

Improved Zero-Knowledge Proofs for Lattice Encryption Schemes, Linking Classical and Lattice Primitives, and Applications

Abstract. It now seems probable that a quantum computer breaking non-trivial instances of the RSA and discrete logarithm problem will be built in the not-too-distant future. Despite this, for many applications such as electronic voting or anonymous authentication where simple encryption and signature schemes are not sufficient, there are currently no practical protocols that one could use to defend against such quantum adversaries. One might think that such protection is not yet needed, but this creates vulnerabilities against an adversary who collects data today and is then able to break all the privacy and anonymity properties of current protocols using a quantum computer in the future.

In the absence of efficient lattice schemes for advanced protocols, in this work we begin the investigation of “hybrid” schemes that try to maintain the efficiency of classical protocols, but with the forward-security of lattice-based ones. Our main result is a proposal of one such protocol – a group signature scheme that combines the efficiency of today’s number theoretic schemes, while using lattice encryption to guarantee privacy to its members even after quantum computers are developed. As intermediate steps, we provide a new and more efficient protocol for proving the validity of a lattice-based ciphertexts, and then extend this to a protocol that proves that a classical commitment and a lattice-based encryption contain the same secret. We believe that these two protocols are also interesting in their own right. From these two primitives, our group signature construction then follows a generic folklore approach, for which, to the best of our knowledge, we provide the first formal security proof.

Keywords: Verifiable Encryption, Group Signatures, Zero-Knowledge Proofs for Lattices.

1 Introduction

With quantum computing looming in the (possibly) near future, there has been a push over the last few years towards constructing efficient primitives based on the hardness of lattice problems, which are believed to be resilient against quantum algorithms. Unfortunately, except for encryption and digital signatures, lattice schemes seem to be a lot less practical (especially space-wise) than their classical counterparts. And so if one wishes to use quantum-secure protocols, his choices are either to keep using the current protocols until they are no longer secure, or to migrate to the less-efficient lattice protocols.

We believe that the decision of whether to switch or not should actually depend on the primitive itself. For example, if one were using a statistically-hiding commitment scheme to commit to sensitive data, then quantum computing does not really pose any threat to the data. Once quantum computers are deployed, the commitment will of course be invalid (since it will no longer be binding), but the message will still be hidden. On the other hand, a number theoretic commitment scheme that is only computationally-hiding will not protect the message contents against quantum computing. In short, we believe that in schemes where quantum computers do not pose a threat to privacy, migrating to quantum-secure constructions is not as urgent.

Some more involved schemes are comprised of both types of primitives – ones that would leak private data and others that information-theoretically hide it. In the absence of efficient lattice constructions for such schemes, our proposal is to construct them using a “hybrid” approach between classical and lattice cryptography. The idea is that in a “pre-quantum” world, they would be more efficient than pure lattice-based schemes. And once quantum computers are built, all the private information contained in the ciphertexts or proofs that are generated while using the schemes would still remain secret.

A good example of a scheme that would benefit from such a construction is a group signature. The “raison d’être” of a group signature scheme is to hide the identity of a user who is signing a message in the name of the group - and so it is imperative that the user’s identity remains hidden even after quantum computers come about and the signature scheme can no longer be used. In a group signature scheme, while every group member is able to anonymously sign, the master group authority still has a way to unmask a user (in case of disputes or misbehavior, for example). A high level overview for how such a scheme is realized is for a group member to produce a signature using a key he was given by the master authority, then encrypt his identity, and then produce a zero-knowledge proof linking the ciphertext to the key used in the signature (i.e., his identity). A verifier is able to ascertain that the signature and the zero-knowledge proof are valid, but does not learn anything about the identity of the user. The master authority, however, can simply decrypt the ciphertext containing the user’s identity.

It would be preferable, therefore, if the encryptions of identities were lattice-based, whereas everything else in the scheme would information-theoretically hide the identity of the user. In this way, assuming that lattice encryption schemes are secure against quantum attacks, a signer’s identity would remain hidden.

The main result of this work is a realization of such “hybrid” schemes. The schemes themselves are comprised of several components, which we believe are interesting in their own rights. The first is a zero-knowledge proof of plaintext knowledge of lattice encryption schemes, and the second is a technique to “link” classical commitments to lattice encryptions by proving in zero-knowledge that they contain the same message.

1.1 Improved Proofs of Plaintext Knowledge for Lattice Schemes

In a zero knowledge proof of plaintext knowledge, the encryptor wants to prove in zero knowledge that the ciphertext is of the correct form and that he knows the message. Efficient constructions of these primitives are known based on number-theoretic hardness assumptions such as discrete log, strong RSA, etc..

Ciphertexts in lattice-based encryption schemes generally have the form $t = Ae \bmod q$, where A is some public matrix and e is the unique vector with small coefficients that satisfies the equation (in this general example, we are lumping the message with e). A proof that t is a valid ciphertext (and also a proof of plaintext knowledge), therefore involves proving that one knows short a short e such that $Ae = t$. It is currently known how to accomplish this task in two ways. The first uses a “Stern-type” protocol [35] in which every run has soundness error $2/3$ [24]. It does not seem possible to improve this protocol since some steps in it are inherently combinatorial and non-algebraic.

A second possible approach is to use the Fiat-Shamir approach for lattices using rejection sampling introduced in [25,27]. But while the latter leads to fairly efficient Fiat-Shamir signatures, there are some barriers to obtaining a proof of knowledge. What one is able to extract from a prover are short vectors r', z' such that $Ae' = tc$ for some integer c , which implies that $Ae'c^{-1} = t$. Unfortunately, this does not imply that $e'c^{-1}$ is short unless $c = \pm 1$. This is the main way in which lattice-based Fiat-Shamir proofs differ from traditional schemes like the discrete-log based Schnorr protocol. In the latter, it is enough to extract any discrete log, whereas in lattice protocols, one must extract a *short* vector. Thus, the obvious approach at Fiat-Shamir proofs of knowledge for lattices (i.e. using challenge vectors from $\{0, 1\}$) also leads to protocols with soundness error $1/2$.

Things do not improve for lattice-based proofs of knowledge even if one considers ideal lattices. Even if A and e are a matrix and a vector of polynomials in the ring $\mathbb{Z}_q[X]/(X^n + 1)$, and c is now a polynomial (as in the Ring-LWE based Fiat-Shamir schemes in [26,27]), then one can again extract only a short e' such that $Ae'c^{-1} = t$. In this case, not only is c^{-1} not necessarily short, but it does not even necessarily exist (since the polynomial $X^n + 1$ can factor into up to n terms.) At this point, we are not aware of any techniques that reduce the soundness error of protocols that prove plaintext knowledge of lattice encryptions, except by (parallel) repetition.

In a recent work, Damgard et al. [18] gave an improved *amortized* proof of plaintext knowledge for LWE encryptions. Their very nice protocol allows to prove knowledge of $O(k)$ plaintexts (where k is the security parameter) in essentially the same time as just one plaintext. But it seems that the ideas behind the protocol do not reduce the time requirement for proving just one plaintext, nor do they apply to Ring-LWE based encryption schemes. In particular, Ring-LWE based schemes are able to encrypt $O(n)$ plaintext bits into one (or two) polynomial, which is often all that is needed. Yet, the techniques in [18] do not seem to be helpful here. The reason is that the challenge matrix required in [18] needs to be of a particular form and cannot simply be a ring element in $\mathbb{Z}_q[X]/(X^n + 1)$.

In this paper, we first show that one can reduce the soundness error of lattice-based zero-knowledge proofs of knowledge for ciphertext validity from $1/2$ to $1/(2n)$, which in practice decreases the number of required iterations of the protocol by a factor of approximately 10. Interestingly, our techniques only work for ideal lattices, and we do not know how to adapt them to general ones. The key observation is that when working over the ring $\mathbb{Z}[X]/(X^n + 1)$, the quantity $2/(X^i - X^j)$ for all $0 \leq i \neq j < n$ is a polynomial with coefficients in $\{-1, 0, 1\}$ (this is proved in Section 3.1).

This immediately allows us to prove that, given A and t , we know a vector of short polynomials e such that $Ae = 2t$. While this is not quite the same as proving that $Ae = t$, it is good enough for most applications, since it still allows us to prove knowledge of the plaintext. This result immediately gives improvements in all schemes

that require such a proof of knowledge for Ring-LWE based encryption schemes such as the ring-version of the “dual” encryption scheme [20], the “two element” scheme of Lyubashevsky et al. [28], and NTRU [22,34].

1.2 Linking Lattice and Classical Commitments

A main step in constructing our “hybrid” schemes is to prove that two primitives, one based on classical cryptography and the other on lattices, are committing to the same message (and that the prover knows the message). In our application, we will use the perfectly hiding Pedersen commitment as the classical, and a Ring-LWE encryption scheme as the lattice-based primitive.

While Pedersen commitments and lattice-based encryption schemes work over rather different rings, we show that we can still perform operations “in parallel” on the two commitments. For example, if the message is $\mu_0, \mu_1, \dots, \mu_{n-1}$, then it is encrypted in Ring-LWE schemes as the polynomial $\mu = \mu_0 + \mu_1 X + \dots + \mu_{n-1} X^{n-1}$, and each μ_i is encrypted individually using a Pedersen commitment. We will then want to prove that a Ring-LWE encryption of μ encrypts the same thing as n Pedersen commitments of μ_i . Even though the two computations are performed over different rings, we show that by mimicking polynomial multiplications over a polynomial ring by appropriate additions and multiplications of coefficients in exponents, we can use our previously mentioned proof of plaintext knowledge to both prove knowledge of μ and show that the Pedersen commitments are committing to the coefficients μ_i . One of the reasons that we are able to provide such a proof is that the terms dealing with μ (and μ_i) in the proof of knowledge are worked on “over the integers” – that is, modular reduction never needs to be performed on these terms. We describe this protocol in detail in Section 4.

1.3 Applications to Group Signatures and Credentials

Group signatures [12] are schemes that allow members of a group to sign messages on behalf of the group without revealing their identity. In case of a dispute, the group manager can lift the signer’s anonymity and reveal his identity. Currently known group signatures based on lattice assumptions should mainly be seen as proofs of concepts, rather than practically useful schemes. The schemes by Gordon et al. [21] and Camenisch et al. [10] have signature size linear in the number of group members. The scheme due to Laguillaumie et al. [23] performs much better asymptotically with signature sizes logarithmic in the number of group members, but, as the authors admit, instantiating it with practical parameters would lead to very large keys and signature sizes. This is in contrast to classical number-theory based solutions, where both the key and the signature size are constant for arbitrarily many group members.

As previously mentioned, quantum computers may have very different impacts for security and privacy: For controversial topics, privacy might be a long term concern, and a user might only be willing to use a scheme if he is convinced that his identity will remain secret, say, fifty years from now. The implications of signature forgeries being broken may be less significant, because one can simply stop accepting classical signatures as soon as quantum computers become practical.

Following this observation, we propose a “hybrid” group signature scheme, where unforgeability holds under classical assumptions, whereas privacy is proved under lattice-based ones. This allows us to combine the flexible tools that are available in the classical framework with the strong privacy guarantees of lattice problems. As a result, we obtain a group signature scheme that is private under quantum-resistant assumptions, where the key and the signature size are only logarithmic in the number of group members. For practical choices of parameters and realistic numbers of group members, the sizes will even be independent of the number of users. Furthermore, by basing our scheme on ring-LWE and not standard LWE, we partially solve an open problem stated in [23].

Our construction follows a generic approach that we believe is folklore, as it underlies several direct constructions in the literature [2,8] and a variant was described explicitly by Chase and Lysyanskaya [11]. When joining the group, a user obtains a certificate from the group manager that is a signature on his identity under the group manager’s public key. To sign a message, the user now encrypts his identity under the manager’s public encryption key, and then issues a signature proof of knowledge that he possesses a valid signature on the plaintext contained in this encryption. Our construction follows this general paradigm, but with some modifications to better fit the specifics of our proof of plaintext knowledge for lattice encryption. To the best of our knowledge, however, the construction was never proved secure, so our proof can be seen as a contribution of independent interest.

2 Preliminaries

2.1 Notation

We denote algorithms by sans-serif letters such as A, B .

If \mathcal{S} is a set, we write $s \xleftarrow{\$} \mathcal{S}$ to denote that s was drawn uniformly at random from \mathcal{S} . Similarly, we write $a \xleftarrow{\$} A$ if a was computed by a randomized algorithm A , and $d \xleftarrow{\$} D$ for a probability distribution D , if d was drawn according to D .

We write $\Pr[\mathcal{E} : \Omega]$ to denote the probability of event \mathcal{E} over the probability space Ω . For instance, $\Pr[x = y : x, y \xleftarrow{\$} D]$ denotes the probability that $x = y$ if x, y were drawn according to a distribution D .

If v is a vector, we denote by $v_{\ll l}$ an anti-cyclic shift of a vector v by l positions, corresponding to a multiplication by X^l in $R_q = \mathbb{Z}_q[X]/(X^n + 1)$. That is, $v_{\ll l} = (v_0, \dots, v_{n-1})_{\ll l} = (-v_{n-l}, \dots, -v_{n-1}, v_0, \dots, v_{n-l-1})$. Furthermore, we identify the vectors (a_0, \dots, a_{n-1}) with the polynomial $a_0 + a_1X + \dots + a_{n-1}X^{n-1}$.

Throughout the paper, λ denotes the main security parameter and ε denotes the empty string.

2.2 Commitment Schemes and Pedersen Commitments

Informally, a commitment scheme is a tuple $(\text{CSetup}, \text{Commit}, \text{COpen})$, where CSetup generates commitment parameters, which are then used to commit to a message m using Commit . A commitment cmt can then be verified using COpen . Informally, a commitment scheme needs to be binding and hiding. The former means that no cmt can be opened to two different messages, while the latter guarantees that cmt does not leak any information about the contained m . A formal definition can be found in Appendix A.1.

We next recap an efficient commitment scheme that was introduced by Pedersen [30]. Let therefore $\{\mathbb{G}(\lambda)\}_{\lambda \in \mathbb{N}}$ be a family of prime order groups such that the discrete logarithm problem is hard in $\mathbb{G}(\lambda)$ for security parameter λ . Let $\tilde{q} = \tilde{q}(\lambda)$ be the order of $\mathbb{G}(\lambda)$.

