's forgeries is relative to a false statement, i.e., False_0 is verified if in \mathbf{H}_0 there exists $i \in [n]$ for which $\text{Dec}(\text{sk}^e, \tau_i^*, C_i^*) = m'_i$ and $m'_i \neq \Delta(m_i^*)$. Define Collision_0 to be the event that there exists $i \in [n]$ and $j \in [q]$, for which $\text{H}(hk, m_i^*) = \text{H}(hk, m_j)$ where m_j is the j -th signature oracle's query made by the adversary.

Let $\epsilon_0 := \mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0 \wedge \neg \text{Collision}_0]$.

Claim. $\epsilon - \epsilon_0 \in \text{negl}(\lambda)$.

Proof. The claim is proved in two steps, first we prove $|\epsilon - \mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0]| \in \text{negl}(\lambda)$ and then we prove $|\mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0] - \mathbb{P}[\text{Forge}_0 \wedge \neg \text{False}_0 \wedge \neg \text{Collision}_0]| \in \text{negl}(\lambda)$. By adaptive computational soundness of the NIWI argument system, we must have that $\mathbb{P}[\text{False}_0] \in \text{negl}(\lambda)$. In fact, from an adversary A provoking False_0 , we can easily construct an adversary breaking adaptive soundness by simply emulating the entire experiment for A and outputting 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 $\mathbf{Forge}_i \wedge \neg \mathbf{False}_i \wedge \neg \mathbf{Collision}_i$. During the experiment the adversary has leakage oracle access to $\alpha = (\mathbf{\Delta}, \mathbf{r}, (s_j, z_j)_{j \in [q]})$ where s_j is the randomness used by \mathbf{Enc} and z_j is the randomness used by \mathbf{P} (the proving algorithm of the NIWI). Notice that, because of the oblivious sampling of the parameters, the randomness r_0 such that $\mathbf{vk} = \mathbf{KGen}(1^\lambda; r_0)$ can be computed efficiently as function of the verification key \mathbf{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 $(\mathbf{pk}, (\mathbf{sk}^s, \mathbf{sk}^e)) \leftarrow \mathfrak{S} \mathbf{Gen}(1^\lambda)$ and $\vartheta, \psi \leftarrow \mathfrak{S} \mathbf{Setup}(1^\lambda)$ and $\mathbf{crs} \leftarrow \mathfrak{S} \mathbf{Init}(1^\lambda)$.

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

The claim follows by the oblivious sampling properties of $\tilde{\mathbf{Setup}}, \tilde{\mathbf{Gen}}$ and $\tilde{\mathbf{Init}}$. 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 δ_i are replaced by equivocable commitments, i.e. $(Com_i, r'_i) \leftarrow \mathfrak{S} \mathbf{ECom}_1(\vartheta, \psi)$ 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 δ_i and the fake randomness r'_i , as $r_i(\mathbf{\Delta}) := \mathbf{EOpen}(\psi, \delta_i, r'_i)$. Given r_i the signature computation is identical. We write $\mathbf{r}(\mathbf{\Delta}) = (r_0(\mathbf{\Delta}), \dots, r_d(\mathbf{\Delta}))$ the vector of randomness \mathbf{r} computed as function of $\mathbf{\Delta}$ (with the help of ψ and (r'_0, \dots, r'_d)).

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

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 $u_j \leftarrow \mathfrak{S} \mathbf{Sample}(\mathbf{pk}, \mathbf{sk}^s, t_j)$ where $t_j = \mathbf{H}(hk, m)$.

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

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

Let $\mathbf{ForgeTag}_i$ 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 \mathbf{Good}_i defined as follow:

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

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

Proof. Notice that $\varepsilon'_3 - \varepsilon_3 \leq \Pr[\text{ForgeTag}_i]$, so we bound the probability of this event. Consider the adversary \mathbf{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 $\text{Sample}(\text{pk}, \text{sk}^s, \cdot)$. Eventually, the adversary \mathbf{A} outputs n forgeries and if all of them are valid and for different messages then the adversary \mathbf{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 \mathbf{B} wins with probability $\Pr[\text{ForgeTag}_i] / n$ 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 $C_j, \text{aux} \leftarrow \text{FakeEnc}(\text{pk}, (t_j, u_j), \text{sk}^s)$ and compute the randomness for the leakage oracle as $s_j \leftarrow \text{EquivEnc}(\tau, \text{aux}, \Delta(m))$. To stress that the randomness s_j can be computed as function of Δ , we write $s_j(\Delta) := \text{EquivEnc}(\tau, \text{aux}_j, \Delta(m_j))$ where m_j is the message queried at the j -th signature oracle call.

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

Proof. The dual mode property of the ABM-Enc scheme implies that, for any $t_j \in \{0, 1\}^\lambda$ and $u_j \leftarrow \text{Sample}(\text{pk}, \text{sk}^s, t_j)$ the distribution that computes the ciphertext C_j using Enc and the distribution that compute it with 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

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

where $\mathbf{r}'(m) = \sum_{i=0}^d r'_i \cdot m^i$ 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 $w := (\Delta(m), \text{EOpen}(\psi, \Delta(m), \mathbf{r}'(m)), \text{EquivEnc}(\tau, \text{aux}, \Delta(m)))$ and z'_j . In particular, $z_j(\Delta) := z'_j$ where z'_j is sampled from the distribution $\{z : \pi_j = \text{P}(\text{crs}, (\vartheta, \text{pk}, \tau_j, \mathbf{Com}(m), C_j), (0, r'_j, s'_j))\}$. The state used to answer the leakage query f is set to:

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

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

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

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

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

$$\Pr_{\pi \leftarrow \text{s P}(\text{crs}, x, w)} [\pi \notin \Pi_{w'}] = 0.$$

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 $\{z \mid \pi = \text{P}(\text{crs}, x, w'; z)\}$ as the distribution is not empty.

