Lattice-Based Group Signatures and Zero-Knowledge Proofs of Automorphism Stability*

Rafael del Pino[1**], Vadim Lyubashevsky[2], and Gregor Seiler[2,3]

[1] ENS Paris
[2] IBM Research – Zurich
[3] ETH Zurich

Abstract. We present a group signature scheme, based on the hardness of lattice problems, whose outputs are more than an order of magnitude smaller than the currently most efficient schemes in the literature. Since lattice-based schemes are also usually non-trivial to efficiently implement, we additionally provide the first experimental implementation of lattice-based group signatures demonstrating that our construction is indeed practical – all operations take less than half a second on a standard laptop.

A key component of our construction is a new zero-knowledge proof system for proving that a committed value belongs to a particular set of small size. The sets for which our proofs are applicable are exactly those that contain elements that remain stable under Galois automorphisms of the underlying cyclotomic number field of our lattice-based protocol. We believe that these proofs will find applications in other settings as well.

The motivation of the new zero-knowledge proof in our construction is to allow the efficient use of the selectively-secure signature scheme (i.e. a signature scheme in which the adversary declares the forgery message before seeing the public key) of Agrawal et al. (Eurocrypt 2010) in constructions of lattice-based group signatures and other privacy protocols. For selectively-secure schemes to be meaningfully converted to standard signature schemes, it is crucial that the size of the message space is not too large. Using our zero-knowledge proofs, we can strategically pick small sets for which we can provide efficient zero-knowledge proofs of membership.

1 Introduction

Commitments and zero-knowledge proofs of knowledge (ZKPoK) of committed values are a key ingredient in many privacy-based protocols. It is also often useful to prove various relations among the committed values, or that the committed values themselves have some particular characteristics. An example of the latter is proving that the commitment is to an element that belongs to a particular, possibly small, subset. Even if the subset stays fixed, this is not a trivial problem to solve efficiently for lattice-based commitments, and we are not aware of any previous practical solutions to this problem.

In this paper, we use the lattice-based commitment scheme [BDL+18] over cyclotomic rings (e.g. over $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^d+1)$) and consider sets that contain elements that remain stable under a

** Work done while at IBM Research – Zurich

certain subgroup of automorphisms of the cyclotomic number field (in our example $\mathbb{Q}[X]/(X^d+1)$). For the particular example of $\mathbb{Z}_q[X]/(X^d+1)$, this allows us to construct sets of size $q, q^2, q^4, \ldots$ for which we can build a ZKPoK showing that the commitment is to an element in this set.[1]

An application of our new proof system is towards constructing more practical lattice-based group signatures [CvH91,BMW03]. A group signature scheme consists of three parties – a trusted setup authority, a group manager (sometimes also called the opener), and group members. The setup authority generates a group public key and secret keys for all the group members. Using their secret keys, the group members can sign messages in a way that anyone can verify that a message was signed by a member of the group, but the identity of the signer remains secret (one should not even be able to tell that two messages were signed by the same member) to everyone except for the opener. The opener should be able to recover the identity of any signer.

Public key	Secret key	Signature	Key Generation	Signing	Verification
123 KB	146 KB	581 KB	429 ms	405 ms	169 ms

Table 1. User key, signature size, and running time of our C implementation on an Intel Skylake i7-6600U processor. The instantiation is of a CPA-anonymous version of our scheme with maximum group of 2^{80}. The CCA-anonymous version would have signatures approximately 20% longer.

Prior Work. A common way of constructing group signatures is via the sign-and-encrypt approach. The group public key that the setup authority creates is the public key to some signature scheme, and the secret key of a user with identity i is a signature of i. To sign a message, the group member produces a non-interactive ZKPoK that he has the authority's signature of some identity i.[2] Furthermore, the group member encrypts his identity i using the opener's public key, and gives another ZKPoK of the fact that the encryption is of the same identity as was used in the proof.