To avoid confusion in the following, all elements with order \tilde{q} are denoted with a tilde in the following.

CSetup. This algorithm chooses $\tilde{g}, \tilde{h} \xleftarrow{\$} \mathbb{G}(\lambda)$ and outputs $cpars = (\tilde{g}, \tilde{h})$.

Commit. To commit to a message $m \in \mathcal{M} = \mathbb{Z}_{\tilde{q}}$, this algorithm first chooses $r \xleftarrow{\$} \mathbb{Z}_{\tilde{q}}$. It then outputs the pair $(\widetilde{cmt}, o) = (m\tilde{g} + r\tilde{h}, r)$.

COpen. Given a commitment \widetilde{cmt} , an opening o , a public key $cpars$ and a message m , this algorithm outputs accept if and only if $\widetilde{cmt} \stackrel{?}{=} m\tilde{g} + o\tilde{h}$.

Theorem 2.1. *Under the discrete logarithm assumption for \mathbb{G} , the given commitment scheme is perfectly hiding and computationally binding.*

2.3 Semantically Secure Encryption and NTRU

A semantically secure (or IND-CPA secure) encryption scheme is a tuple $(\text{EncKG}, \text{Enc}, \text{Dec})$ of algorithms, where EncKG generates public/private key pair, Enc can be used to encrypt a message m under the public key, and the message can be recovered from the ciphertext by Dec using the secret key. Informally, while $\text{Dec}(\text{Enc}(m)) = m$ should always hold, only knowing the ciphertext and the public key should not leak any information about the contained message. A formal definition is given in Appendix A.2.

In our protocols we will be showing improved zero knowledge proofs of plaintext knowledge for lattice-based encryption schemes, as well as showing how to link messages being encrypted by these schemes to Pedersen commitments. Our improved proof of knowledge protocol will work for any Ring-LWE based scheme where the basic encryption operation consists of taking public key polynomial(s) a_i and computing the ciphertext(s) $b_i = a_i s + e_i$ where s and e_i are polynomials with small norms. Examples of such schemes include the ring-version of the “dual” encryption scheme [20], the “two element” scheme of Lyubashevsky et al. [28], and the NTRU encryption scheme [22,34].

For simplicity, in this paper we will only be working over the rings $R = \mathbb{Z}[X]/\langle X^n + 1 \rangle$ and $R_q = R/qR$, for some prime q . Also for simplicity, we will use NTRU as our encryption scheme because its ciphertext has only one element and is therefore simpler to describe in protocols. The NTRU scheme was first proposed by Hoffstein et al. [22], and we will be using a modification of it due to Stehlé and Steinfeld [34].

Definition 2.2. The discrete Normal distribution on \mathbb{Z}^m centered at v with standard deviation σ is defined by the density function $D_{v,\sigma}^m(x) = \rho_{v,\sigma}^m(x)/\rho_\sigma(\mathbb{Z}^m)$, where $\rho_{v,\sigma}^m(x) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^m e^{-\frac{\|x-v\|^2}{2\sigma^2}}$ is the continuous normal distribution on \mathbb{R}^m and $\rho_\sigma(\mathbb{Z}^m) = \sum_{z \in \mathbb{Z}^m} \rho_{0,\sigma}^m(z)$ is the scaling factor required to obtain a probability distribution.

In abuse of notation we will sometimes write $u \stackrel{\$}{\leftarrow} D_{v,\sigma}$ instead of $u \stackrel{\$}{\leftarrow} D_{v,\sigma}^n$ for a polynomial $u \in R_q$ if there is no risk of confusion.

In the following, let p be a prime less than q and $\sigma, \alpha \in \mathbb{R}$.

Message space. The message space is given by $\mathcal{M} = \{y \in R : \|y\|_\infty < p\}$

KeyGen. Sample f', g from $D_{\mathbb{Z}^n, \sigma}$, set $f = pf' + 1$, and resample, if $f \bmod q$ or $g \bmod q$ are not invertible. Output the public key $h = pg/f$ and the secret key f . Note here that h is invertible.

Encrypt. To encrypt a message $m \in \mathcal{M}$, set $s, e \stackrel{\$}{\leftarrow} D_\alpha$ and return the ciphertext $c = hs + pe + m \in R_q$.

Decrypt. To decrypt c with secret key f , compute $c' = fc \in R_q$ and output $m' = c' \bmod p$.

If the value of σ is large enough (approximately $\tilde{O}(n\sqrt{q})$), then g/f is uniformly random in R_q [34], and the security of the above scheme is based on the Ring-LWE problem. For smaller values of σ , however, the scheme is more efficient and can be based on the assumption that $h = g/f$ is indistinguishable from uniform. This type of assumption, while not based on any worst-case lattice problem, has been around since the introduction of the original NTRU scheme over fifteen years ago. Our protocol works for either instantiation.

To obtain group signatures, we will need our encryption scheme to additionally be a commitment scheme. In other words, there should not be more than one way to obtain the same ciphertext. For the NTRU encryption scheme, this will require that we work over a modulus q such that the polynomial $X^n + 1$ splits into two irreducible polynomials of degree $n/2$.

Lemma 2.3. Suppose that $q = 3 \bmod 8$ and let $\#S, \#E$, and $\#M$ be the domain sizes of the parameters s, e , and m in the ciphertext $c = hs + pe + m$. Additionally suppose that for all $m \in \mathcal{M}$, $\|m\|_\infty < p/2$. Then the probability that for a random h , there exists a ciphertext that can be obtained in two ways is at most $\frac{(2\#M+1) \cdot (2\#S+1) \cdot (2\#E+1)}{q^{n/2}}$.

The proof can be found in Appendix C.1. Note that the above lemma applies to NTRU public keys h that are uniformly random. If $h = pg/f$ is not random, then the ability to come up with two plaintexts for the same ciphertext would constitute a distinguisher for the assumed pseudorandomness of h .

2.4 Rejection Sampling

For a protocol to be zero-knowledge, the prover's responses must not depend on its secret inputs, as otherwise they could not be simulated. However, in our protocols, the prover's response will be from a discrete normal distribution which is shifted depending on the secret key. To correct for this, we employ rejection sampling [26,27], where a potential response is only output with a certain probability, and otherwise the protocol is aborted.

Informally, the following theorem states that for sufficiently large σ the rejection sampling procedure outputs results that are independent of the secret. The technique only requires a constant number of iterations before a value is output, and furthermore the output is also statistically close for every secret v with norm at most T . For concrete parameters we refer to the original work of Lyubashevsky [27, Theorem 4.6].

Theorem 2.4. Let V be a subset of \mathbb{Z}^ℓ in which all elements have norms less than T , and let h be a probability distribution over V . Then, for any constant M , there exists a $\sigma = \tilde{\Theta}(T)$ such that the output distributions of the following algorithms A, F are statistically close:

$$\begin{aligned} \text{A} : v \stackrel{\$}{\leftarrow} h; \quad z \stackrel{\$}{\leftarrow} D_{v,\sigma}^\ell; \quad \text{output } (z, v) \text{ with probability } \min\left(D_{\sigma(z)}^\ell / (MD_{v,\sigma}^\ell(z)), 1\right) \\ \text{F} : v \stackrel{\$}{\leftarrow} h; \quad z \stackrel{\$}{\leftarrow} D_{0,\sigma}^\ell; \quad \text{output } (z, v) \text{ with probability } 1/M \end{aligned}$$

Moreover, the probability that A outputs something is exponentially close to that of F, i.e., $1/M$.

2.5 Zero-Knowledge Proofs and Σ -Protocols

On a high level, a zero-knowledge proof of knowledge (ZKPoK) is a two party protocol between a *prover* and a *verifier*, which allows the former to convince the latter that it knows some secret piece of information, without revealing anything about the secret apart from what the claim itself already reveals. For a formal definition we refer to Bellare and Goldreich [3].

The ZKPoKs constructed in this paper will be instantiations of the following definition, which is a straightforward generalization of the standard notion of Σ -protocols [13,15]:

Definition 2.5. *Let (P, V) be a two-party protocol, where V is PPT, and let $\mathcal{R}, \mathcal{R}'$ be a binary relation such that $\mathcal{R} \subseteq \mathcal{R}'$. Then (P, V) is called a Σ' -protocol for $\mathcal{R}, \mathcal{R}'$ with challenge set \mathcal{C} , public input y and private input w , if and only if it satisfies the following conditions:*

- *Three-move form: The protocol is of the following form: The prover P computes a commitment τ and sends it to V . The verifier V then draws a challenge $c \xleftarrow{\$} \mathcal{C}$ and sends it to P . The prover sends a response s to the verifier. Depending on the protocol transcript (τ, c, s) , the verifier finally accepts or rejects the proof. The protocol transcript (τ, c, s) is called accepting, if the verifier accepts the protocol run.*
- *Completeness: Whenever $(y, w) \in \mathcal{R}$, the verifier V accepts with probability at least $1 - \alpha$.*
- *Special soundness: There exists a PPT algorithm E (the knowledge extractor) which takes two accepting transcripts $(\tau, c', s'), (t, c'', s'')$ satisfying $c' \neq c''$ as inputs, and outputs w' such that $(y, w') \in \mathcal{R}'$.*
- *Special honest-verifier zero-knowledge (HVZK): There exists a PPT algorithm S (the simulator) taking $y \in L(\mathcal{R})$ and $c \in \mathcal{C}$ as inputs, that outputs (τ, s) so that the triple (τ, c, s) is indistinguishable from an accepting protocol transcript generated by a real protocol run.*

This definition differs from the standard definition of Σ -protocols in two ways. First, we allow the honest prover to fail in at most an α -fraction of all protocol runs, whereas the standard definition requires perfect completeness, i.e., $\alpha = 0$. However, this relaxation is crucial in our construction that is based on rejection sampling [26,27], where the honest prover sometimes has to abort the protocol to achieve zero-knowledge. Second, we introduce a second relation $\mathcal{R}' \supseteq \mathcal{R}$, such that provers knowing a witness in \mathcal{R} are guaranteed privacy, but the verifier is only ensured that the prover knows a witness for \mathcal{R}' . This has already been used in [1] and informally also in, e.g., [19,16]. If the *soundness gap* between \mathcal{R} and \mathcal{R}' is sufficiently small, the implied security guarantees are often enough for higher-level applications.

We want to stress that previous results showing that a Σ -protocol is always also an honest-verifier ZKPoK with knowledge error $1/|\mathcal{C}|$ directly carry over to the modified definition whenever $1 - \alpha > 1/|\mathcal{C}|$. Zero-knowledge against arbitrary verifiers can be achieved by applying standard techniques such as Damgård et al. [17,14].

Finally, it is a well known result that negligible knowledge and completeness errors in λ can be achieved, e.g., by running the protocol λ times in parallel and accepting if and only if at least $\lambda(1 - \alpha)/2$ transcripts were valid, if there exists a constant c such that $(1 - \alpha)/2 > 1/|\mathcal{C}| + c$.

Some of the Σ -protocol presented in this paper will further satisfy the following useful properties:


- *Quasi-unique responses:* No PPT adversary A can output (y, τ, c, s, s') with $s \neq s'$ such that $V(y, \tau, c, s) = V(y, \tau, c, s') = \text{accept}$.
- *High-entropy commitments:* For all $(y, w) \in \mathcal{R}$ and for all τ , the probability that an honestly generated commitment by P takes on the value τ is negligible.

3 Proving Knowledge of Ring-LWE Secrets

Before stating our basic protocol for efficiently proving knowledge of short s, e such that $2y = as + e$, we prove a technical lemma that is at the heart of the knowledge extractor of our protocols.

3.1 A Technical Lemma

The following lemma guarantees that certain binomials in $\mathbb{Z}[X]/(X^n + 1)$ can be inverted, and their inverses have only small coefficients.



Protocol 3.2: Proof of knowledge of LWE-secrets s, e such that $y = as + e$.

Lemma 3.1. *Let n be a power of 2 and let $0 < i, j < 2n - 1$. Then $2(X^i - X^j)^{-1} \bmod (X^n + 1)$ only has coefficients in $\{-1, 0, 1\}$.*

Proof. Without loss of generality, assume that $j > i$. Using that $X^n = -1 \bmod (X^n + 1)$ we have that $2(X^i - X^j)^{-1} = -2X^{n-i}(1 - X^{j-i})^{-1}$. It is therefore sufficient to prove the claim for $i = 0$ only.

Now remark that, for every $k \geq 1$ it holds that: $(1 - X^j)(1 + X^j + X^{2j} + \dots + X^{(k-1)j}) = 1 - X^{kj}$.

Let us write $j = 2^{j'} j''$, with j'' a positive odd integer and $0 \leq j' \leq \log_2(n)$, and let us choose $k = 2^{\log_2(n) - j'}$ (recall that n is a power of 2). We then have $jk = nj''$, and $X^{kj} = (-1)^{j''} = -1 \bmod (X^n + 1)$, hence $1 - X^{kj} = 2 \bmod (X^n + 1)$. Therefore, we have

$$\begin{aligned} 2(1 - X^j)^{-1} &= 1 + X^j + X^{2j} + \dots + X^{(k-1)j} \bmod (X^n + 1) \\ &= 1 \pm X^{j \bmod n} \pm X^{2j \bmod n} \pm \dots \pm X^{(k-1)j \bmod n} \bmod (X^n + 1). \end{aligned}$$

Finally, in this equation, no two exponents are equal, since otherwise this means that n divides jk' with $1 \leq k' < k$, which is impossible by definition of k . \square

3.2 The Protocol

We next present our basic protocol. Let therefore be $y = as + e$, where the LWE-secrets $s, e \leftarrow^{\$} D_\alpha$ are chosen from a discrete Gaussian distribution with standard deviation α . Protocol 3.2 now allows a prover to convince a verifier that it knows short s' and e' such that $2y = as' + e'$, i.e., the verifier is ensured that the prover knows short secrets for twice the public input. Here, by short we mean the following: An honest prover will always be able to convince the verifier whenever $\|s\|, \|e\| \leq \tilde{\mathcal{O}}(\sqrt{n}\alpha)$, which is the case with overwhelming probability if they were generated honestly. On the other hand, the verifier is guaranteed that the prover knows LWE-secrets with norm at most $\tilde{\mathcal{O}}(n^2\alpha)$. This soundness gap on the size of the witnesses is akin to those in, e.g., [16, 1].