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

The Predictor $\text{P}^{\mathcal{O}_\Delta(\cdot)}$. The predictor runs the same as the previous hybrid, with the difference that Δ is not sampled by the predictor as part of the signing key, but can instead be accessed via $\mathcal{O}_\Delta(\cdot)$. In particular, all signature queries are handled as in \mathbf{H}_4 . Moreover, whenever the adversary \mathbf{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 \mathbf{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 $\mathbf{y}_i^* = \text{Dec}(\text{sk}^e, \tau_i^*, C_i^*)$;
3. For $j \in [\mu]$ sample a polynomial δ_j^* in $\mathbb{Z}_p[X]$ of degree d such that $\delta_j(m_i^*) = y_{i,j}^*$ for all $i \in [n]$;
4. Outputs $\Delta^* = (\delta_1^*, \dots, \delta_\mu^*)$.

Lemma 1. $\Pr [\text{P}^{\mathcal{O}_\Delta(\cdot)} = \Delta] \leq \exp((d+1)\mu \log p - \ell)$.

Proof. Notice that $|\Delta| = ((d+1)\mu) \log p$ so, without any extra information, the guessing probability of P is bound to $\exp(-(d+1)\mu \log p)$. On the other hand, the size of the leakage is ℓ bits so, by Lemma 5, the guessing probability of P can increase at most of a multiplicative factor of 2^ℓ .

Lemma 2. *If there exists a PPT adversary \mathbf{A} such that $\text{Adv}_{\text{one-more}}^{\text{SS}, \mathbf{A}}(\lambda) = \varepsilon$ then $\Pr [\text{P}^{\mathcal{O}_\Delta(\cdot)} = \Delta] \geq \exp(-((d-n) \cdot \mu) \log p) \cdot \varepsilon_4'$.*

Proof. Conditioning on the event 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 $y_{i,j}^* = \delta_j(m_i^*)$ holds. The predictor P in this case sample uniformly at random μ polynomials that evaluates as Δ 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 $\varepsilon - \mathbf{negl}(\lambda)$ moreover, putting together the bounds of Lemma 1 and Lemma 2, and by taking the logarithms:

$$(n - d)\mu \log p + \log(\varepsilon - \mathbf{negl}(\lambda)) \leq -(d + 1)\mu \log p + \ell.$$

By easy calculation we can derive that the following equation holds:

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

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 $\varepsilon \in \mathbf{negl}(\lambda)$ for the equation above to hold.

Concrete Instantiation. We instantiate the ABM-Scheme with the construction \mathcal{ABM}_{DDH} of [18] based on DDH assumption, the NIWI argument system with Groth-Sahai [22] and the trapdoor commitment with the Pedersen’s commitment scheme. A ciphertext C of \mathcal{ABM}_{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 $n = \mathbf{poly}(\lambda)$ 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{ABM}_{DDH} we need to use the external Diffie-Hellman (SXDH) assumption.

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

Setup. The algorithm \mathbf{Setup} parses \mathbf{prm} as $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, G_1, G_2, G_T)$, picks at a random $[\mathbf{g}]_1 \leftarrow \mathfrak{s} \mathbb{G}_1^\mu$, $\alpha \leftarrow \mathfrak{s} \mathbb{Z}_p$ and $[h]_1 \leftarrow \mathfrak{s} \mathbb{G}_1$, sets $[\mathcal{M}]_1 \leftarrow (1, \alpha)^T \cdot [\mathbf{g}^T, h]_1$, sets $[h]_1 = [h, \alpha \cdot h]_1^T$ be the last column of $[\mathcal{M}]_1$, and sets $[\alpha]_2$. It outputs the verification key $\vartheta = ([\mathcal{M}]_1, [\alpha]_2) \in (\mathbb{G}_1^{2, \mu+1} \times \mathbb{G}_2)$.

Commit. The algorithm \mathbf{Com} on input $[\mathcal{M}]_1, [\alpha]_2$ and a message $\mathbf{m} \in \mathbb{Z}_p^\mu$, samples $r \leftarrow \mathfrak{s} \mathbb{Z}_p$ and sets $\mathbf{Com} = [\mathcal{M}] \cdot (\mathbf{m}^T, r)^T \in \mathbb{G}_1^2$. The opening of the commitment is the r .