To create a practical scheme using the above approach, one typically needs to have a very efficient *standard model* signature scheme that is used by the setup authority to sign user identities.[3] While there exist efficient standard model signature schemes based on classical assumptions (e.g. [CL02]) which can be used for constructions of fairly compact group signatures, the non-existence of such signatures based on lattice assumptions, or any other post-quantum hard problem, is the main culprit in the fact that the only "efficiency" lattice-based group signatures have is asymptotic (c.f. [GKV10,LLNW16,LNWX18]).

Lattice-based signature schemes in the standard model are built based on Boyen's framework [Boy10]. There have been efficiency improvements to this scheme (e.g. [DM15,KY16]) that used polynomial lattices, but they still appear to be unsuitable for producing practical (group) signatures. The only group signature appearing prior to our work that proposes concrete parameters uses different techniques, and the signatures in it are on the order of 50MB [LLNW16].

[1] More precisely, because the ZKPoK of the commitment scheme in [BDL+18] is "approximate", we are able to prove that a small multiple of the commitment opens to the same small multiple of a member of the set. For our application, this is good enough.

[2] The ZKPoK is a Fiat-Shamir transformation of a Σ-protocol, and so the message that the group member signs is simply added into the random oracle input.

[3] The reason that signature schemes using cryptographic hash functions are not suitable is that their lack of algebraic structure makes it very difficult to construct efficient proofs of knowledge that prove something about the identity i when it is an input to the random oracle.

While lattice-based signatures in the standard model are inefficient, there is a much more efficient *selectively-secure* lattice-based digital signature scheme that is implicit from the works of [ABB10,Boy10]. A selectively-secure signature scheme is one in which the adversary declares the message that he will forge on prior to seeing the public key. A scheme like this can be converted to a regular signature scheme with a reduction loss of $1/|S|$, where S is the message space simply by guessing the message that the Adversary will forge on. Thus for small message spaces, this becomes a signature scheme with a meaningful reduction from hard lattice problems.

There have been several previous papers that utilized the above-mentioned selectively secure scheme for group signatures and related applications [NZZ15,BCN17,BCN18]. In those papers, the techniques for proving that the identity i is in a particular set resulted in either a significant increase in the proof size and/or a very noticeable loss in the tightness of the proof.

Roughly-speaking, the reason that the construction in [BCN17,BCN18] is less efficient than ours is that in order to prove that the message is in a small set, the space of the messages and the challenges is restricted to a small-dimensional sub-ring. In order to have negligible soundness error, it is thus necessary to either increase the size of the coefficients of the challenge or to repeat the protocol several times – both of these solutions end up increasing the size of the signatures. Our technique, on the other hand, does not require to reduce the degree of the challenge. Additionally, the construction in [BCN17,BCN18] requires the identity to have small coefficients, whereas the proof of knowledge has "slack" and proves that the identity has somewhat larger coefficients – this further decreases the tightness of the reduction. The construction in the current paper uses a commitment scheme in which the messages need not have small coefficients [BDL$^+$18] and so the slack in the zero-knowledge proof (which affects the randomness used in the commitment) does not affect the size of the message coefficients.

Concurrently with our paper, Katz et al. [KKW18] presented a construction of a group signature scheme based only on the assumption that AES-256 and SHA-256 behave as random oracles. For small group sizes (approximately 2^{13}), the sizes of the signatures are in fact smaller than ours (while the signing time is still around 8 times longer). For larger group sizes, however, our signatures are smaller. Additionally the opening procedure of [KKW18] may be prohibitive for large groups as it is linear in the group size.

Our Contribution. In the present work we show how our new proofs for stability under automorphisms allow for a fairly natural, at a high level, group signature construction based on the hardnes of lattice problems. In particular, the set of identities will be exactly those elements in $\mathcal{R}_q$ that are preserved under some set of automorphisms. The size of these sets can be small (as small as q), and so we will only lose a factor of the group size in the reduction. The idea for the ZKPoK will then be to do the proof of knowledge with the *commitments of* i rather than with i (thus not revealing the identity) and prove that our commitments are to elements in the appropriate set of identities – for this we will use the module-homomorphic properties of our commitment scheme – i.e. if $i \cdot s = u$, for small s, then $Com(i;\mathbf{r}) \cdot s = \begin{bmatrix} 0 \\ u \end{bmatrix} + Com(0;\mathbf{r}')$. The encryption to the opener can be done using the main idea from the verifiable encryption scheme from [LN17]. A point of note is that the selectively-secure signature scheme requires that the messages come from a set S such that the difference of any two elements from the set is invertible. This is compatible with our definition of sets because they turn out to be subfields of the original ring $\mathcal{R}_q$.