To be able to simulate aborts when proving the zero-knowledge property of the protocol, we must not send the prover's first message in the plain, but commit to it and later open it in the last round of the Σ -protocol. We therefore make use of an auxiliary commitment scheme (aCSetup, aCommit, aCOpen), and assume that honestly generated commitment parameters are given as common input to both parties. We do not make any assumptions on the auxiliary commitment scheme. However, if it is computationally binding, the resulting protocol is only sound under the respective assumption, and similarly if it is computationally hiding. For simplicity, the reader may just think of the scheme as a random oracle.

Theorem 3.3. *Protocol 3.2 is an HVZK Σ' -protocol for the following relations:*

$$\begin{aligned} \mathcal{R} &= \{((s, e), (a, y)) : y = as + e \quad \wedge \quad \|s\|, \|e\| \leq \tilde{\mathcal{O}}(\sqrt{n}\alpha)\} \\ \mathcal{R}' &= \{((s, e), (a, y)) : 2y = as + e \quad \wedge \quad \|s\|, \|e\| \leq \tilde{\mathcal{O}}(n^2\alpha)\} \end{aligned}$$

The protocol has a knowledge error of $1/(2n)$, a completeness error of $1 - 1/M^2$, and high-entropy commitments.

Proof. We need to prove the properties of Definition 2.5.

Completeness. First note that by Theorem 2.4, the prover will respond with probability $1/M$ for each secret, i.e., with an overall probability of $1/M^2$. If the prover does not abort, we have that:

$$a\mathbf{s}_s + \mathbf{s}_e = a(\mathbf{r}_s + X^c s) + (\mathbf{r}_e + X^c e) = X^c(as + e) + (a\mathbf{r}_s + \mathbf{r}_e) = X^c y + \mathbf{t}.$$

For the norms we have that $\|\mathbf{s}_s\| \leq \|\mathbf{r}_s\| + \|s\| \leq \tilde{O}(n\alpha)$ with overwhelming probability, as the standard deviations of \mathbf{r}_s is $\tilde{O}(\sqrt{n}\alpha)$, and similarly for \mathbf{s}_e .

Honest-verifier zero-knowledge. Given a challenge c , the simulator outputs $(\text{aCommit}(0), c, \perp)$ with probability $1/M^2$. With probability $1 - 1/M^2$, the simulator S proceeds as follows: It chooses $\mathbf{s}_s, \mathbf{s}_e \xleftarrow{\$} D_{\tilde{O}(\sqrt{n}\alpha)}$, and computes $\mathbf{t} = a\mathbf{s}_s + \mathbf{s}_e - X^c y$, and $(c_{\text{aux}}, d_{\text{aux}}) \xleftarrow{\$} \text{aCommit}(\mathbf{t})$. Finally, S outputs $(c_{\text{aux}}, c, (\mathbf{t}, d_{\text{aux}}, (\mathbf{s}_s, \mathbf{s}_e)))$.

It follows from Theorem 2.4 that if no abort occurs the distribution of $\mathbf{s}_e, \mathbf{s}_s$ does not depend on s, e , and thus simulated and real protocol transcripts are indistinguishable. In case that an abort occurs, the indistinguishability follows from the hiding property of aCommit and the fact that aborts are equally likely for every c .

Special soundness. Assume that we are given $(c_{\text{aux}}, c', (\mathbf{t}', d'_{\text{aux}}, (\mathbf{s}'_s, \mathbf{s}'_e)))$ and $(c_{\text{aux}}, c'', (\mathbf{t}'', d''_{\text{aux}}, (\mathbf{s}''_s, \mathbf{s}''_e)))$ passing the checks performed by the verifier. From the binding property of the auxiliary commitment scheme we get that $\mathbf{t}' = \mathbf{t}'' =: \mathbf{t}$. Now, by subtracting the verification equations we get: $(X^{c'} - X^{c''})y = a(\mathbf{s}'_s - \mathbf{s}''_s) + (\mathbf{s}'_e - \mathbf{s}''_e)$. Multiplying by $2(X^{c'} - X^{c''})^{-1}$ yields:

$$2y = a \frac{2(\mathbf{s}'_s - \mathbf{s}''_s)}{X^{c'} - X^{c''}} + \frac{2(\mathbf{s}'_e - \mathbf{s}''_e)}{X^{c'} - X^{c''}} =: a\hat{s} + \hat{e}.$$

Furthermore, we get that $\|\hat{s}\| \leq \|\mathbf{s}'_e - \mathbf{s}''_e\| \sqrt{n} \left\| \frac{2}{X^{c'} - X^{c''}} \right\| \leq \tilde{O}(n^2\alpha)$, where in the second inequality we used Lemma 3.1, and similarly for \hat{e} .

High-entropy commitments. This directly follows from the security of the auxiliary commitment scheme. \square

By Section 2.5, both the completeness and the knowledge error can be made negligible if $n > M^2$.

4 Proving Equality among Classical and Lattice-Based Primitives

In the following we show how our basic protocol from Section 3 can be used to link number-theory and lattice-based primitives via zero-knowledge proofs of knowledge. We exemplify this by showing how to prove that the messages contained in Pedersen commitments correspond to the plaintext in an encryption under the secure version of NTRU. We want to stress that in particular the choice of the encryption scheme is arbitrary, and it is easy to exchange it against other schemes, including standard NTRU [22] or Ring-LWE encryption [28].

Let $y = hs + pe + m \in R_q$ be the NTRU encryption of a message $m \in \{0, 1\}^n$, and let $p > 2n^2$ be coprime with q . Let further \tilde{g}, \tilde{h} be a Pedersen commitment parameters, cf. Section 2.2, and let $\widetilde{cmt}_i = m_i \tilde{g} + r_i \tilde{h}$ for $i = 0, \dots, n-1$ be commitments to coefficients of m , where the order of \tilde{g} and \tilde{h} is $\tilde{q} > 2n^2$.

Then Protocol 4.1 can be used to prove, in zero-knowledge, that the commitments and the ciphertext are broadly well-formed and consistent, i.e., contain the same message. More precisely, the protocol guarantees the verifier that the prover knows the plaintext encrypted in $2y$, and that the coefficients of the respective message are all smaller than p . Furthermore, it shows that the messages are the same that are contained in $2\widetilde{cmt}_i$, i.e., $2y$ and the $2\widetilde{cmt}_i$ are consistent.

Prover	Verifier
$\mathbf{r}_s, \mathbf{r}_e \xleftarrow{\$} D_{\tilde{\mathcal{O}}(\sqrt{n}\alpha)}$ $\mathbf{r}_m \xleftarrow{\$} D_{\tilde{\mathcal{O}}(\sqrt{n})}$ $\mathbf{r}_{r,i} \xleftarrow{\$} \mathbb{Z}_q$ for $i = 0, \dots, n-1$ $\mathbf{t} = h\mathbf{r}_s + p\mathbf{r}_e + \mathbf{r}_m$ $\tilde{\mathbf{t}}_i = \mathbf{r}_{m,i}\tilde{g} + \mathbf{r}_{r,i}\tilde{h}$ for $i = 0, \dots, 2n-1$ $(c_{\text{aux}}, d_{\text{aux}}) = \text{aCommit}(\mathbf{t}, (\tilde{\mathbf{t}}_i)_{i=0}^{n-1})$	$c \xleftarrow{\$} \mathcal{C} = \{0, \dots, 2n-1\}$
$\mathbf{s}_s = \mathbf{r}_s + X^c \mathbf{s}$ $\mathbf{s}_e = \mathbf{r}_e + X^c \mathbf{e}$ $\mathbf{s}_m = \mathbf{r}_m + X^c \mathbf{m}$ $\mathbf{s}_r = \mathbf{r}_r + (r_0, \dots, r_{n-1}) \ll c$	$X^c y + \mathbf{t} \stackrel{?}{=} h\mathbf{s}_s + p\mathbf{s}_e + \mathbf{s}_m$ $(\widetilde{cmt}_0, \dots, \widetilde{cmt}_{n-1}) \ll c + (\tilde{\mathbf{t}}_0, \dots, \tilde{\mathbf{t}}_{n-1}) \stackrel{?}{=} \mathbf{s}_m \tilde{g} + \mathbf{s}_r \tilde{h}$ $\text{aCOpen}((\mathbf{t}, \tilde{\mathbf{t}}_0, \dots, \tilde{\mathbf{t}}_{n-1}), c_{\text{aux}}, d_{\text{aux}}) \stackrel{?}{=} \text{accept}$ $\ \mathbf{s}_s\ , \ \mathbf{s}_e\ \leq \tilde{\mathcal{O}}(n\alpha)$ $\ \mathbf{s}_m\ \leq \tilde{\mathcal{O}}(n)$
$\text{accept with probability } \frac{D_{\tilde{\mathcal{O}}(\sqrt{n}\alpha)}(\mathbf{s}_e)}{MD_{e, \tilde{\mathcal{O}}(\sqrt{n}\alpha)}(\mathbf{s}_e)} \frac{D_{\tilde{\mathcal{O}}(\sqrt{n}\alpha)}(\mathbf{s}_s)}{MD_{s, \tilde{\mathcal{O}}(\sqrt{n}\alpha)}(\mathbf{s}_s)} \frac{D_{\tilde{\mathcal{O}}(\sqrt{n})}(\mathbf{s}_m)}{MD_{m, \tilde{\mathcal{O}}(\sqrt{n})}(\mathbf{s}_m)}$ $(\mathbf{t}, (\tilde{\mathbf{t}}_i)_{i=0}^{n-1}), d_{\text{aux}}, (\mathbf{s}_s, \mathbf{s}_e, \mathbf{s}_m, \mathbf{s}_r)$	

Protocol 4.1: Proof that Pedersen commitments and NTRU encryption contain the same plaintext.

Theorem 4.2. *Protocol 4.1 is an HVZK Σ' -protocol for the following relations:*

$$\mathcal{R} = \left\{ ((m, s, e, (r_i)_{i=0}^{n-1}), (\tilde{g}, \tilde{h}, (\widetilde{cmt}_i)_{i=0}^{n-1}, h, p, y)) : y = hs + pe + m \wedge \bigwedge_{i=0}^{n-1} \widetilde{cmt}_i = m_i \tilde{g} + r_i \tilde{h} \right. \\ \left. \wedge \|m\|_\infty \leq 1 \wedge \|s\|, \|e\| \leq \tilde{\mathcal{O}}(\sqrt{n}\alpha) \right\},$$

$$\mathcal{R}' = \left\{ ((m, s, e, (r_i)_{i=0}^{n-1}), (\tilde{g}, \tilde{h}, (\widetilde{cmt}_i)_{i=0}^{n-1}, h, p, y)) : 2y = hs + pe + m \wedge \bigwedge_{i=0}^{n-1} 2\widetilde{cmt}_i = m_i \tilde{g} + r_i \tilde{h} \right. \\ \left. \wedge \|m\|_\infty \leq 2n^2 \wedge \|s\|, \|e\| \leq \tilde{\mathcal{O}}(n^2\alpha) \right\}$$

The protocol has a knowledge error of $1/(2n)$, and a completeness error of $1 - 1/M^3$.

Furthermore, if for the auxiliary commitment scheme a commitment does not only bind the user to the message, but also to the opening information, the protocol has quasi-unique responses and high-entropy commitments.

A detailed proof is given in Appendix C.2. By the remark in Section 2.5, both the completeness and the knowledge error can be made negligible if $n > M^3$.

5 Application to Group Signatures

We next show how Protocol 4.1 can be used to construct a group signature scheme whose signature size is only logarithmic in the number of group members. As discussed in Section 1.3, our scheme will be private under lattice assumptions, but only unforgeable under the discrete logarithm assumption, which is sufficient for many applications such as e-voting, where long-term privacy is vital.

However, before presenting the actual signature scheme, we will prove secure a generic construction that is folklore, and underlies several direct constructions in the literature [2,8] and a variant was described explicitly by Chase and Lysyanskaya [11]. However, we provide the first formal proof that this construction actually yields secure group signature schemes.

The resulting construction will satisfy the following definition of group signatures providing full (CCA) anonymity, and is a variant of the definition put forth by Bellare et al. [4].

<p>Experiment $\mathbf{Exp}_A^{\text{anon}-b}(\lambda)$:</p> <p>$(gpk, gok, \mathbf{gsk}) \xleftarrow{\\$} \text{GKG}(1^\lambda, 1^n)$</p> <p>$(st, i_0^*, i_1^*, m^*) \xleftarrow{\\$} A^{\text{GOpen}(gok, \cdot, \cdot)}((gpk, \mathbf{gsk}), \varepsilon)$</p> <p>$\sigma^* \xleftarrow{\\$} \text{GSign}(\mathbf{gsk}[i_b], m^*)$</p> <p>$b' \xleftarrow{\\$} A^{\text{GOpen}(gok, \cdot, \cdot)}(\sigma^*, st)$</p> <p>If $(m^*, \sigma^*) \notin \mathcal{Q}_{\text{GOpen}}$ return b' else return 0</p>	<p>Experiment $\mathbf{Exp}_A^{\text{trace}}(\lambda)$:</p> <p>$(gpk, gok, \mathbf{gsk}) \xleftarrow{\\$} \text{GKG}(1^\lambda, 1^n)$</p> <p>$(m, \sigma) \xleftarrow{\\$} A^{\text{GSign}(\mathbf{gsk}[\cdot, \cdot], \mathbf{gsk}[\cdot])}(gpk, gok)$</p> <p>$i \xleftarrow{\\$} \text{GOpen}(gok, m, \sigma)$</p> <p>If $\text{GVerify}(gpk, m, \sigma) = 1 \wedge i \notin \mathcal{Q}_{\mathbf{gsk}}$ $\wedge (i, m) \notin \mathcal{Q}_{\text{GSign}}$ then return 1 else return 0</p>
--	---

Fig. 1. The anonymity (left) and traceability (right) experiments for group signatures. The sets $\mathcal{Q}_{\text{GOpen}}$, $\mathcal{Q}_{\text{GSign}}$, $\mathcal{Q}_{\mathbf{gsk}}$ contain all queries (m, σ) , (i, m) , and i that A submitted to its GOpen , GSign , and \mathbf{gsk} oracles, respectively.