Fig. 4: The commitment scheme \mathcal{COM}

5 A Signature Scheme based on KEA

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

The second building block is a quasi-adaptive non-interactive perfect zero-knowledge argument of knowledge \mathcal{NIZK}_{ext} . The argument system is adaptive weak knowledge sound⁵. Roughly speaking, the NIZK is a two-fold version of the scheme of Kiltz and Wee. For space reason we defer the details of the NIZK in Appendix C where we define also the \mathcal{D}_k -KerMDH and the KerLin₂ assumptions (see Escala *et al.* [12]), here we state the following theorem:

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

The Signature Scheme. The signature scheme \mathcal{SS}_2 is described in Fig. 5. We

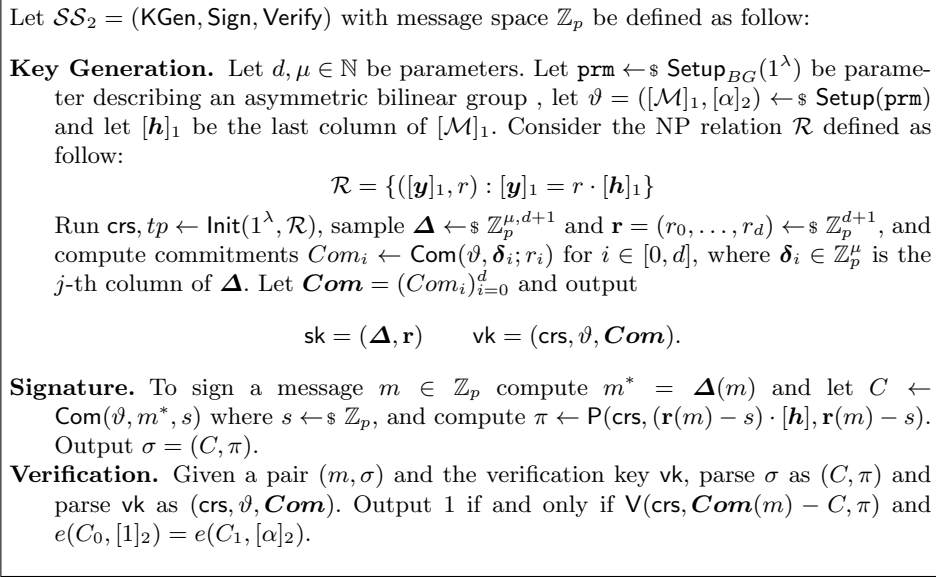


Fig. 5: The signature scheme \mathcal{SS}_2 .

show that the scheme is correct. For any tuple m, σ where σ is a valid signature for m with the verification key $\text{vk} = (\text{crs}, \vartheta, \mathbf{Com})$, let parse σ as (C, π) , we have:

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

⁵ We reverse the order of the quantifiers in the usual definition of knowledge soundness. Namely, for each adversary A there exists an extractor Ext . See more details in Appendix C.

The last equation follows because $[\mathbf{h}]_1$ 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.

Theorem 4. *Let $\mu, d \in \mathbb{N}$ and $\mu > 8$. If the $(\mu + 1)$ -KE* assumption and the KerLin₂ assumption hold over Setup_{BG} then, for any $0 \leq \ell \leq ((d + 1) \log \lambda) - \lambda$, the signature scheme \mathcal{SS}_2 described Fig. 5 is $(\ell, 1)$ -fully-leakage one-more unforgeable with leak-free key generation.*

We give an intuition of the proof. In particular, we explain how to use the knowledge of the exponent assumption of Def. 1. The main idea is to define a sampler that, roughly speaking, executes the fully-leakage one-more unforgeability experiment. More in details, the sampler \mathcal{S} samples all the randomness needed, including the secret key, the randomness for the signatures and the random tape of the adversary, with the only exception of the parameters of the KEA* assumption. The sampler proceeds with executing the experiment up to the moment before the adversary outputs its forgeries. Eventually the sampler outputs the full view of the adversary including the queried signatures, the leakage and the random tape of the adversary, let View be such value.

At this point we can deterministically execute the adversary feeding it with the view produced by the sampler. This adversary produces n commitment values (one for each forgery) for which, thanks to the knowledge of the exponent assumption he must know the opening.

Notice we do not incur in any problem of recursive composition of extractors. In fact the adversary outputs all its commitments at once. More in details, given the adversary code, for any $i \in [n]$, we can define the adversary A_i which outputs only the i -th commitment of A . Using the knowledge of the exponent assumption, for any index i , there must exist an extractor Ext_i for the adversary A_i . Crucially, the computational complexity of the extractor Ext_i depends only on A_i and not on Ext_j for an index $j \neq i$.

The proof continues showing that the extracted values are indeed evaluations of the polynomial Δ sampled by the sampler. To argue this we use the adaptive weak knowledge soundness of the NIZK. We give more details about this step in the formal proof.

Now, consider the predictor that on input the random variable View first runs the extractors $\text{Ext}_1, \dots, \text{Ext}_n$ obtaining n evaluation points of the polynomial Δ and then guesses a random polynomial that interpolates the evaluation points. The probability that this predictor guesses the polynomial Δ is roughly $\varepsilon p^{-(d-n)\mu}$ where ε is the winning probability of the adversary A . On the other hand, we prove that, thanks to perfect hiding and perfect zero-knowledge, no predictor can guess the polynomial Δ with probability more than $2^\ell p^{-(d+1)\mu}$. We complete the proof by noticing that the two bounds are in contradiction when ε is noticeable in the security parameter.