Instantiating our scheme with concrete parameters gives group signatures of around 580 KB, which is almost a 2 order of magnitude reduction from [LLNW16] and about an order of magnitude

reduction over the concurrent construction in [BCN18].[4] Our main technique should also applicable to a variety of other privacy applications that require similar proofs of knowledge. For example, one should be able to apply these techniques in a very similar manner to the constructions of anonymous credentials as in [BCN17].

To demonstrate the practicality of our group signature scheme, we have implemented it in C. On a laptop with an Intel Skylake i7 processor, the implementation needs 428.7 ms to generate a group public key and one member secret key. Signing a message takes 404.5 ms and the signature can be verified in 169.1 ms. For the signing keys of the group members one needs to sample preimages of a linear map from a discrete Gaussian distribution. This can, in theory, be done with the GPV sampling algorithm from [GPV08], but it requires computing the Gram-Schmidt decomposition of a basis which is a prohibitively expensive operation in the high dimensions required for our scheme. We have therefore implemented the Fast Fourier Orthogonalization algorithm from [DP16] adapted to cyclotomic fields which computes a compact LDL^* decomposition of the basis that is used in a Fast Fourier Nearest Plane algorithm, also from [DP16], to sample preimages. This was done before in the Falcon signature scheme [PFH$^+$18], but contrary to that implementation, ours supports arbitrary precision complex arithmetic since double precision is not enough for our larger moduli.

In Sections 1.1 and 1.2, we give high-level sketches of our main results – the proof of stability under automorphisms (the full details of which are in Sections 3 and 4) and the construction of the group signature scheme (the full details of which are in Section 5).

1.1 Commitments and Proofs of Automorphism Stability

We will use a particular instantiation of the commitment scheme from [BDL$^+$18] where the common reference string public key is

$$\begin{bmatrix}\mathbf{a}_1\\ \mathbf{a}_2\end{bmatrix} = \begin{bmatrix}1 & a_1 & a_2\\ 0 & 1 & a_3\end{bmatrix} \in \mathcal{R}_q^{2\times 3} \tag{1}$$

and the commitment to a polynomial $\mu \in \mathcal{R}_q$ requires us to pick a random polynomial $\mathbf{r} \in \mathcal{R}_q^3$ with small coefficients and output the commitment

$$Com(\mu; \mathbf{r}) = \begin{bmatrix}t_1\\ t_2\end{bmatrix} = \begin{bmatrix}\mathbf{a}_1\\ \mathbf{a}_2\end{bmatrix} \cdot \mathbf{r} + \begin{bmatrix}0\\ \mu\end{bmatrix}. \tag{2}$$

Using the "Fiat-Shamir with Aborts" zero-knowledge proof technique [Lyu09,Lyu12], one can prove the knowledge of a polynomial vector $\bar{\mathbf{r}}$ with coefficients somewhat larger than those in $\mathbf{r}$, and a

[4] Table 1 of the conference version of [BCN18] gives a signature size of 1.72MB for 80-bits of security. This security, however, has only been calculated for the *traceability* part of the security of group signatures (i.e. it's not possible to produce a signature that cannot be traced by the opener to a particular user) which is based on the Ring-SIS problem. The *anonymity* security notion (i.e. the identity of the signing group member should remain secret) is based on the hardness of the Ring-LWE problem, and it does not appear that this has been accounted for in the parameter setting. In particular, the Ring-LWE instance with $-1/0/1$ secret/noise coefficients in [BCN18] is over the ring $\mathbb{Z}_q[X]/(X^d+1)$ where $q \approx 2^{115}, d = 2048$ (for comparison, our ring has $q \approx 2^{80}$ and $d = 4096$, both of which significantly increase the complexity of the problem). By our calculation, d would need to be increased from 2048 to 8192 for the claimed security in [BCN18], and this would increase the signature size by approximately a factor of 4, making it a little more than an order of magnitude larger than in the current work.

polynomial c with $-1/0/1$ coefficients such that

$$c \cdot \begin{bmatrix} t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \end{bmatrix} \cdot \bar{\mathbf{r}} + \begin{bmatrix} 0 \\ c\mu \end{bmatrix}. \tag{3}$$

Even though $\|\bar{\mathbf{r}}\| > \|\mathbf{r}\|$ and there is an extra term c present, opening the commitment with $\bar{\mathbf{r}}$ and c can still be binding if the parameters are appropriately set.