Definition 5.1. A group signature scheme is a tuple $(\text{GKG}, \text{GSign}, \text{GVerify}, \text{GOpen})$ where:

- On input $1^\lambda, 1^n$, the key generation algorithm GKG outputs a group public key gpk , an opening key gok , and a vector of n signing keys \mathbf{gsk} where $\mathbf{gsk}[i]$ is given to user $i \in \{1, \dots, n\}$.
- On input $gsk = \mathbf{gsk}[i]$ and message $m \in \mathcal{M}$, the signing algorithm GSign outputs a group signature σ .
- On input gpk, m, σ , the verification algorithm GVerify outputs `accept` or `reject`.
- On input gok, m, σ , the opening algorithm GOpen outputs the identity of the purported signer $i \in \{1, \dots, n\}$ or \perp to indicate failure.

The algorithms satisfy the following properties:

- **Correctness:** Verification accepts whenever keys and signatures are honestly generated, i.e., for all $\lambda, n \in \mathbb{N}$, all $i \in \{1, \dots, n\}$, and all $m \in \mathcal{M}$

$$\Pr[\text{GVerify}(gpk, m, \sigma) = \text{accept} : (gpk, gok, \mathbf{gsk}) \xleftarrow{\$} \text{GKG}(1^\lambda, 1^n), \sigma \xleftarrow{\$} \text{GSign}(\mathbf{gsk}[i], m)] = 1.$$

- **Anonymity:** One cannot tell which signer generated a particular signature, even when given access to an opening oracle. Referring to Figure 1, for all PPT A there exists a negligible function negl such that

$$|\Pr[\mathbf{Exp}_A^{\text{anon}-0}(\lambda) = 1] - \Pr[\mathbf{Exp}_A^{\text{anon}-1}(\lambda) = 1]| \leq \text{negl}(\lambda).$$

- **Traceability:** One cannot generate a signature that cannot be opened or that opens to an honest user. Referring to Figure 1, for all PPT A there exists a negligible function negl such that

$$\Pr[\mathbf{Exp}_A^{\text{trace}}(\lambda)] \leq \text{negl}(\lambda).$$

5.1 Building Blocks

The construction is based on weakly unforgeable standard signatures, and signature proofs of knowledge. In the following, we recap the respective definitions.

Informally, a signature scheme is a triple $(\text{SKG}, \text{SSign}, \text{SVerify})$, where SKG generates a signing/verification key pair (ssk, spk) , SSign can be used to sign a message m using the signing key, and SVerify can be used to check the validity of a signature only using the public verification key. It should hold that honestly computed signatures are always valid, and that no adversary can come up with a valid signature on a new message after having received signatures on message that he chose before obtaining spk . A formal definition can be found in Appendix A.3.

We adapt definitions of Chase-Lysyanskaya to allow for signatures in the random-oracle model that are simulated by programming the random oracle and extracted through rewinding. We also add definition of simulation soundness, meaning that adversary cannot produce new signatures for false statements even after seeing simulated signatures on arbitrary statements.

Definition 5.2. A signature of knowledge scheme for a language $\mathcal{L} \subseteq \{0, 1\}^*$ with witness relation $R_{\mathcal{L}}$ is a tuple $(\text{SoKSetup}, \text{SoKSign}, \text{SoKVerify}, \text{SoKSim})$ where:

- On input 1^λ , the setup algorithm SoKSetup outputs common parameters sokp .
- On input sokp, x, w such that $(x, w) \in R_{\mathcal{L}}$ and message $m \in \mathcal{M}$, the signing algorithm SoKSign outputs a signature of knowledge sok .
- On input $\text{sokp}, x, m, \text{sok}$, the verification algorithm SoKVerify outputs accept or reject.
- The stateful simulation algorithm SoKSim can be called in three modes. When called as $(\text{sokp}, st) \xleftarrow{\$} \text{SoKSim}(\text{setup}, 1^\lambda, \varepsilon)$, it produces simulated parameters sokp , possibly keeping a trapdoor in its internal state st . When run as $(h, st') \xleftarrow{\$} \text{SoKSim}(\text{ro}, Q, st)$, it produces a response h for a random oracle query Q . When run as $(\text{sok}, st') \xleftarrow{\$} \text{SoKSim}(\text{sign}, x, m, st)$, it produces a simulated signature of knowledge sok without using a witness.

For ease of notation, let $\text{StpSim}(1^\lambda)$ be the algorithm that returns the first part of $\text{SoKSim}(\text{setup}, 1^\lambda, st)$, let $\text{ROSim}(Q)$ be the algorithm that returns the first part of $\text{SoKSim}(\text{ro}, Q, st)$, let $\text{SSim}(x, w, m)$ be the algorithm that returns the first part of $\text{SoKSim}(\text{sign}, x, m, st)$ if $(x, w) \in R_{\mathcal{L}}$ and returns \perp otherwise, and let $\text{SSim}'(x, m)$ be the algorithm that returns the first part of $\text{SoKSim}(\text{sign}, x, m, st)$ without checking language membership. The experiment keeps a single synchronized state for SoKSim across all invocations of these derived algorithms.

The algorithms satisfy the following properties:

- **Correctness:** Verification accepts whenever parameters and signatures are correctly generated, i.e., for all $\lambda \in \mathbb{N}$, all $(x, w) \in R_{\mathcal{L}}$, and all $m \in \mathcal{M}$, there exists a negligible function negl such that

$$\Pr \left[\begin{array}{l} \text{SoKVerify}(\text{sokp}, x, m, \text{sok}) = \text{reject} : \\ \text{sokp} \xleftarrow{\$} \text{SoKSetup}(1^\lambda), \text{sok} \xleftarrow{\$} \text{SoKSign}(\text{sokp}, x, w, m) \end{array} \right] \leq \text{negl}(1^\lambda).$$

- **Simulatability:** No adversary can distinguish whether it is interacting with a real random oracle and signing oracle, or with their simulated versions. Formally, for all PPT A there exists a negligible function negl such that

$$\left| \Pr[b = 1 : \text{sokp} \xleftarrow{\$} \text{StpSim}(1^\lambda), b \xleftarrow{\$} A^{\text{ROSim}(\cdot), \text{SSim}(\cdot, \cdot, \cdot)}(\text{sokp})] - \Pr[b = 1 : \text{sokp} \xleftarrow{\$} \text{SoKSetup}(1^\lambda), b \xleftarrow{\$} A^{H(\cdot), \text{SoKSign}(\text{sokp}, \cdot, \cdot)}(\text{sokp})] \right| \leq \text{negl}(\lambda).$$

- **Extractability:** The only way to produce a valid signature of knowledge is by knowing a witness. Formally, for all PPT A there exists an extractor SoKExt_A and a negligible function negl such that

$$\Pr \left[\begin{array}{l} \text{SoKVerify}(\text{sokp}, x, m, \text{sok}) = \text{accept} \wedge (x, w, m) \notin \mathcal{Q} \wedge (x, w) \notin R_{\mathcal{L}} : \\ \text{sokp} \xleftarrow{\$} \text{StpSim}(1^\lambda; \rho_S), (x, m, \text{sok}) \xleftarrow{\$} A^{\text{ROSim}(\cdot), \text{SSim}(\cdot, \cdot, \cdot)}(\text{sokp}; \rho_A), \\ w \xleftarrow{\$} \text{SoKExt}_A(\text{sokp}, x, m, \text{sok}, \rho_S, \rho_A) \end{array} \right] \leq \text{negl}(\lambda),$$

where \mathcal{Q} is the set of queries (x, w, m) that A submitted to its SSim oracle.

- **Simulation-soundness:** No adversary can produce a new signature on a false statement, even after seeing a signature on an arbitrary statement. Formally, for all PPT A there exists a negligible function negl such that

$$\Pr \left[\begin{array}{l} \text{SoKVerify}(\text{sokp}, x, m, \text{sok}) = \text{accept} \wedge (x', m', \text{sok}') \neq (x, m, \text{sok}) \wedge x \notin \mathcal{L} : \\ \text{sokp} \xleftarrow{\$} \text{StpSim}(1^\lambda), (x, m, st) \xleftarrow{\$} A^{\text{ROSim}(\cdot)}(\text{sokp}), \text{sok} \xleftarrow{\$} \text{SSim}'(x, m), \\ (x', m', \text{sok}') \xleftarrow{\$} A^{\text{ROSim}(\cdot)}(\text{sok}, st) \end{array} \right] \leq \text{negl}(\lambda).$$

5.2 Generic Construction

A folklore construction of group signatures is to have a user's signing key be a standard signature on his identity i , and to have a group signature on message m be an encryption of his identity together with a signature of knowledge on m that the encrypted identity is equal to the identity in his signing key. The construction appeared implicitly [2,8] or explicitly [11] in the literature, but was never proved secure.

To obtain full anonymity, this generic construction would probably require CCA security from encryption scheme, but our NTRU variant is only semantically secure. We could apply a generic CCA-yielding transformation using random oracles or non-interactive zero-knowledge proofs of knowledge (NIZK), but this would make the signature of knowledge hopelessly inefficient. Instead, we take inspiration from the Naor-Yung construction [29,32] by using a semantically secure scheme to encrypt the user's identity twice under two different public keys and letting the signature of knowledge prove that both ciphertexts encrypt the same plaintext. Moreover, our proof system for LWE encryption in Section 3 does not quite prove knowledge of (m, ρ) such that $C = \text{Enc}(epk, m, \rho)$, but rather of (m', ρ') such that $2C = \text{Enc}(epk, m', \rho')$. We therefore give a generic construction that deviates slightly from the general idea, but that is sufficient and that we can efficiently instantiate with our protocol from Section 3.

Let $(\text{EncKG}, \text{Enc}, \text{Dec})$ be an encryption scheme with message space \mathcal{M} and a homomorphism (Φ, φ, ψ) such that φ is an injective function and such that for all epk and for all $m \in \mathcal{ID} \subseteq \mathcal{M}$

$$C = \text{Enc}(epk, m; \rho) \Rightarrow \Phi(C) = \text{Enc}(epk, \varphi(m); \psi(\rho)) .$$

Let $(\text{SKG}, \text{SSign}, \text{SVerify})$ be a standard signature scheme, and let $(\text{SoKSetup}, \text{SoKSign}, \text{SoKVerify}, \text{SoKSim})$ be a signature of knowledge scheme for the language \mathcal{L} with witness relationship

$$R_{\mathcal{L}} = \{((spk, epk_1, epk_2, C_1, C_2), (i, sig, \rho_1, \rho_2)) : \\ \text{SVerify}(spk, i, sig) = \text{accept} \wedge \Phi(C_1) = \text{Enc}(epk_1, \varphi(i); \rho_1) \wedge \Phi(C_2) = \text{Enc}(epk_2, \varphi(i); \rho_2)\} ,$$

consider the following group signature scheme with user identities $i \in \mathcal{ID}$:

- $\text{GKG}(1^\lambda, 1^n)$: The group manager generates signature keys $(spk, ssk) \xleftarrow{\$} \text{SKG}(1^\lambda)$, encryption keys $(epk_1, esk_1) \xleftarrow{\$} \text{EncKG}(1^\lambda)$, $(epk_2, esk_2) \xleftarrow{\$} \text{EncKG}(1^\lambda)$, and parameters $sokp \xleftarrow{\$} \text{SoKSetup}(1^\lambda)$. He computes $\text{gsk}[i] \xleftarrow{\$} \text{SSign}(ssk, i)$ for $i \in \mathcal{ID}$ and outputs $\text{gpk} = (spk, epk_1, epk_2, sokp)$, $\text{gok} = esk_1$, and gsk .
- $\text{GSign}(\text{gsk}, m)$: Signer i computes two ciphertexts $C_1 \leftarrow \text{Enc}(epk_1, i; \rho_1)$ and $C_2 \leftarrow \text{Enc}(epk_2, i; \rho_2)$, computes a signature of knowledge $sok \xleftarrow{\$} \text{SoKSign}(sokp, (spk, epk_1, epk_2, C_1, C_2), (i, sig, \psi(\rho_1), \psi(\rho_2)), m)$ and outputs $\sigma = (C_1, C_2, sok)$.
- $\text{GVerify}(\text{gpk}, m, \sigma)$: To verify a group signature, one checks that $\text{SoKVerify}(sokp, (spk, epk_1, epk_2, C_1, C_2), m, sok) = \text{accept}$.
- $\text{GOpen}(\text{gok}, m, \sigma)$: The opener decrypts $d \leftarrow \text{Dec}(esk_1, C_1)$ and returns $i \in \mathcal{ID}$ such that $\varphi(i) = d$ (by inverting φ if possible, otherwise by exhaustive search).

Theorem 5.3. *The group signature scheme sketched above is anonymous if the encryption scheme is semantically secure and perfectly complete, and the signature of knowledge scheme is simulatable and simulation-sound.*

Theorem 5.4. *The group signature scheme is traceable if the underlying signature scheme is weakly unforgeable, the encryption scheme is perfectly complete, and the signature of knowledge scheme is simulatable and extractable.*

The proofs of the last two theorems are omitted here and are given in Appendices C.3 and C.4.

5.3 Signatures of Knowledge from Σ -Protocols

We now show a construction of the required signatures of knowledge in the random-oracle model from a signature scheme and an encryption scheme with Σ -protocol proofs. More particularly, we require that the signature scheme can prove knowledge of a signature on a committed message, while the encryption scheme can prove that an encrypted plaintext is equal to a committed message.