Proof (of Thm. 4). Let A be an adversary such that $\text{Adv}_{A, \mathcal{SS}_2}^{\text{one-more}^*}(\lambda, \ell, 1) = \varepsilon$ for parameter ℓ as described in the statement of the theorem. Let $\mathbf{H}_0(\lambda)$ be the experiment $\mathbf{Exp}_{\mathcal{SS}, A}^{\text{one-more}^*}(\lambda)$. Denote with $((m_1^*, (C_1^*, \pi_1^*)), \dots, (m_n^*, (C_n^*, \pi_n^*)))$ the

list of forgeries of A . During the experiment the adversary has oracle access to $\alpha = (\mathbf{\Delta}, \mathbf{r}, (s_j, z_j)_{j \in [q]})$ 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). The proof proceeds with an hybrid argument. In particular, the proof has seven main hybrid experiments named $\mathbf{H}_0, \dots, \mathbf{H}_7$ and other sub-hybrids that we name with $\mathbf{H}_{i,j}$ for $i \in \{2, 3\}$ and $j \in [n]$. Let Forge_i (resp. $\text{Forge}_{i,j}$) be the event that \mathbf{H}_i (resp. $\mathbf{H}_{i,j}$) returns 1, so that $\mathbf{P}[\text{Forge}_0] = \varepsilon$.

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

The hybrid \mathbf{H}_1 executes three steps: (1) it creates the parameters $(\text{prm}_{BG}, [\mathcal{M}]_1, [\alpha]_2)$, (2) it executes the sampler $(\text{vk}, \text{View}, r_A) \leftarrow \mathcal{S}([\mathcal{M}]_1, [\alpha]_2)$, (3) it runs $A(\text{vk}; r_A)$ and answers all the oracle queries using the information in View . 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 $\text{Ext}_1, \dots, \text{Ext}_i$ and runs the same as the hybrid \mathbf{H}_1 but, also, it runs the extractors and outputs 1 if and only if the extracted values match the commitments C_1^*, \dots, C_i^* . More in details, the hybrid $\mathbf{H}_{2,i}$ first creates the parameters $(\text{prm}_{BG}, [\mathcal{M}]_1, [\alpha]_2)$, then it executes the sampler $(\text{vk}, \text{View}, r_A) \leftarrow \mathcal{S}([\mathcal{M}]_1, [\alpha]_2)$, then it runs $A(\text{vk}; r_A)$ and answers all the oracle queries using the information in View . Eventually, A outputs its forgeries $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$ where $\sigma_i^* = (C_i^*, \pi_i^*)$, and for $j = 1, \dots, i$ the hybrid $\mathbf{H}_{1,i}$ computes $\mathbf{x}_i \leftarrow \text{Ext}_i([\mathcal{M}]_1, [\alpha]_2, (\text{vk}, \text{View}, r_A))$ 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 $j = 1, \dots, i$ we have $C_j^* = [\mathcal{M}]_1 \cdot \mathbf{x}_j$.

Claim. There exist PPT extractors $\text{Ext}_1, \dots, \text{Ext}_n$ such that for any $i > 1$, $|\varepsilon_{1,i-1} - \varepsilon_{1,i}| \in \text{negl}(\lambda)$. 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 $A'_i([\mathcal{M}]_1; r')$:

1. Parse r' as $(\text{vk}, \text{View}, r_A)$;

2. Run $A(\text{vk}; r_A)$ and answers all the oracle queries using the information in **View**;
3. Eventually, A outputs its forgeries $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$;
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 $[\mathbf{y}]_1 := C_i^*$.

For any PPT Ext_i the two hybrids diverge when $[\mathcal{M}]_1 \cdot \mathbf{x}_i \neq [\mathbf{y}]_1$, where \mathbf{x}_i is the output of the extractor, but the signature σ_i^* verifies correctly. Notice that the verification algorithm checks that $e([y_0]_1, [\alpha]_2) = e([y_1]_1, [1]_2)$, where $\mathbf{y} = (y_0, y_1)$ and so $[\mathbf{y}]_1 \in \text{Span}([1, \alpha]_1)$. Therefore:

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

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 n different PPT extractors $\text{Ext}_1, \dots, \text{Ext}_n$ plus i different PPT extractors $\text{Ext}'_1, \dots, \text{Ext}'_i$ and runs the same as the hybrid $\mathbf{H}_{2.n}$ but also for any $j = 1, \dots, i$ it computes $w_i \leftarrow \text{Ext}'_i(\text{crs}, tp, r')$ where $r' = (\Delta, \mathbf{r}, [\mathbf{g}, h], \alpha)$ 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 $j = 1, \dots, n$ we have $C_j^* = [\mathcal{M}]_1 \cdot \mathbf{x}_j$ and,
- (c) for any $j = 1, \dots, i$ check $\mathbf{Com}(m_i^*) - C_i^* = w_i \cdot [h, \alpha h]$.

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

The claim follows by the weak knowledge soundness of \mathcal{NIZK}_{ext} .

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 $A'_i(\text{crs}; r')$:

1. Parse r' as $(\Delta, \mathbf{r}, [\mathbf{g}^T, h], \alpha)$, define ϑ as described in **Setup**, compute \mathbf{Com} using Δ and \mathbf{r} , and set the verification key $\text{vk} = (\text{crs}, \vartheta, \mathbf{Com})$ and $\text{sk} = (\Delta, \mathbf{r})$;
2. Run $A(\text{vk}; r_A)$ and answer all the signature oracle queries using $\text{Sign}(\text{sk}, \cdot)$ and the leakage oracle queries with α ;
3. Eventually, A outputs its forgeries $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$;
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 $(\mathbf{Com}(m_i^*) - C_i^*)$ and the proof π_i^* .