In our work, we will additionally show how for certain sets $S \subset \mathcal{R}_q$, we can prove (3) and additionally show that $\mu \in S$. The sets for which we are able to show this are those that are preserved under the automorphisms of the cyclotomic number field $K_m = \mathbb{Q}[X]/(\Phi_m(X))$. For example, if $\Phi_m(X) = X^d + 1$ (where $m = 2d$ is a power of 2), then the $\phi(m) = d$ automorphisms are $\sigma_j : X \to X^j$ for all odd integers $0 < j < 2d$.

We give a protocol for a proof of knowledge as for (3) which additionally allows us to prove that $\sigma_j(\mu) \equiv \mu (\mathrm{mod}\, q)$. Our proof is derived from a generalization of a zero-knowledge proof of linear relations of commitments in (2) from [BDL+18]. In particular, we can show how to prove linear relations for messages μ_i for commitments under *distinct* public keys $\begin{bmatrix} \mathbf{a}_1^{(i)} \\ \mathbf{a}_2^{(i)} \end{bmatrix}$. For proving that a commitment in (2) is closed under an automorphism σ then requires proving that the two commitments $\begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$ and $\begin{bmatrix} \sigma(t_1) \\ \sigma(t_2) \end{bmatrix}$, under the respective public keys $\begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \end{bmatrix}$ and $\begin{bmatrix} \sigma(\mathbf{a}_1) \\ \sigma(\mathbf{a}_2) \end{bmatrix}$, are both to the same message – which implies that $\mu = \sigma(\mu)$. The communication complexity of this protocol (of the non-interactive version) involves sending essentially one extra vector of the same size as $\bar{\mathbf{r}}$ for every automorphism. It is therefore beneficial to not have to prove stability under too many σ_j.

Galois Theory gives us the exact group structure of the automorphisms and specifies which subsets of K_m are preserved under them. It furthermore allows us to determine the minimum set of automorphisms that are needed to generate the group. For example, an element $v \in \mathbb{Q}[X]/(X^d+1)$ is a constant if and only if $\sigma_5(v) = \sigma_{m-1}(v) = v$. Thus, proving stability under σ_5 and σ_{m-1} would prove that we have committed to a constant μ. As another example, $v \in \mathbb{Q}[X]/(X^d + 1)$ is of the form $\alpha + \beta X^{d/2}$ for $\alpha, \beta \in \mathbb{Q}$ if and only if $\sigma_5(v) = v$. Thus for this set of size q^2, it is only necessary to prove stability under one automorphism.

The situation in our case is made more complicated due to the fact that we give proofs that $\sigma_j(v) \equiv v (\mathrm{mod}\, q)$, while Galois Theory only tells us about stability of sets with coefficients over $\mathbb{Q}$ (i.e. without reduction modulo q). So one could fathom that $\sigma_j(v) = v$ modulo q but $\sigma_j(v) \neq v$. We show, however, that one can find primes q such that subsets of $\mathbb{Z}_q[X]/(\Phi_m(X))$ have the same properties under automorphisms as subsets of $\mathbb{Z}[X]/(\Phi_m(X))$. In particular, we can build subsets of size q^i for all $i \mid \phi(m)$. For the particular case of rings of the form $\mathbb{Z}_q[X]/(X^d+1)$, this implies that one can have a generating set of 1 or 2 automorphisms for particular sets of size $q, q^2, q^4, \ldots, q^{d/2}$. We also give concrete descriptions of these sets and show efficient procedures for generating elements in them.