Let $(\text{EncKG}, \text{Enc}, \text{Dec})$ be an encryption scheme with message space \mathcal{M} , let $(\text{CSetup}, \text{Commit}, \text{COpen})$ be a commitment scheme, and let (Φ, φ, ψ) and (Ψ, φ, χ) be homomorphisms as before such that for all $cpar$ s:

$$\text{COpen}(cpar, cmt, o) = \text{accept} \Rightarrow \text{COpen}(cpar, \Psi(cmt), \chi(o)) = \text{accept} .$$

Let (P_s, V_s, S_s) be a Σ -protocol for the language \mathcal{L}_s with

$$R_{\mathcal{L}_s} = \{((spk, cpars, cmt), (sig, m, o)) : \\ SVerify(spk, m, sig) = \text{accept} \wedge COpen(cpars, m, cmt, o) = \text{accept}\}$$

and let (P_e, V_e, S_e) be a Σ -protocol for the language \mathcal{L}_e with

$$R_{\mathcal{L}_e} = \{((epk, C, cpars, cmt), (m, \rho, o)) : \\ \Phi(C) = \text{Enc}(epk, m; \rho) \wedge COpen(cpars, m, \Psi(cmt), o) = \text{accept}\}.$$

Let \mathcal{C}_s and \mathcal{C}_e be the challenge spaces for these respective Σ -protocols, and let $H : \{0, 1\}^* \rightarrow \mathcal{C}_s \times \mathcal{C}_e$. Consider the following construction of a signature of knowledge scheme for the language \mathcal{L} :

- **SoKSetup** (1^λ) : Return $sokp = cpars \xleftarrow{\$} CSetup(1^\lambda)$.
- **SoKSign** $(sokp, x, w, m)$: Create a commitment $(cmt, o) \xleftarrow{\$} \text{Commit}(i, cpars)$. Compute the first round of one signature Σ -protocol and two encryption Σ -protocols: $(\mathfrak{t}_s, st_s) \xleftarrow{\$} P_s((spk, cpars, cmt), (sig, m, o))$ and $(\mathfrak{t}_j, st_j) \xleftarrow{\$} P_e((epk_j, C_j, cpars, cmt), (\varphi(i), \psi(\rho_j), \chi(o)))$ for $j = 1, 2$. Generate the challenges $(c_s, c_e) \leftarrow H(sp, cpars, cmt, epk_1, C_1, epk_2, C_2, \mathfrak{t}_s, \mathfrak{t}_1, \mathfrak{t}_2, m)$. Generate the responses $\mathfrak{s}_s \leftarrow P_s(c_s, st_s)$ and $\mathfrak{s}_j \leftarrow P_e(c_e, st_e)$ for $j = 1, 2$ and output the signature of knowledge $sok = (\mathfrak{t}_s, \mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{s}_s, \mathfrak{s}_1, \mathfrak{s}_2)$.
- **SoKVerify** $(sokp, x, m, sok)$: Recompute the challenges $(c_s, c_e) \leftarrow H(sp, cpars, cmt, epk_1, C_1, epk_2, C_2, \mathfrak{t}_s, \mathfrak{t}_1, \mathfrak{t}_2, m)$. Return **accept** if $V_s((spk, cpars, cmt), \mathfrak{t}_s, c_s, \mathfrak{s}_s) = \text{accept}$ and $V_e((epk_j, C_j, cpars, cmt), \mathfrak{t}_j, c_e, \mathfrak{s}_j) = \text{accept}$ for $j = 1, 2$. Otherwise, return **reject**.
- **SoKSim**: The simulation algorithm keeps in its state its random tape, an initially empty table HT to keep track of previous random-oracle queries, and a counter ctr initialized to zero. The simulator's random tape ρ includes random-oracle responses $h_1, \dots, h_{q_H+q_S} \xleftarrow{\$} \mathcal{C}_s \times \mathcal{C}_e$, where q_H and q_S are upper bounds on the number of random-oracle and signing queries that an adversary can make. When called as $\text{SoKSim}(\text{setup}, 1^\lambda, \varepsilon)$, it generates commitment parameters $cpars \xleftarrow{\$} CSetup(1^\lambda)$ and returns $(cpars, st = (\rho, HT, ctr, cpars))$. When run as $\text{SoKSim}(\text{ro}, Q, st)$, it checks whether the query Q was made before. If so, it returns $h_{HT[Q]}$. Otherwise, it increases the counter ctr , sets $HT[Q] \leftarrow ctr$, and returns h_{ctr} . When run as $\text{SoKSim}(\text{sign}, (spk, epk_1, epk_2, C_1, C_2), m, st)$, the simulator first creates a commitment $(cmt, o) \xleftarrow{\$} \text{Commit}(1, cpars)$. It then increases the counter ctr and parses h_{ctr} as (c_s, c_e) . It runs the simulators S_s, S_e to obtain simulated protocol transcripts $(\mathfrak{t}_s, \mathfrak{s}_s) \xleftarrow{\$} S_s((spk, cpars, cmt), c_s)$ and $(\mathfrak{t}_j, \mathfrak{s}_j) \xleftarrow{\$} S_e((epk_j, C_j, cpars, cmt), c_e)$ for $j = 1, 2$. If $HT[spk, cpars, cmt, epk_1, C_1, epk_2, C_2, \mathfrak{t}_s, \mathfrak{t}_1, \mathfrak{t}_2, m]$ is not defined, then set it to h_{ctr} , else abort.

Theorem 5.5. *The above scheme is correct if the proof systems (P_s, V_s) and (P_e, V_e) are complete.*

Theorem 5.6. *The above scheme is simulatable in the random-oracle model if the commitment scheme is hiding and the proof systems (P_s, V_s) and (P_e, V_e) are special HVZK and have high-entropy commitments.*

Theorem 5.7. *The above scheme is extractable in the random-oracle model if the commitment scheme is binding and the proof systems (P_s, V_s) and (P_e, V_e) are special-sound and have super-polynomial challenge spaces.*

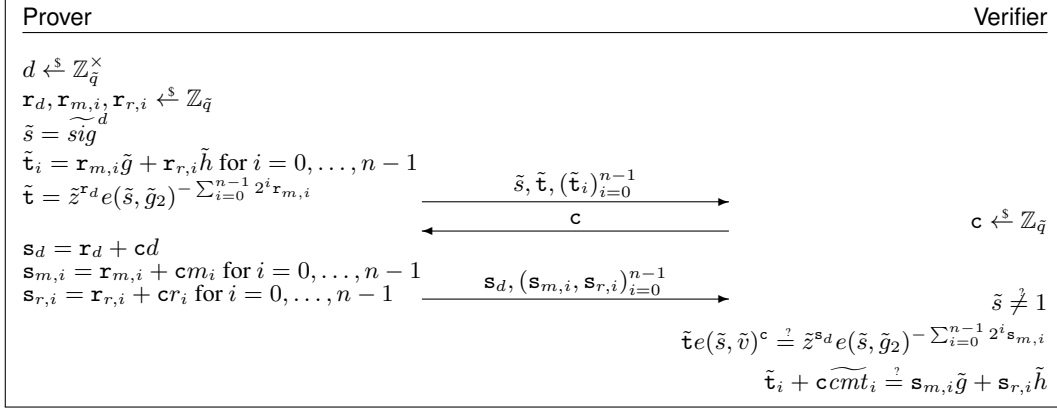
Theorem 5.8. *The above scheme is simulation-sound if the underlying commitment scheme is binding and the underlying Σ -protocols (P_s, V_s, S_s) and (P_e, V_e, S_e) are special-sound, have quasi-unique responses, and have super-polynomial challenge spaces.*

Due to length limitations, the proofs of the previous theorems can be found in Appendices C.5 to C.8.

5.4 Σ -Protocols for Boneh-Boyen Signatures and the Group Signature Scheme

In the following we briefly recap the weakly unforgeable version of the Boneh-Boyen signature scheme [6,7]. We assume that the reader is familiar with bilinear pairings and the strong Diffie-Hellman (SDH) assumption, and only recap them in Appendix A.4.

The Boneh-Boyen signature scheme is defined as follows for a bilinear group generator BGen:



Protocol 5.10: Proof of possession of a signature on m , which is also contained in a set of Pedersen commitments.

SKG. This algorithm first computes $(\tilde{q}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow^s \text{BGGen}(1^\lambda)$. It chooses $\tilde{g}_1 \leftarrow^s \mathbb{G}_1$, $\tilde{g}_2 \leftarrow^s \mathbb{G}_2$, $x \leftarrow^s \mathbb{Z}_{\tilde{q}}^\times$, and defines $\tilde{v} = \tilde{g}_2^x$ and $\tilde{z} = e(\tilde{g}_1, \tilde{g}_2)$. It outputs $spk = ((\tilde{q}, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e), \tilde{g}_1, \tilde{g}_2, \tilde{v}, \tilde{z})$ and $ssk = x$.

SSign. To sign a message $m \in \mathbb{Z}_{\tilde{q}}$ with secret key $ssk = x$, this algorithm outputs the signature $\widetilde{sig} = \tilde{g}_1^{\frac{1}{x+m}}$.

SVerify. Given a signature public key spk , a message $m \in \mathbb{Z}_{\tilde{q}}$ and a signature \widetilde{sig} , this algorithm outputs `accept` if $\tilde{v}\tilde{g}_2^m = 1$ in case $\widetilde{sig} = 1$, and if $e(\widetilde{sig}, \tilde{v}\tilde{g}_2^m) = \tilde{z}$ in case $\widetilde{sig} \neq 1$. In all other cases, it outputs `reject`.

Lemma 5.9. *If the SDH assumption holds for BGGen, then the above scheme is a weakly unforgeable signature scheme according to Definition A.3.*

We next show how a user can prove possession of a Boneh-Boyen signature on a message m , while keeping both, the message and the signature, private. In addition, the proof will additionally show that the m is also contained in a set of Pedersen commitments $\widetilde{cmt}_i = m_i\tilde{g} + r_i\tilde{h}$ such that $m = \sum_{i=0}^{n-1} 2^i m_i$, cf. Section 2.2.

The idea underlying Protocol 5.10 is similar to that in Camenisch et al. [9]: The prover first rerandomizes the signature to obtain a value s , which it sends to the verifier. Subsequently, the prover and the verifier run a standard Schnorr proof for the resulting statement.

Theorem 5.11. *Protocol 5.10 is a perfectly HVZK proof of knowledge for the following relations:*

$$\mathcal{R} = \left\{ \left((\widetilde{sig}, m, r, (m_i, r_i)_{i=0}^{n-1}), (spk, (\widetilde{cmt}_i)_{i=0}^{n-1}) \right) : \text{SVerify}(spk, m, \widetilde{sig}) = \text{accept} \wedge \right. \\ \left. \widetilde{cmt}_i = m_i\tilde{g} + r_i\tilde{h} \wedge \tilde{v}\tilde{g}_2^m \neq 1 \wedge m = \sum_{i=0}^{n-1} 2^i m_i \right\}$$

and \mathcal{R}' is defined as \mathcal{R} but without the condition that $\tilde{v}\tilde{g}_2^m \neq 1$.

The protocol is perfectly complete, and has a knowledge error of $1/\tilde{q}$. Furthermore, the protocol has quasi unique responses and high-entropy commitments.

The proof of this theorem is straightforward and can be found in Appendix C.9.

The Group Signature Scheme. Combining Protocols 4.1 and 5.10 now directly gives a group signature by the construction from Section 5.3. The \mathcal{ID} is given by $\{0, 1\}^\ell$ (which can be identified with $\{0, \dots, 2^\ell - 1\}$), where $\ell = \min(n, \lfloor \log_2 \tilde{q} \rfloor - 1)$, n is the dimension of the ring being used and \tilde{q} is the order of the groups of the commitment- and the signature schemes. All homomorphisms there can just be instantiated with a multiplication by 2. Note here that for the given identity space and the way our constructions work, multiplying by 2 can be seen as multiplication over the integers, as no modular reductions take place for all elements from \mathcal{ID} . Note further that NTRU can be made perfectly complete by truncating the discrete Gaussian distributions for key generation and encryption appropriately without touching security.

References

1. Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, April 2012.
2. Giuseppe Ateniese, Jan Camenisch, Marc Joye, and Gene Tsudik. A practical and provably secure coalition-resistant group signature scheme. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 255–270. Springer, August 2000.
3. Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *CRYPTO*, pages 390–420, 1992.
4. Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 614–629. Springer, May 2003.
5. Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 390–399. ACM Press, October / November 2006.
6. Dan Boneh and Xavier Boyen. Short signatures without random oracles. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 56–73. Springer, May 2004.
7. Dan Boneh and Xavier Boyen. Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology*, 21(2):149–177, April 2008.
8. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 41–55. Springer, August 2004.
9. Jan Camenisch, Maria Dubovitskaya, and Gregory Neven. Oblivious transfer with access control. In *ACM Conference on Computer and Communications Security*, pages 131–140, 2009.
10. Jan Camenisch, Gregory Neven, and Markus Rückert. Fully anonymous attribute tokens from lattices. In Ivan Visconti and Roberto De Prisco, editors, *SCN 12*, volume 7485 of *LNCS*, pages 57–75. Springer, September 2012.
11. Melissa Chase and Anna Lysyanskaya. On signatures of knowledge. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 78–96. Springer, August 2006.
12. David Chaum and Eugène van Heyst. Group signatures. In Donald W. Davies, editor, *EUROCRYPT’91*, volume 547 of *LNCS*, pages 257–265. Springer, April 1991.
13. Ronald Cramer. *Modular Design of Secure yet Practical Cryptographic Protocols*. PhD thesis, CWI and University of Amsterdam, 1997.
14. Ivan Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In *EUROCRYPT*, pages 418–430, 2000.
15. Ivan Damgård. On Σ -Protocols. Lecture on Cryptologic Protocol Theory; Faculty of Science, University of Aarhus, 2010.
16. Ivan Damgård and Eiichiro Fujisaki. A statistically-hiding integer commitment scheme based on groups with hidden order. In *ASIACRYPT*, pages 125–142, 2002.
17. Ivan Damgård, Oded Goldreich, Tatsuaki Okamoto, and Avi Wigderson. Honest verifier vs dishonest verifier in public coin zero-knowledge proofs. In *CRYPTO*, pages 325–338, 1995.
18. Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662, 2012.
19. Eiichiro Fujisaki and Tatsuaki Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In *CRYPTO*, pages 16–30, 1997.
20. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
21. S. Dov Gordon, Jonathan Katz, and Vinod Vaikuntanathan. A group signature scheme from lattice assumptions. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 395–412. Springer, December 2010.
22. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288, 1998.
23. Fabien Laguillaumie, Adeline Langlois, Benoît Libert, and Damien Stehlé. Lattice-based group signatures with logarithmic signature size. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 41–61. Springer, December 2013.
24. San Ling, Khoa Nguyen, Damien Stehlé, and Huaxiong Wang. Improved zero-knowledge proofs of knowledge for the isis problem, and applications. In *Public Key Cryptography*, pages 107–124, 2013.
25. Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Public Key Cryptography*, pages 162–179, 2008.
26. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
27. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
28. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013. Preliminary version appeared in *EUROCRYPT 2010*.
29. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd ACM STOC*, pages 427–437. ACM Press, May 1990.
30. Torben P. Pedersen. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In *CRYPTO*, pages 129–140, 1991.

31. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.
32. Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th FOCS*, pages 543–553. IEEE Computer Society Press, October 1999.
33. Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *CRYPTO*, pages 239–252, 1989.
34. Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, pages 27–47, 2011.
35. Jacques Stern. A new identification scheme based on syndrome decoding. In *CRYPTO*, pages 13–21, 1993.

A Definitions

A.1 Commitment Schemes

In the following we formally define commitment schemes.

Definition A.1. A commitment scheme is a triple $(\text{CSetup}, \text{Commit}, \text{COpen})$ such that:

- On input 1^λ , the key generation algorithm CSetup outputs commitment parameters $cpars$.
- The commitment algorithm Commit takes as inputs a message m from a message space \mathcal{M} and commitment parameters $cpars$, and outputs a commitment/opening pair (cmt, o) .
- The opening algorithm COpen takes parameters $cpars$, a message m , a commitment cmt , and opening o , and outputs `accept` or `reject`.

Furthermore, the algorithms satisfy the following properties:

- Correctness: COpen outputs `accept` whenever the inputs were computed by an honest party, i.e.,

$$\Pr \left[\text{COpen}(cpars, m, cmt, o) = \text{accept} : \begin{array}{l} cpars \xleftarrow{\$} \text{CSetup}(1^\lambda), m \in \mathcal{M}, \\ (cmt, o) \xleftarrow{\$} \text{Commit}(cpars, m) \end{array} \right] = 1$$

- Binding: A commitment cannot be opened to different messages. A scheme is said to be computationally binding if no PPT adversary can come up with a commitment and two different openings, i.e., for every PPT adversary A there exists a negligible function negl such that:

$$\Pr \left[\begin{array}{l} b = b' = \text{accept} \wedge m \neq m' : \\ cpars \xleftarrow{\$} \text{CSetup}(1^\lambda), (cmt, m, o, m', o') \xleftarrow{\$} A(cpars), \\ b = \text{COpen}(cpars, m, cmt, o), \\ b' = \text{COpen}(cpars, m', cmt, o') \end{array} \right] \leq \text{negl}(\lambda)$$

A scheme is said to be perfectly binding if this holds unconditionally, i.e., $\text{negl}(\lambda) = 0$.

- Hiding: A commitment computationally hides the committed message: for every probabilistic polynomial time (PPT) adversary A there is a negligible function negl such that:

$$\Pr \left[b = b' : \begin{array}{l} cpars \xleftarrow{\$} \text{CSetup}(1^\lambda), (m_0, m_1, \text{aux}) \xleftarrow{\$} A_1(cpars), b \xleftarrow{\$} \{0, 1\}, \\ (cmt, o) = \text{Commit}(m_b, cpars), b' \xleftarrow{\$} A_2(cmt, \text{aux}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

A scheme is said to be perfectly hiding if a commitment information-theoretically does not contain any information about the contained message, i.e., $\text{negl}(\lambda) = 0$.

A.2 Semantically Secure Encryption

Definition A.2. A semantically secure public key encryption scheme is a triple $(\text{EncKG}, \text{Enc}, \text{Dec})$ satisfying the following:

- On input 1^λ , EncKG outputs a secret key/public key pair (esk, epk) .
- The encryption algorithm Enc takes the public key epk and a message m from some message space \mathcal{M} , and outputs a ciphertext c .

- The decryption algorithm Dec takes the secret key esk and a ciphertext c and returns a message $m' \in \mathcal{M} \cup \{\perp\}$.

Furthermore, the algorithms have to satisfy the following properties:

- Correctness: Dec returns m if c is an honest encryption of m with overwhelming probability. That is, there exists a negligible function negl such that:

$$\Pr \left[m = m' : (esk, epk) \xleftarrow{\$} \text{EncKG}(1^\lambda), m \in \mathcal{M}, c \xleftarrow{\$} \text{Enc}(epk, m), m' = \text{Dec}(esk, c) \right] \geq 1 - \text{negl}(\lambda)$$

For every $m \in \mathcal{M}$ we have that $\text{Dec}(esk, \text{Enc}(epk, m)) = m$ whenever $(esk, epk) \xleftarrow{\$} \text{EncKG}(1^\lambda)$

- Semantic security: A scheme is called semantically (or IND-CPA) secure, if for every PPT adversary $A = (A_1, A_2)$ there exists a negligible function negl such that:

$$\Pr \left[b = b' : (esk, epk) \xleftarrow{\$} \text{EncKG}(1^\lambda), b \xleftarrow{\$} \{0, 1\}, (m_0, m_1, \text{aux}) \xleftarrow{\$} A_1(epk) \right. \\ \left. c = \text{Enc}(epk, m_b), b' \xleftarrow{\$} A_2(\text{aux}, c) \right] \leq \text{negl}(\lambda).$$

A.3 Weakly Unforgeable Signatures

Definition A.3. A standard signature scheme is a tuple $(\text{SKG}, \text{SSign}, \text{SVerify})$ where:

- On input 1^λ , the key generation algorithm SKG outputs a public key spk and a signing key ssk .
- On input ssk and message $m \in \mathcal{M}$, the signing algorithm SSign outputs a signature sig .
- On input spk , m , sig , the verification algorithm SVerify outputs accept or reject .

The algorithms satisfy the following properties:

- Correctness: SVerify outputs accept whenever the keys and signature are honestly generated, i.e., for all $\lambda \in \mathbb{N}$ and all $m \in \mathcal{M}$

$$\Pr[\text{SVerify}(\text{spk}, m, \text{sig}) = \text{accept} : (\text{spk}, \text{ssk}) \xleftarrow{\$} \text{SKG}(1^\lambda), \text{sig} \xleftarrow{\$} \text{SSign}(\text{ssk}, m)] = 1.$$

- Weak unforgeability: The signature scheme is weakly unforgeable if for every PPT adversary A there is a negligible function negl such that

$$\Pr \left[\text{SVerify}(\text{spk}, m, \text{sig}) = \text{accept} \wedge m \notin \{m_1, \dots, m_s\} : \right. \\ \left. \begin{array}{l} (m_1, \dots, m_s) \xleftarrow{\$} A(1^\lambda), (\text{spk}, \text{ssk}) \xleftarrow{\$} \text{SKG}(1^\lambda), \\ \text{For } i = 1, \dots, s \text{ do } \text{sig}_i \xleftarrow{\$} \text{SSign}(\text{ssk}, m_i), \\ (m, \text{sig}) \xleftarrow{\$} A(\text{spk}, m_1, \dots, m_s) \end{array} \right] \leq \text{negl}(\lambda).$$

A.4 Bilinear Pairings and the SDH Assumption

Let therefore q be prime, and $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ be groups of order q . A mapping $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is called a bilinear pairing, if it satisfies the following:

1. Bilinearity: $\forall u \in \mathbb{G}_1, \forall v \in \mathbb{G}_2, \forall a, b \in \mathbb{Z}$ it holds that $e(u^a, v^b) = e(u, v)^{ab}$,
2. Non-degeneracy: if g_1 and g_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively, then $e(g_1, g_2)$ generates \mathbb{G}_T , and
3. Computability: all group operations and e can be efficiently computed.

The scheme presented by Boneh and Boyen is provable secure under the strong Diffie-Hellman assumption recapitulated next.

Definition A.4. Let BGGen , the bilinear group generator, be a PPT algorithm that, on input 1^λ , outputs a prime q , groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of order q , and a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. We say that the strong Diffie-Hellman (SDH) assumption holds for BGGen , if for every PPT adversary A and every polynomially bounded function k there exists a negligible function negl such that the following holds:

$$\Pr \left[(a_1, a_2) = (c, g_1^{\frac{1}{x+c}}) \in \mathbb{Z} \times \mathbb{G}_1 : \begin{array}{l} (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \xleftarrow{\$} \text{BGGen}(1^\lambda), g_1 \xleftarrow{\$} \mathbb{G}_1, g_2 \xleftarrow{\$} \mathbb{G}_2, \\ x \xleftarrow{\$} \mathbb{Z}_q^\times, (a_1, a_2) \xleftarrow{\$} A(g_1, g_1^x, \dots, g_1^{(x, k(\lambda))}, g_2, g_2^x) \end{array} \right] \leq \text{negl}(\lambda).$$

B Generalized Forking Lemma

We make use of the forking lemma of Bellare and Neven [5], which is a generalized variant of the original forking lemma due to Pointcheval and Stern [31].

Lemma B.1. *Let q be an integer and let \mathcal{C} be a set. Let A be a randomized algorithm that on input in, h_1, \dots, h_q returns a pair (I, out) , where $0 \leq I \leq q$. Let IG be a randomized input generator algorithm. The accepting probability of A is*

$$acc = \Pr [I \geq 1 : in \xleftarrow{\$} IG, h_1, \dots, h_q \xleftarrow{\$} \mathcal{C}, (I, out) \xleftarrow{\$} A(in, h_1, \dots, h_q)] .$$

The forking algorithm $\text{Fork}_A(in)$ associated to A first chooses random coins ρ for A and $h_1, \dots, h_q \xleftarrow{\$} \mathcal{C}$. It then runs $A(in, h_1, \dots, h_q; \rho)$ to an integer I and side output out . If $I = 0$ then it returns $(0, \varepsilon, \varepsilon)$, otherwise it chooses fresh random $h'_1, \dots, h'_q \xleftarrow{\$} \mathcal{C}$ and runs A again on input $(in, h_1, \dots, h_{I-1}, h'_1, \dots, h'_q)$ and on the same random tape ρ until it outputs (I', out') . If $I = I'$ and $h_I \neq h'_I$, then Fork_A outputs $(1, out, out')$, otherwise it returns $(0, \varepsilon, \varepsilon)$. Let

$$frk = \Pr [b = 1 : in \xleftarrow{\$} IG, (b, out, out') \xleftarrow{\$} \text{Fork}_A(in)] .$$

Then

$$frk \geq acc \cdot \left(\frac{acc}{q} - \frac{1}{\#\mathcal{C}} \right) .$$

Alternatively,

$$acc \leq \frac{q}{\#\mathcal{C}} + \sqrt{q \cdot frk} .$$

C Proofs

C.1 Proof of Lemma 2.3

Proof. Because $q = 3 \pmod 8$, the polynomial $X^n + 1$ splits into two irreducible polynomials modulo q . Now suppose that there are two possible plaintexts $hs + pe + m$ and $hs' + pe + m'$ that correspond to the same NTRU ciphertext c . Then $h(s - s') + p(e - e') + (m - m') = 0$. Since $X^n + 1$ splits into two irreducible polynomials of degree $n/2$ modulo q , we have that for a fixed $s - s', e - e', m - m'$, $\Pr_h[h(s - s') + p(e - e') + (m - m') = 0] \leq \frac{1}{q^{n/2}}$. The claim in the lemma is obtained by applying the union bound. \square

C.2 Proof of Theorem 4.2

Proof. We need to prove the properties of Definition 2.5.

Completeness. First note that by Theorem 2.4, the prover will respond with probability $1/M$ for each secret, i.e., with an overall probability of $1/M^3$. Assuming that the prover does not abort, we have that:

$$\begin{aligned} h\mathbf{s}_s + p\mathbf{s}_e + \mathbf{s}_m &= h(\mathbf{r}_s + X^c s) + p(\mathbf{r}_e + X^c e) + (\mathbf{r}_m + X^c m) \\ &= X^c(hs + pe + m) + (h\mathbf{r}_s + p\mathbf{r}_e + \mathbf{r}_m) = X^c y + \mathbf{t} . \end{aligned}$$

For the norms it is easy to see that $\|\mathbf{s}_s\| \leq \|\mathbf{r}_s\| + \|s\| \leq \|\mathbf{r}_s\| + \sqrt{n}\alpha \leq n\alpha$ with overwhelming probability because the standard deviation of \mathbf{r}_s is much smaller than $n\alpha$, and similar for \mathbf{s}_e and \mathbf{s}_m .

What remains to show are the verification equations for the Pedersen commitments. First, note that as q, \tilde{q} are much larger than the standard deviation of $\mathbf{r}_s, \mathbf{r}_m$, the computation of \mathbf{s}_s and \mathbf{s}_m is over the integers, i.e., with overwhelming probability there is no modular reduction. Furthermore, multiplying by X^c corresponds to anti-cyclically shifting the coefficients in the respective polynomial.

$$\begin{aligned} \mathbf{s}_m \tilde{g} + \mathbf{s}_r \tilde{h} &= ((\mathbf{r}_{m,0}, \dots, \mathbf{r}_{m,n-1}) + (m_0, \dots, m_{n-1})_{\ll c}) \tilde{g} + ((\mathbf{r}_{r,0}, \dots, \mathbf{r}_{r,n-1}) + (r_0, \dots, r_{n-1})_{\ll c}) \tilde{h} \\ &= (\mathbf{r}_{m,0}, \dots, \mu_{m,n-1}) \tilde{g} + (\mathbf{r}_{r,0}, \dots, \mathbf{r}_{r,n-1}) \tilde{h} + (m_0, \dots, m_{n-1})_{\ll c} \tilde{g} + (r_0, \dots, r_{n-1})_{\ll c} \tilde{h} \\ &= (\tilde{\mathbf{t}}_0, \dots, \tilde{\mathbf{t}}_{n-1}) + \widetilde{cmt}_{\ll c} \end{aligned}$$

Finally, by the completeness of the auxiliary commitment scheme, aCOpen will always output accept for an honest prover.