For any PPT Ext'_i the two hybrids diverge when conditions (a) and (b) holds but $\mathbf{Com}(m_i^*) - C_i^* \neq w_i \cdot [h, \alpha h]$ happens. Clearly, condition (a) implies that $\mathbf{Com}(m_i^*) - C_i^* \in \text{Span}([h, \alpha h])$ 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 conditions are changed as follow:

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

Claim. $|\varepsilon_{3,n} - \varepsilon_4| \in \text{negl}(spar)$.

The claim follows by a simple reduction to the DLOG problem.

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

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

1. Run the hybrid \mathbf{H}_4 with the parameter set to $[\mathbf{g}]_1$, 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 $\mathbf{\Delta} \leftarrow \$_{\mathbb{Z}_p^{\mu, d+1}}$ 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 $\mathbf{x}'_i \neq \mathbf{\Delta}(m_i^*)$.
3. Parse \mathbf{x}_i as $(x_{i,1}, \dots, x_{i,\mu+1})$ and $\mathbf{\Delta}(m_i^*)$ as (y_1, \dots, y_μ) , output the vectors

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

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

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

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 conditions 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 \mathbf{P} is executed with input the output produced by \mathcal{S} . Eventually, the predictors (which runs internally \mathbf{A}) receives n forgeries $(m_1^*, \sigma_1^*), \dots, (m_n^*, \sigma_n^*)$. The predictor checks the winning conditions (a),(b),(d) of the hybrid \mathbf{H}_4 and does as follow:

- If one of the conditions is invalid then output \perp ,
- else for $j \in [\mu]$ sample a polynomial δ_j^* in $\mathbb{Z}_p[X]$ of degree d such that $\delta_j(m_i^*) = \mathbf{x}'_{i,j}$ for $i \in [n]$,
- outputs $\Delta^* = (\delta_1^*, \dots, \delta_\mu^*)$.

Recall that the advantage of \mathbf{A} in the one-more unforgeability game is ε .

Lemma 3. $\Pr[\mathbf{P}(\mathcal{S}([\mathcal{M}]_1, [\alpha]_2)) = \Delta] \geq \exp(((n-d) \cdot \mu) \log p) \cdot (\varepsilon - \text{negl}(\lambda))$.

Proof. By the triangular inequality and the claims above we have that $\varepsilon_5 \geq \varepsilon - \text{negl}(\lambda)$. When the event Forge_5 happens then $\Delta(m_i^*) = \mathbf{x}'_i$ for $i \in [n]$ so the event that $\Delta^* = \Delta$ is equivalent to the event that the predictor \mathbf{P} correctly guesses the remaining $d-n$ zeros of the polynomials δ_i for $i \in [\mu]$ which is equal to $1/p^{\mu(d-n)} = \exp(((n-d) \cdot \mu) \log p)$.

Lemma 4. For any prm $\leftarrow \text{Setup}_{BG}(1^\lambda)$, any $([\mathcal{M}]_1, [\alpha]_2) \in \mathbb{G}_1^{2 \cdot \mu + 1} \times \mathbb{G}_2$ and any predictor \mathbf{P}' we have $\Pr[\mathbf{P}'(\mathcal{S}([\mathcal{M}]_1, [\alpha]_2)) = \Delta] \leq \exp(-(d+1)\mu \log p + \ell)$.

Proof. We define two samplers \mathcal{S}_1 and \mathcal{S}_2 , we prove that their output distributions $(\text{vk}, (\sigma_1, \dots, \sigma_q, \text{Leak}), r_A)$ are equivalent to the distribution of \mathcal{S} , and moreover, the components $\text{vk}, \sigma_1, \dots, \sigma_q, r_A$ are independent of Δ as sampled by \mathcal{S} . Both the sampler \mathcal{S}_1 and \mathcal{S}_2 are not efficiently computable, however, this is not a problem as we are proving that their distributions are identically distributed to the distribution of \mathcal{S} .

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

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

Claim. For any parameter $\text{prm} \leftarrow \text{Setup}_{BG}(1^\lambda)$ 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.

Proof. We notice that for any \mathbf{m} the commitment to \mathbf{m} is uniformly distributed over $\text{Span}([1, \alpha])$. Therefore, for any Com_i (resp. C_i), it always exists such r_i (resp. s_i), and moreover, once Δ and 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 \mathbf{S} of NIZK and, moreover, the randomness z_i is uniformly sampled over the set⁶

$$\{z_i : \pi_i = \text{P}(\text{crs}, (\mathbf{r}(\Delta)(m_i) - \mathbf{s}(\Delta)) \cdot [\mathbf{h}], (\mathbf{r}(\Delta)(m_i) - \mathbf{s}(\Delta)))\}$$

where $\mathbf{r}(\Delta)$ is the vector of the randomness as computed by \mathcal{S}_1 .