1.2 Group Signatures

We now give a high level overview of how one would use the techniques to construct a group signature scheme. The master group public key of the setup authority will be a public key for the selectively secure signature scheme from [ABB10] adapted to polynomial rings:

$$[\ \mathbf{a} \mid \mathbf{b}\ ], u = \mathbf{a} \cdot \mathbf{s}_1' + \mathbf{b} \cdot \mathbf{s}_2' + \mathbf{a}_2 \cdot \mathbf{s}_3' \tag{4}$$

where $\mathbf{a} = [a\ a']$ for a uniformly-random a, a' and $\mathbf{b} = [b_1\ b_2] = \mathbf{a} \cdot \begin{bmatrix} r_1 & r_2 \\ e_1 & e_2 \end{bmatrix}$ where r_i, e_i are polynomials in $\mathcal{R}_q$ with small coefficients such that (a, b_1, b_2) are indistinguishable from random based on the hardness of the Ring-LWE problem. The group member identities are polynomials $i \in S \subseteq \mathcal{R}_q$ where the set S is preserved under some set of automorphisms of $\mathcal{R}_q$. The secret key of a user with identity i consists of vectors $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3$ that are generated by the setup authority using his secret trapdoor key $\mathbf{R} = \begin{bmatrix} r_1 & r_2 \\ e_1 & e_2 \end{bmatrix}$. The setup authority first picks a short vector $\mathbf{s}_3$ from a particular distribution, and then "pre-image samples" short vectors $\mathbf{s}_1, \mathbf{s}_2$ such that

$$[\ \mathbf{a}\ |\ \mathbf{b} + i \cdot [\ 1\ \sqrt{q}]\ ] \cdot \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \end{bmatrix} = u + \mathbf{a}_2 \cdot \mathbf{s}_3. \tag{5}$$

The matrix $[\ 1\ \sqrt{q}\ ]$ is sometimes referred to as the "gadget matrix"[5] that allows for efficient pre-image sampling of short vectors $\mathbf{s}_1, \mathbf{s}_2$ as in (5) for all $i \neq 0$. The procedure for computing such vectors in a way that produces a distribution independent of the secret key $\mathbf{R}$ is described in [MP12]. When $i = 0$, the setup authority can output $\mathbf{s}'_1, \mathbf{s}'_2, -\mathbf{s}'_3$ as the key.[6] The purpose of the $\mathbf{a}_2 \cdot \mathbf{s}_3$ part of the construction is only necessary for the security proof – it's unclear if it truly adds any security in practice. The performance downside of including this term is fairly small — one needs to do an extra sampling of $\mathbf{s}_3$ and multiplication by $\mathbf{a}_2$ in the key generation, and the size of the solution to the Ring-SIS problem in the security proof is a small (virtually inconsequential) additive factor larger. For the following high-level overview, the reader can just take $\mathbf{s}_3, \mathbf{a}_2 = \mathbf{0}$.

Signing. The high level idea for signing is for the user with identity i to prove knowledge of $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3$ that satisfies (5). If the proof of knowledge is a Σ-protocol, then it can be converted into a non-interactive proof using the Fiat-Shamir heuristic, which turns the Σ-protocol into a signature scheme if one inputs the message into the random oracle. The main difficulty in all group signature constructions lies in doing this proof without revealing i.

To hide i in our proof, the signer will commit to i and $i\sqrt{q}$ using the commitment scheme from Section 1.1 and publish his commitments as part of the signature. The main observation is that

$$\left[\begin{bmatrix} \mathbf{0} \\ \mathbf{a} \end{bmatrix} \ |\ \begin{bmatrix} \mathbf{0} \\ \mathbf{b} \end{bmatrix} + [Com(i; \mathbf{r})\ Com(i\sqrt{q}; \mathbf{r}')] \right] \cdot \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \end{bmatrix} = \begin{bmatrix} 0 \\ u + \mathbf{a}_2 \cdot \mathbf{s}_3 \end{bmatrix} + \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \end{bmatrix} \cdot \tilde{\mathbf{r}}. \tag{6}$$

The signer will give an approximate ZKPoK of the short randomnesses $\mathbf{r}, \mathbf{r}'$ that open the commitments to $i, i\sqrt{q}$ and also that $i \in S$.[7] In other words, he'll prove knowledge of

[5] The gadget matrix for polynomials is more generally defined as $[1\ q^{1/b}\ q^{2/b}\ \ldots\ q^{(b-1)/b}]$ for some b. In this work, we take $b = 2$. Also, we write $\sqrt{q}$ instead of $\lceil \sqrt{q} \rfloor$ for improved readability.