Honest-verifier zero-knowledge. The proof is a straightforward generalization of that for Theorem 3.3 and the basic Schnorr protocol [33], and therefore omitted.

Soundness. Assume that we have two accepting transcripts for the same first message but different challenges, i.e., that we have $(c_{\text{aux}}, c', ((\mathbf{t}', (\tilde{\mathbf{t}}'_i)_{i=0}^{n-1}), d'_{\text{aux}}, (\mathbf{s}'_s, \mathbf{s}'_e, \mathbf{s}'_m, \mathbf{s}'_r)))$ and $(c_{\text{aux}}, c'', ((\mathbf{t}'', (\tilde{\mathbf{t}}''_i)_{i=0}^{n-1}), d''_{\text{aux}}, (\mathbf{s}''_s, \mathbf{s}''_e, \mathbf{s}''_m, \mathbf{s}''_r)))$ that pass the checks performed by the verifier.

Then, by the binding property of the auxiliary commitment scheme, we have that $\mathbf{t}' = \mathbf{t}'' =: \mathbf{t}$ and $(\tilde{\mathbf{t}}'_i)_{i=0}^{n-1} = (\tilde{\mathbf{t}}''_i)_{i=0}^{n-1} =: (\tilde{\mathbf{t}}_i)_{i=0}^{n-1}$.

Now, by subtracting the verification equations, we get that:

$$\begin{aligned} (\widetilde{cmt}_0, \dots, \widetilde{cmt}_{n-1})_{\ll c''} - (\widetilde{cmt}_0, \dots, \widetilde{cmt}_{n-1})_{\ll c'} &= (\mathbf{s}''_m - \mathbf{s}'_m)\tilde{g} + (\mathbf{s}''_r - \mathbf{s}'_r)\tilde{h}, \\ (X^{c''} - X^{c'})y &= h(\mathbf{s}''_s - \mathbf{s}'_s) + p(\mathbf{s}''_e - \mathbf{s}'_e) + (\mathbf{s}''_m - \mathbf{s}'_m). \end{aligned}$$

By the symmetry of anti-cyclic shifts and multiplication by powers of X , this is easily seen to be equivalent to:

$$(\widetilde{cmt}_0, \dots, \widetilde{cmt}_{n-1})C = (\mathbf{s}''_m - \mathbf{s}'_m)\tilde{g} + (\mathbf{s}''_r - \mathbf{s}'_r)\tilde{h}, \quad (1)$$

$$yC = h(\mathbf{s}''_s - \mathbf{s}'_s) + p(\mathbf{s}''_e - \mathbf{s}'_e) + (\mathbf{s}''_m - \mathbf{s}'_m), \quad (2)$$

where the i^{th} column of the matrix C contains the coefficients of $(X^{c'} - X^{c'') \bmod X^n + 1}$. Note that as $\|\mathbf{s}'_m\|, \|\mathbf{s}''_m\|$ are small, the coefficients of $\mathbf{s}'_m - \mathbf{s}''_m$ in R_q and $\mathbb{Z}_{\tilde{q}}$ are the same over the integers, as no modular reduction takes place here.

By Lemma 3.1 we have that $2(X^{c'} - X^{c'')$ is invertible in $\mathbb{Z}/(X^n + 1)$, and the inverse only has coefficients in $\{1, 0, 1\}$. Let d be the inverse, and D denote the corresponding matrix. Since 2 is coprime with q and \tilde{q} , D can also be seen as a matrix with entries in R_q and $\mathbb{Z}_{\tilde{q}}$. By multiplying (1) and (2) by D from the right, we get that $\hat{m} = 2(\mathbf{s}'_m - \mathbf{s}''_m)d$, $\hat{r} = 2(\mathbf{s}'_r - \mathbf{s}''_r)D$, $\hat{e} = 2(\mathbf{s}'_e - \mathbf{s}''_e)d$ and $\hat{s} = 2(\mathbf{s}'_s - \mathbf{s}''_s)d$ satisfy the following:

$$\begin{aligned} h\hat{s} + p\hat{e} + \hat{m} &= 2y \\ \hat{m}_i g + \hat{r}_i h &= 2\mathbf{t}_i \quad \text{for } 0 \leq i < n, \end{aligned}$$

and the norms of $\hat{e}, \hat{m}, \hat{s}$ satisfy $\|\hat{m}\|_{\infty} \leq n(\|\mathbf{s}'_m\|_{\infty} + \|\mathbf{s}''_m\|_{\infty}) \leq 2n^2$, and similarly for \hat{e} and \hat{s} . Note again that as $q, \tilde{q} > 2n^2$, \hat{m} is the same in R_q and $\mathbb{Z}_{\tilde{q}}$ over the integers.

Quasi-unique responses. Assume that a PPT adversary outputs $(c_{\text{aux}}, c, (\mathbf{t}', \tilde{\mathbf{t}}'_0, \dots, \tilde{\mathbf{t}}'_{n-1}), d'_{\text{aux}}, (\mathbf{s}'_s, \mathbf{s}'_e, \mathbf{s}'_m, \mathbf{s}'_r))$ and $(c_{\text{aux}}, c, (\mathbf{t}'', \tilde{\mathbf{t}}''_0, \dots, \tilde{\mathbf{t}}''_{n-1}), d''_{\text{aux}}, (\mathbf{s}''_s, \mathbf{s}''_e, \mathbf{s}''_m, \mathbf{s}''_r))$ that pass the tests performed by the verifier. By the strong binding requirements on the auxiliary commitment scheme, we get that $\mathbf{t}' = \mathbf{t}'' =: \mathbf{t}$, $\tilde{\mathbf{t}}'_i = \tilde{\mathbf{t}}''_i =: \tilde{\mathbf{t}}_i$ for all i , and $d'_{\text{aux}} = d''_{\text{aux}}$. What remains to show is that $\mathbf{s}_s, \mathbf{s}_e, \mathbf{s}_m, \mathbf{s}_r$ are unique as well.

We have that $0 = (\mathbf{s}'_m - \mathbf{s}''_m)\tilde{g} + (\mathbf{s}'_r - \mathbf{s}''_r)\tilde{h}$. Note that for every i , $\mathbf{s}'_{m,i} \neq \mathbf{s}''_{m,i}$ if and only if $\mathbf{s}'_{r,i} \neq \mathbf{s}''_{r,i}$. In this case we get that $\log_{\tilde{g}} \tilde{h} = -\frac{\mathbf{s}'_m - \mathbf{s}''_m}{\mathbf{s}'_r - \mathbf{s}''_r}$, which any PPT adversary can only compute with negligible probability under the assumed DLOG assumption.

We further get that $\mathbf{s}'_s = \mathbf{s}''_s$, $\mathbf{s}'_m = \mathbf{s}''_m$ and $\mathbf{s}'_e = \mathbf{s}''_e$ in R_q , as otherwise we had that $h(\mathbf{s}'_s - \mathbf{s}''_s) + p(\mathbf{s}'_e - \mathbf{s}''_e) + (\mathbf{s}'_m - \mathbf{s}''_m) = 0$. However, if an adversary could come up with such values, it would be able to distinguish uniformly random h from NTRU public keys, as by Lemma 2.3 there does not exist a non-trivial solution to the last equation with overwhelming probability for truly random h .

High-entropy commitments. This directly follows from the security of the auxiliary commitment scheme. \square

C.3 Proof of Theorem 5.3

Proof. We prove the theorem by describing a sequence of indistinguishable games, where the adversary's view in the final game is completely independent of the hidden bit b chosen by the experiment.

Game 0: This is the original game as described in Figure 1.

- Game 1:** We now use the simulator SoKSim to generate the signature of knowledge sok^* in the challenge signature $\sigma^* = (C_1^*, C_2^*, sok^*)$. Meaning, we use StpSim to generate the parameters $sokp$, we use ROSim to respond to A's random-oracle queries, and we compute sok^* as $SSim((spk, epk_1, epk_2, C_1, C_2), (i, \text{gsk}[i_b^*], \rho_1, \rho_2), m^*)$. It is easy to see that any adversary A distinguishing Game 0 from Game 1 can be used to break the simulatability of the signature of knowledge scheme.
- Game 2:** Instead of computing sok^* using $SSim$, we generate it as $SSim'((spk, epk_1, epk_2, C_1, C_2), m^*)$, i.e., without first checking language membership. Since we only call $SSim'$ on actual language elements anyway, this change is purely conceptual, so this game is identical to the previous game from the point of view of the adversary.
- Game 3:** In the computation of the target signature $\sigma^* = (C_1^*, C_2^*, sok^*)$, we compute C_2^* as $\text{Enc}(epk_2, 1)$ instead of $\text{Enc}(epk_2, i_b^*)$. Any adversary distinguishing this game from the previous one can be used to break the semantic security of the encryption scheme. Note that we continue to simulate the opening oracle using esk_1 , as in the real scheme, so we never need to decrypt any ciphertexts under epk_2 . Also note that now sok^* is actually a signature of knowledge on a false statement, because C_1^* and C_2^* encrypt different user identities.
- Game 4:** We now respond to A's opening queries by decrypting C_2 using esk_2 instead of decrypting C_1 using esk_1 . This does not change the view of the adversary as long as all pairs (C_1, C_2) included in A's opening queries decrypt to the same plaintext. The simulation-soundness guarantees that, except with negligible probability, all group signatures include actual language elements, meaning that there exist i, ρ_1, ρ_2 such that $\Phi(C_1) = \text{Enc}(epk_1, \varphi(i); \rho_2)$ and $\Phi(C_2) = \text{Enc}(epk_2, \varphi(i); \rho_2)$. Due to the perfect completeness of the encryption scheme, this means that both ciphertexts decrypt to the same plaintext d . Due to the injectivity of φ , they therefore lead to the same i such that $\varphi(i) = d$.
- Game 5:** We now also replace C_1^* with an encryption of 1 instead of b . Since we no longer decrypt anything under esk_1 during the simulation, this game is indistinguishable from the previous one by the semantic security of the encryption scheme. Note that at this point, the simulation is completely independent of the bit b , so that A has exactly the same probability to output 1 for either choice of b . \square

C.4 Proof of Theorem 5.4

Proof. We distinguish between the case that A wins because $\text{GOpen}(gok, m, \sigma) = \perp$, and the case that it does so because $\text{GVerify}(gpk, m, \sigma) = 1 \wedge i \notin Q_{\text{gsk}} \wedge (i, m) \notin Q_{\text{GSign}} \wedge i \neq \perp$.

In the former case, simulatability and extractability together mean that, except with negligible probability, there must exist i, sig, ρ_1 such that $\text{SVerify}(spk, i, sig) = \text{accept}$ and $\Phi(C_1) = \text{Enc}(epk_1, \varphi(i); \rho_1)$. By the perfect completeness of the encryption scheme, such a ciphertext cannot fail to decrypt under esk_1 to some plaintext d . The only way for opening to fail is if there does not exist $i' \in \mathcal{ID}$ such that $d = \varphi(i')$, which means that $i \notin \mathcal{ID}$, meaning that sig must be a forgery.

In the latter case, consider the following sequence of games.

- Game 0:** This is the real traceability game.
- Game 1:** We now simulate random-oracle responses and the signatures of knowledge included in issued group signatures with SoKSim. This game is indistinguishable from Game 0 by the simulatability of the signature of knowledge scheme.
- Game 2:** We now use the extractor SoKExt on the signature of knowledge in the forgery to obtain (sig, i) such that $\text{SVerify}(spk, i, sig) = \text{accept}$. By the extractability property, the extractor succeeds in doing so with non-negligible probability.
- Game 3:** We now simulate all signatures of knowledge in GSign responses using $SSim'$, i.e., without checking the witness. This game is identical to the previous one since we call $SSim$ only on valid language elements. Note that the simulation no longer requires to compute standard signatures on the user identities i .

Given an adversary A with non-negligible probability of winning in Game 3, consider the following adversary B against the weak unforgeability of the standard signature scheme. Algorithm B guesses an identity $i^* \xleftarrow{\$} \{1, \dots, n\}$ and asks for signatures on all identities $1, \dots, n$ except i^* . It runs A, using the challenge public key spk as part of the group public key gpk , and simulating random-oracle responses and signatures of knowledge

in group signatures with SoKSim. When A outputs its forgery, B extracts (sig, i) . If $i = i^*$, B outputs (sig, i) as its own forgery. \square

C.5 Proof of Theorem 5.5

Proof. The SoKVerify algorithm only rejects if V_s or one of the two executions of V_e reject. Given that they each reject at most with negligible probability α , SoKVerify rejects with probability at most 3α . \square

C.6 Proof of Theorem 5.6

Proof. We prove the theorem through a sequence of indistinguishable games, where the first game is identical to $A^{H(\cdot), \text{SoKSign}(\cdot, \cdot)}(sokp)$ with real parameters $sokp$ and the last is identical to $A^{\text{ROSim}(\cdot), \text{SSim}(\cdot, \cdot)}(sokp)$ with simulated parameters.

Game 0: This is the original game $A^{H(\cdot), \text{SoKSign}(\cdot, \cdot)}(sokp)$ with $sokp \stackrel{\$}{\leftarrow} \text{SoKSetup}(1^\lambda)$.

Game 1: Generate $sokp \stackrel{\$}{\leftarrow} \text{StpSim}(1^\lambda)$ and respond to A's random-oracle queries, as well as the random-oracle queries made by the SoKSign oracle, using $\text{ROSim}(\cdot)$. This game is identical to the previous game since both generate parameters through $\text{CSetup}(1^\lambda)$ and both return truly random responses to random-oracle queries.

Game 2: If the random-oracle query caused by an invocation of the SoKSign oracle was made previously, either by the adversary, or by another SoKSign query, then the simulation is aborted. This happens with negligible probability due to the high-entropy commitments of the proof systems (P_s, V_s) and (P_e, V_e) .