Claim. For any parameter $\text{prm} \leftarrow \text{Setup}_{BG}(1^\lambda)$ 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 $z_i(\Delta)$ to stress that z_i is computed as function of Δ , for each leakage oracle query f the answer of f is computed as $f(\Delta, \mathbf{r}(\Delta), (s_i(\Delta), z_i(\Delta))_{i \leq q})$.

Claim. For any \mathbf{P}' we have $\Pr[\mathbf{P}'(\mathcal{S}_2([\mathcal{M}]_1, [\alpha]_2)) = \Delta] \leq \exp(-(d+1)\mu \log p + \ell)$.

Proof. Let q be the number of signature queries made by \mathbf{A} and let Leak the concatenation of all the leakage performed by \mathbf{A} . For any predictor \mathbf{P}'

$$\Pr[\mathbf{P}'(\mathcal{S}_2([\mathcal{M}]_1, [\alpha]_2)) = \Delta] = \Pr[\mathbf{P}'(\text{Leak}) = \Delta] \quad (1)$$

$$\begin{aligned} &= \sum_L \Pr[\mathbf{P}'(L) = \Delta \mid \text{Leak} = L] \Pr[\text{Leak} = L] \\ &\leq 2^\ell \max_D \Pr[\Delta = D]. \end{aligned} \quad (2)$$

Where Eq. 1 holds because \mathbf{vk}, r_A and the signatures $\sigma_1, \dots, \sigma_q$ are sampled independently from Δ , while Eq. 2 holds applying the chain rule (Lemma 5). Finally we notice that Δ is sampled uniformly at random so the statement of the claim follows.

⁶ Namely, the set of assignment for the randomness z_i for for which the execution of P with randomness z_i and the appropriate tuple instance and witness does compute exactly the proof π_i .

By putting together the first two claims we have that the probability of guessing Δ by a predictor given in input the output produced by \mathcal{S}_2 is the same as it gets in input the output produced by \mathcal{S}_1 , by the last claim, therefore, the lemma follows.

Returning to the proof of the theorem, we can put together the inequalities of Lemma 3 and Lemma 4, and by taking the logarithms we have:

$$-d\mu \log p + \ell \geq -(d - n)\mu \log p + \log(\varepsilon - \mathbf{negl}(\lambda))$$

By adding $d\mu \log p$ to both sides we derive that $\ell \geq n\mu \log p + \log(\varepsilon - \mathbf{negl}(\lambda))$, and by the fact that $n > \frac{\ell}{s \cdot \gamma} + 1$ and $\gamma = 1$ we derive that $-\log(\varepsilon - \mathbf{negl}(\lambda)) > s \geq \lambda$. For the equation above to hold, necessarily, ε is negligible in λ .

6 Acknowledgements

I would like to thank Dario Fiore for a conversation we had on his paper [16]. Also, I would like to thank Dennis Hofheinz which suggested to me the paper of Fujisaki on ABM Encryption.

References

1. M. Abe and S. Fehr. Perfect NIZK with adaptive soundness. In *TCC 2007*, pages 118–136, 2007.
2. J. Alwen, Y. Dodis, M. Naor, G. Segev, S. Walfish, and D. Wichs. Public-key encryption in the bounded-retrieval model. In *EUROCRYPT 2010*, pages 113–134, 2010.
3. J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In *CRYPTO 2009*, pages 36–54, 2009.
4. G. Ateniese, A. Faonio, and S. Kamara. Leakage-resilient identification schemes from zero-knowledge proofs of storage. In *15th IMA International Conference on Cryptography and Coding*, pages 311–328, 2015.
5. M. Bellare and A. Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In *CRYPTO 2004*, pages 273–289, 2004.
6. N. Bitansky, R. Canetti, O. Paneth, and A. Rosen. On the existence of extractable one-way functions. In *46th ACM STOC*, pages 505–514, 2014.
7. E. Boyle, G. Segev, and D. Wichs. Fully leakage-resilient signatures. In *EUROCRYPT 2011*, pages 89–108, 2011.
8. 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, 2010.
9. I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In *CRYPTO'91*, pages 445–456, 1992.
10. Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs. Cryptography against continuous memory attacks. In *51st FOCS*, pages 511–520, 2010.
11. Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *EUROCRYPT 2004*, pages 523–540.

12. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. Villar. An algebraic framework for Diffie-Hellman assumptions. In *CRYPTO 2013, Part II*, pages 129–147, 2013.
13. A. Faonio and J. B. Nielsen. Fully leakage-resilient codes. In *PKC 2017, Part I*, pages 333–358, 2017.
14. A. Faonio, J. B. Nielsen, and D. Venturi. Mind your coins: Fully leakage-resilient signatures with graceful degradation. In *ICALP 2015, Part I*, pages 456–468, 2015.
15. 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.
16. D. Fiore and A. Nitulescu. On the (in)security of SNARKs in the presence of oracles. In *TCC 2016-B, Part I*, pages 108–138, 2016.
17. E. Fujisaki. All-but-many encryption - A new framework for fully-equipped UC commitments. In *ASIACRYPT 2014, Part II*, pages 426–447, 2014.
18. E. Fujisaki. All-but-many encryption. *J. Cryptology*, 31(1):226–275, 2018.
19. K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic analysis: Concrete results. In *CHES 2001*, pages 251–261, 2001.
20. R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *EUROCRYPT 2013*, pages 626–645, 2013.
21. J. Groth. Short pairing-based non-interactive zero-knowledge arguments. In *ASIACRYPT 2010*, pages 321–340, 2010.
22. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *EUROCRYPT 2008*, pages 415–432, 2008.
23. D. Hofheinz. All-but-many lossy trapdoor functions. In *EUROCRYPT 2012*, pages 209–227, 2012.
24. C. S. Jutla and A. Roy. Shorter quasi-adaptive NIZK proofs for linear subspaces. In *ASIACRYPT 2013, Part I*, pages 1–20, 2013.
25. E. Kiltz and H. Wee. Quasi-adaptive NIZK for linear subspaces revisited. In *EUROCRYPT 2015, Part II*, pages 101–128, 2015.
26. P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *CRYPTO'96*, pages 104–113, 1996.
27. P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *CRYPTO'99*, pages 388–397, 1999.
28. T. Malkin, I. Teranishi, Y. Vahlis, and M. Yung. Signatures resilient to continual leakage on memory and computation. In *TCC 2011*, pages 89–106, 2011.
29. M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In *CRYPTO 2009*, pages 18–35, 2009.
30. J. B. Nielsen, D. Venturi, and A. Zottarel. Leakage-resilient signatures with graceful degradation. In *PKC 2014*, pages 362–379, 2014.
31. 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, 2013.
32. T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *CRYPTO'91*, pages 129–140, 1992.

A Information-Theoretic Lemmas

We state a lemma from Dodis *et al.* [11]. In their paper the lemma is stated in terms of the conditional average min-entropy of a random variable X given another random variable Y . To keep the notation lighter, we instead give a (weaker) version of the chain rule for the average conditional min-entropy (Lemma 2.2 in [11]) which is sufficient for our purpose.

Lemma 5. *Let X, Y be random variables. If Y has at most 2^ℓ possible values, then*

$$\mathbb{E}_{y \leftarrow Y} \left[\max_{\mathbf{P}} \Pr [\mathbf{P}(y) = X] \right] \leq 2^\ell \max_x \Pr [X = x].$$

In a successive work, Alwen *et al.* [3] generalized the notion of average conditional min-entropy to predictor participating in interactive experiments (instead of predictors that gets on input the random variable Y).

Lemma 6. *Let X be a random variable and $\mathcal{O}_X(\cdot)$ be a leakage oracle instantiated with X which outputs at most ℓ bits, then for any predictor \mathbf{P} with oracle access to \mathcal{O}_X :*

$$\Pr \left[\mathbf{P}^{\mathcal{O}_X(\cdot)} = X \right] \leq 2^\ell \max_x \Pr [X = x].$$

B Commitment Schemes

B.1 Security properties

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

Binding Property. Consider the following probability:

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

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

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

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

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 $\text{Com}(\vartheta, m_0)$ and $\text{Com}(\vartheta, m_1)$ are indistinguishable.

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*

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

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

C.2 Quasi-Adaptive NIZK argument systems

Experiment $\underline{\text{Exp}}_{A, \mathcal{NIZK}, \mathcal{D}_R}^{\text{snd-QANIZK}}(\lambda)$:	Experiment $\underline{\text{Exp}}_{A, \text{Ext}, \mathcal{NIZK}, \mathcal{D}_R}^{\text{knw-QANIZK}}(\lambda)$:
<ol style="list-style-type: none"> 1. $\text{prm} \leftarrow \text{Setup}_{BG}(\lambda), \rho \leftarrow \mathcal{D}_R(\text{prm})$ $\text{crs}, \text{tp} \leftarrow \text{Init}(\text{prm}, \rho)$; 2. $(x, \pi) \leftarrow A(\text{crs})$; 3. Output $(x \notin \mathcal{L}_\rho \wedge \forall(\text{crs}, x, \pi) = 1)$. 	<ol style="list-style-type: none"> 1. $\text{prm} \leftarrow \text{Setup}_{BG}(\lambda), \rho \leftarrow \mathcal{D}_R(\text{prm})$, $\text{crs}, \text{tp} \leftarrow \text{Init}(\text{prm}, \rho)$; 2. $(x, \pi) \leftarrow A(\text{crs}; r), w \leftarrow \text{Ext}(\text{crs}; r)$ and $r \leftarrow \{0, 1\}^\lambda$; 3. Output $((x, w) \notin \mathcal{R}_\rho \wedge \forall(\text{crs}, x, \pi) = 1)$.

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

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

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

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

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

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

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 : $\text{Adv}_{A, \text{Ext}, \mathcal{NIZK}, \mathcal{D}_R}^{\text{knw-QANIZK}}(\lambda) \in \text{negl}(\lambda)$.

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

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

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

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

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

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

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

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

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

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

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

We define zero-knowledge:

Definition 8. *There exists a PPT simulator S such that for any λ , any prm output by $\text{Setup}_{BG}(1^\lambda)$ any ρ output by \mathcal{D}_R any crs, tp output by $\text{Init}(\text{prm}, \rho)$ any $(x, w) \in \mathcal{R}_\rho$ the distributions*

$$\text{P}(\text{crs}, x, w) \text{ and } \text{S}(\text{crs}, tp, x)$$

are the same (where the coin tosses are taken over P and Sim).