Game 3: Respond to all signature queries $\text{SoKSign}(x, w, m)$ with $\text{SSim}(x, w, m)$. We prove that Game 2 is indistinguishable from Game 1 through a hybrid argument on the signature queries of A. Let Game 2. j be the game where A's first j signature queries are responded to with SSim , while all later queries are responded to with SoKSign . Clearly, we have that Game 2.0 is identical to Game 2 and Game 2. q_S to Game 3. Suppose that A can distinguish with non-negligible probability between Game 2. j and Game 2. $(j + 1)$. Then consider the following sequence of sub-games.

Game 2. j .0: Identical to Game 2. j

Game 2. j .1: When responding to the $j + 1$ st SSim query, rather than computing t_s, s_s through P_s , use the simulator S_s using the same challenge value c_s . This game is indistinguishable from Game 2. j .0 by the special honest-verifier zero-knowledge property of (P_s, V_s) .

Game 2. j .2: Similarly, use S_e instead of P_e to generate t_1, s_1, t_2, s_2 . Indistinguishability follows from the honest-verifier zero-knowledge property of (P_s, V_s) .

Game 2. j .3: Finally, generate the commitment cmt as a commitment of 1 instead of of i . Indistinguishability follows from the hiding property of the commitment scheme. \square

C.7 Proof of Theorem 5.7

Proof. Consider the algorithm B that runs the adversary A against a simulated random oracle and signing oracle. More precisely, on input $1^\lambda, h_1, \dots, h_q \stackrel{\$}{\leftarrow} \mathcal{C}$, and random tape (ρ_S, ρ_A) , algorithm B runs $sokp \stackrel{\$}{\leftarrow} \text{StpSim}(1^\lambda; \rho_S)$ and $(x^*, m^*, sok^*) \stackrel{\$}{\leftarrow} A^{\text{ROSim}(\cdot), \text{SSim}(\cdot, \cdot)}(sokp; \rho_A)$, using the values h_1, \dots, h_q as random-oracle responses. Let $Q^* = (spk, cpars, cmt, epk_1, C_1, epk_2, C_2, t_s, \tau_1, \tau_2, m^*)$ be the ‘‘crucial’’ random oracle query that is involved in the computation and verification of sok^* . We can assume without loss of generality that A makes this query before it stops executing, so let $I = HT[Q^*]$ be the index of the returned random-oracle response. If $\text{SoKVerify}(sokp, x^*, m^*, sok^*) = \text{accept}$ and A did not make a signature query (x^*, w, m^*) for any w such that $(x^*, w) \in R_{\mathcal{L}}$, then B outputs $(I, (x^*, m^*, sok^*))$, otherwise it outputs $(0, \varepsilon)$.

Using the generalized forking lemma, recalled in Lemma B.1, for any algorithm B that outputs (I, out^*) with probability acc , the forking algorithm Fork_B returns two accepting outputs out^*, out'^* corresponding to random-oracle responses $h_I \neq h'_I$ with probability

$$frk \geq acc \cdot \left(\frac{acc}{q} - \frac{1}{\#\mathcal{C}} \right).$$

Let $\text{SoKExt}_A(\text{sokp}, x, m, \text{sok}, \rho_S, \rho_A)$ be an algorithm that runs the forking algorithm as described above to obtain two valid signatures of knowledge $(x^*, m^*, \text{sok}^*), (x'^*, m'^*, \text{sok}'^*)$ for two different random-oracle responses $h_I \neq h'_I$. Since both runs of A are identical up to the I -th random-oracle response, the arguments Q^* and Q'^* to A 's I -th random-oracle query are identical as well. Because these arguments contain the full descriptions of the language elements x^*, x_s^*, x_1^*, x_2^* , we know that these are the same for both runs, as well as the message m^* and the first rounds of the Σ -protocols $\mathbf{t}_s^*, \mathbf{t}_1^*, \mathbf{t}_2^*$. Because $\text{SoKVerify}(\text{sokp}, x^*, m^*, \text{sok}^*) = \text{SoKVerify}(\text{sokp}, x^*, m^*, \text{sok}'^*) = \text{accept}$, we therefore have two valid transcripts $(\mathbf{t}_s^*, \mathbf{c}_s^*, \mathbf{s}_s^*), (\mathbf{t}_s^*, \mathbf{c}'_s^*, \mathbf{s}'_s^*)$ for $V_s((\text{spk}^*, \text{cpars}^*, \text{cmt}^*), \cdot)$ and two valid transcripts $(\mathbf{t}_j^*, \mathbf{c}_e^*, \mathbf{s}_j^*), (\mathbf{t}_j^*, \mathbf{c}'_e^*, \mathbf{s}'_j^*)$ for $V_e((\text{epk}_j^*, C_j^*, \text{cpars}^*, \text{cmt}^*), \cdot)$ for each of $j = 1, 2$.

If now $\mathbf{c}_s^* \neq \mathbf{c}'_s^*$ and $\mathbf{c}_e^* \neq \mathbf{c}'_e^*$, then we can use the special soundness of the proof systems and extract witnesses (sig, i, o) and (i', ρ_j, o') for $j = 1, 2$ using the knowledge extractors E_s, E_e . By the binding property of the commitment scheme, we have that $\varphi(i) = i'$, because otherwise $(\varphi(i), \chi(o))$ and (i', o') would be valid openings to different messages of the same commitment cmt^* . From the extracted witnesses, we can therefore compose a valid witness $(i, \text{sig}, \rho_1, \rho_2)$ for $x^* \in \mathcal{L}$.

However, the forking lemma only guarantees that $h_I = (\mathbf{c}_s^*, \mathbf{c}_e^*) \neq h'_I = (\mathbf{c}'_s^*, \mathbf{c}'_e^*)$, which does not necessarily imply that $\mathbf{c}_s^* \neq \mathbf{c}'_s^*$ and $\mathbf{c}_e^* \neq \mathbf{c}'_e^*$. Since all random-oracle responses are chosen at random, the probability that any of the components is equal is

$$\Pr[\mathbf{c}_s^* = \mathbf{c}'_s^* \vee \mathbf{c}_e^* = \mathbf{c}'_e^*] \leq \Pr[\mathbf{c}_s^* = \mathbf{c}'_s^*] + \Pr[\mathbf{c}_e^* = \mathbf{c}'_e^*] = \frac{1}{\#\mathcal{C}_s} + \frac{1}{\#\mathcal{C}_e}.$$

The probability that SoKExt_A succeeds is therefore at least

$$\begin{aligned} \text{frk} - \frac{1}{\#\mathcal{C}_s} - \frac{1}{\#\mathcal{C}_e} &\geq \text{acc} \cdot \left(\frac{\text{acc}}{q} - \frac{1}{\#\mathcal{C}} \right) - \frac{1}{\#\mathcal{C}_s} - \frac{1}{\#\mathcal{C}_e} \\ &\geq \frac{\text{acc}^2}{q} - \frac{3}{\min(\#\mathcal{C}_s, \#\mathcal{C}_e)}. \end{aligned}$$

If acc is non-negligible, then the first term is non-negligible. If the challenge spaces $\mathcal{C}_s, \mathcal{C}_e$ are super-polynomial, then the second term is negligible, so that overall SoKExt_A succeeds with non-negligible probability. \square

C.8 Proof of Theorem 5.8

Proof. Suppose there exists an adversary A that breaks the simulation-soundness with non-negligible probability. Without loss of generality, we assume that A makes both random-oracle queries Q, Q' involved in the verification $\text{SoKVerify}(\text{sokp}, x, m, \text{sok})$ and $\text{SoKVerify}(\text{sokp}, x', m', \text{sok}')$. Let $\text{sok} = (\mathbf{t}_s, \mathbf{t}_1, \mathbf{t}_2, \mathbf{s}_s, \mathbf{s}_1, \mathbf{s}_2)$ and $\text{sok}' = (\mathbf{t}'_s, \mathbf{t}'_1, \mathbf{t}'_2, \mathbf{s}'_s, \mathbf{s}'_1, \mathbf{s}'_2)$.

If $Q = Q'$, then we have that $(x, m, \mathbf{t}_s, \mathbf{t}_1, \mathbf{t}_2) = (x, m, \mathbf{t}_s, \mathbf{t}_1, \mathbf{t}_2)$, since all of these items are part of the random-oracle query. Since $\text{sok} = \text{sok}'$, it must hold that $(\mathbf{s}_s, \mathbf{s}_1, \mathbf{s}_2) \neq (\mathbf{s}'_s, \mathbf{s}'_1, \mathbf{s}'_2)$. We therefore have for at least one of the three Σ -protocols two accepting transcripts with the same commitment and challenge but different responses, which violates the quasi-unique responses of the proof systems.

If $Q \neq Q'$, then by a similar reasoning as in the proof of Theorem 5.7, the forking lemma (Lemma B.1) allows us to extract with non-negligible probability two valid signatures $\text{sok}', \text{sok}''$ for the same statement x' but for different random-oracle responses, if the commitment scheme is binding and the proof systems have super-polynomial challenge spaces. From these, a witness for x' can be extracted, which contradicts that $x' \notin \mathcal{L}$. \square

C.9 Proof of Theorem 5.11

Proof. We need to prove completeness, honest-verifier zero-knowledge and soundness. As the proof is standard, we only sketch it here.

Completeness. This is trivial to verify.

Honest-verifier zero-knowledge. Given a challenge c , the simulator S draws $\tilde{s}' \xleftarrow{\$} \mathbb{G}_1$ and $\mathbf{s}'_d, \mathbf{s}'_{m,i}, \mathbf{s}'_{r,i} \xleftarrow{\$} \mathbb{Z}_{\tilde{q}}$. It computes $\tilde{\mathbf{t}}' = e(\tilde{s}', \tilde{v})^{-c} \tilde{z}^{\mathbf{s}'_d} e(\tilde{s}', \tilde{g}_2)^{-\sum_{i=0}^{n-1} 2^i \mathbf{s}'_{m,i}}$ and $\tilde{\mathbf{t}}'_i = -c \widetilde{\text{cmt}}_i + \mathbf{s}'_{m,i} \tilde{g} + \mathbf{s}'_{r,i} \tilde{h}$. The simulator outputs

$((\tilde{s}', \tilde{\mathbf{t}}', (\tilde{\mathbf{t}}'_i)_{i=0}^{n-1}), \mathbf{c}, (\mathbf{s}'_d, (\mathbf{s}'_{m,i}, \mathbf{s}'_{r,i})_{i=0}^{n-1}))$. By standard arguments one can see that the output distribution is identical to that of real protocol runs for challenge \mathbf{c} .

Soundness. Using standard arguments one can show that, given two accepting protocol transcripts with the same first message but different challenges, one can efficiently extract a tuple $(d', (m'_i)_{i=0}^{n-1}, (r'_i)_{i=0}^{n-1})$ such that $e(\tilde{s}, \tilde{v}) = \tilde{z}^{d'} e(\tilde{s}, \tilde{g}_2)^{-\sum_{i=0}^{n-1} 2^i m'_i}$ and $\widetilde{cmt}_i = m'_i \tilde{g} + r'_i \tilde{h}$ for $i = 0, \dots, n-1$. Let now $m' = \sum_{i=0}^{n-1} 2^i m'_i$. If $d' \neq 0$, setting $\widetilde{sig}' = \tilde{s}^{1/d'}$ yields $e(\widetilde{sig}'^{d'}, \tilde{v}) = \tilde{z}^{d'} e(\widetilde{sig}'^{d'}, \tilde{g}_2)^{-m'}$. This is equivalent to $e(\widetilde{sig}', \tilde{v}) = \tilde{z} e(\widetilde{sig}', \tilde{g}_2)^{-m'}$ or $e(\widetilde{sig}', \tilde{v} \tilde{g}_2^{m'}) = \tilde{z}$, and thus \widetilde{sig}' is a valid signature for m' . If $d' = 0$, we get that $e(\tilde{s}, \tilde{v}) = e(\tilde{s}, \tilde{g}_2)^{-m'}$ or $e(\tilde{s}, \tilde{v} \tilde{g}_2^{m'}) = 1$, and therefore $\tilde{v} \tilde{g}_2^{m'} = 1$ because $\tilde{s} \neq 1$. In this case setting $\widetilde{sig}' = 1$ gives a valid signature for m' .

Quasi-unique responses. Assume that for some $\tilde{\mathbf{t}}, (\tilde{\mathbf{t}}_i)_{i=0}^{n-1}, \mathbf{c}$, a PPT adversary A outputs $(\mathbf{s}_d, (\mathbf{s}_{m,i}, \mathbf{s}_{r,i})_{i=0}^{n-1}), (\mathbf{s}'_d, (\mathbf{s}'_{m,i}, \mathbf{s}'_{r,i})_{i=0}^{n-1})$ that pass the verification checks. Define $\mathbf{s}_m = \sum_{i=0}^{n-1} \mathbf{s}_{m,i}$, and similarly for \mathbf{s}'_m . By dividing the verification equations we get that:

$$1 = \tilde{z}^{\mathbf{s}_d - \mathbf{s}'_d} e(\tilde{s}, \tilde{g}_2)^{\mathbf{s}'_m - \mathbf{s}_m} \quad \text{and} \quad 1 = (\mathbf{s}_{m,i} - \mathbf{s}'_{m,i}) \tilde{g} + (\mathbf{s}_{r,i} - \mathbf{s}'_{r,i}) \tilde{h}.$$

It follows that $\mathbf{s}_{m,i} = \mathbf{s}'_{m,i}$ and $\mathbf{s}_{r,i} = \mathbf{s}'_{r,i}$ for all i with overwhelming probability, as otherwise A could compute $\log_{\tilde{g}} \tilde{h} = -\frac{\mathbf{s}_{m,i} - \mathbf{s}'_{m,i}}{\mathbf{s}_{r,i} - \mathbf{s}'_{r,i}}$. Therefore, also $\mathbf{s}_m = \mathbf{s}'_m$, and thus $\mathbf{s}_d = \mathbf{s}'_d$ because \tilde{z} is a generator.

High-entropy commitments. This follows immediately from the fact that each element sent in the first message is independently and uniformly random in its domain. \square