C.3 Constructions

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

Theorem 5 (Kiltz and Wee, [25]). *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.

The quasi-adaptive NIZK argument system in Fig 6 is a variation of [25]. 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 bilinear group) over matrices $[\mathbf{H}]_1 \in \mathbb{G}^{n,t}$ we say that the distribution \mathcal{D}_R is *witness sampleable* if there exists another efficiently sampleable distribution \mathcal{D}'_R over matrices $\mathbf{H}' \in \mathbb{Z}_p^{n,t}$ such that $[\mathbf{H}]_1 \equiv [\mathbf{H}']_1$. When proving knowledge soundness for witness sampleable distributions we additionally give to the extractor the matrix \mathbf{H}' as extra auxiliary input. We restate the theorem 3 from Sec. 5:

Theorem 6. *The quasi-adaptive argument system \mathcal{NIZK}_{ext} in Fig. 6 is perfect zero-knowledge and if both the \mathcal{D}_k -KerMDH assumption and the 1-KE* assumption hold for Setup_{BG} and, moreover, the distribution \mathcal{D}_R is witness-sampleable and $k+1 > t$, 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 \mathcal{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 $[\mathbf{H}]$ samples $\beta \leftarrow \$_{\mathbb{Z}_q}$ and computes $[\mathbf{P}']_1 \leftarrow \beta \cdot [\mathbf{P}]_1$ and $[\mathbf{C}'] \leftarrow \beta \cdot [\mathbf{C}]_1$ to create a CRS for \mathcal{NIZK}_{ext} . Eventually, A outputs a tuple statement $[\mathbf{y}]$ and a proof π, π' and A' outputs $[\mathbf{y}], \pi$. 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 $\text{Win} := \text{Exp}_{A, \text{Ext}, \mathcal{D}_R}^{\text{snd-QANIZK}}$ and let $\varepsilon := \Pr[\text{Win}]$. Let Sound be the event that the the element $[\mathbf{y}]_1$ output by A is indeed in the column span of $[\mathbf{H}]$ and let Forge be the event that $\text{Win} \wedge \text{Sound}$, meaning that the proof verify and $[\mathbf{y}]_1$ is a valid instance. Let \mathbf{H}_0 be the hybrid experiment that it is equivalent to $\text{Exp}_{A, \text{Ext}, \mathcal{D}_R}^{\text{snd-QANIZK}}$. We define $\varepsilon_i := \Pr[\text{Forge}]$ where the probability is taken over the experiment \mathbf{H}_i . By the argument given above we know that $|\varepsilon - \varepsilon_1| \leq \Pr[\neg \text{Sound}] \leq \text{negl}(\lambda)$.

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

Proof. The two hybrids diverge when the proof (π, π') verify, the instance $[\mathbf{y}]_1$ 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 -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 $e(\pi, [\mathbf{A}]_2) = e([\mathbf{y}]_1, [\mathbf{C}]_2)$ and $e(\pi', [\mathbf{A}]_2) = e([\mathbf{y}]_1, [\beta \cdot \mathbf{C}]_2)$, and therefore $e(\pi \cdot \beta, [\mathbf{A}]_2) = e(\pi', [\mathbf{A}]_2)$. Now, let $[\mathbf{z}]_1 \leftarrow \pi \cdot \beta - \pi'$ we have that $e([\mathbf{z}]_1, [\mathbf{A}]_2) = 0$ but $[\mathbf{z}]_1 \neq 0$, so we clearly break the \mathcal{D}_k -KerMDH assumption. (Notice that in the reduction we can sample $\beta \leftarrow \$_{\mathbb{Z}_p}$).

Recall that the proof π is a vector in \mathbb{G}_1^{k+1} , so we can parse π as $(\pi_1, \dots, \pi_{k+1})$. In the next part of the proof we show that it is possible to extract, one by one, all the discrete logarithms of the components of π .

Let $\mathbf{H}_{2,i}$ be an hybrid that takes as parameter and extractor $\text{Ext}_1, \dots, \text{Ext}_i$, 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 $\tilde{x}_i \leftarrow \text{Ext}([1, \beta]_1, [\beta]_2, r)$ 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 $\pi_j = [\tilde{x}_j]$.

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

Proof. For any index $i \in [1, k+1]$ 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 $\pi_i = [\tilde{x}_i]_1$. We define an adversary for the KE assumption on bilinear group.

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

1. Create all the parameters of the crs, in particular, since \mathcal{D}_R is witness sampleable, first sample $\mathbf{H} \leftarrow \mathbb{Z}_q^{n,t}$, $\mathbf{A} \leftarrow \mathcal{D}_k$ and $\mathbf{K} \leftarrow \mathbb{Z}_q^{n,k}$ and set $[\mathbf{P}]_1 = [\mathbf{M}^T \cdot \mathbf{K}]$, $[\mathbf{P}']_1 = [\beta]_1 \cdot (\mathbf{M}^T \cdot \mathbf{K})$ and similarly $[\mathbf{C}]_2$ and $[\mathbf{C}']_2 = [\beta]_2 \cdot (\mathbf{K} \cdot \mathbf{A})$.
2. Run the adversary \mathbf{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 \mathbf{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} . Notice that, with witness-sampleable distribution

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

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

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

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

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