[6] For $i \neq 0$, the setup authority is able to output many possible valid $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3$ using his trapdoor and the gadget matrix. For $i = 0$, however, the gadget matrix disappears and so the setup authority is only able to return one $\mathbf{s}'_1, \mathbf{s}'_2, -\mathbf{s}_3$ that he "planted" when creating u in (4). For the security proof, it will be necessary that the distribution for all i is the same, and so for this reason, we make the pre-image sampling procedure for all i deterministic. In other words, the randomness used in the sampling will be derived by the setup authority using a keyed PRF whose input depends on i.

[7] Due to the slack in our zero-knowledge protocols, the proofs will be for larger values of $\mathbf{r}, \mathbf{r}'$ than those used in the commitments. But for simplicity of exposition in the introduction of this paper, we will use the same notation.

$$\begin{bmatrix}\mathbf{a}_1\\ \mathbf{a}_2\end{bmatrix} \cdot [\mathbf{r}\ \mathbf{r}'] + \begin{bmatrix}0 & 0\\ ci & ci\sqrt{q}\end{bmatrix} = c \cdot \begin{bmatrix}t_1^{(1)} & t_1^{(2)}\\ t_2^{(1)} & t_2^{(2)}\end{bmatrix}. \tag{7}$$

In parallel, the signer will also prove that

$$\left[\mathbf{a} \mid \mathbf{b} + [t_2^{(1)}\ t_2^{(2)}]\right] \cdot \begin{bmatrix}\mathbf{s}_1\\ \mathbf{s}_2\end{bmatrix} = cu + \mathbf{a}_2 \cdot \mathbf{r}''.^{8} \tag{8}$$

Multiplying (8) by c and combining with (7) produces the equation

$$\left[\mathbf{a} \mid c\mathbf{b} + \mathbf{a}_2 \cdot [\mathbf{r}\ \mathbf{r}'] + c \cdot [i\ i\sqrt{q}]\right] \cdot \begin{bmatrix}c\mathbf{s}_1\\ \mathbf{s}_2\end{bmatrix} = c^2 u + c\mathbf{a}_2 \cdot \mathbf{r}''. \tag{9}$$

We can then show that if an Adversary can produce polynomial vectors $\mathbf{r}, \mathbf{r}', \mathbf{s}_1, \mathbf{s}_2, \mathbf{r}'', c$ with small coefficients that satisfy the above equation, then he is able to solve the Ring-SIS problem. The proof is very similar to the proof of selective security for the signature scheme of [ABB10] . Intuitively, suppose that the Adversary in the impersonation game produces a solution (i.e. the extracted values from the PoK) for (9) for $i = 0$. Then, using the fact that $\mathbf{b} = \mathbf{aR}$ and $u = \mathbf{a} \cdot \mathbf{s}_1' + \mathbf{b} \cdot \mathbf{s}_2' + \mathbf{a}_2 \cdot \mathbf{s}_3'$ and writing $\mathbf{R}' = [\mathbf{r}\ \mathbf{r}']$, (9) can be rewritten as

$$\mathbf{a} \cdot \left(c\mathbf{s}_1 + c\mathbf{R}\mathbf{s}_2 - c^2\mathbf{s}_1' - c^2\mathbf{R}\mathbf{s}_2'\right) + \mathbf{a}_2 \cdot \left(\mathbf{R}'\mathbf{s}_2 - c\mathbf{r}'' - c^2\mathbf{s}_3'\right) = 0, \tag{10}$$

which is a solution to the Ring-SIS problem because the coefficients of all the terms in parentheses are small relative to q.

Of course, the Adversary is not guaranteed to impersonate on identity $i = 0$, but may choose an arbitrary $i' \in S$. To handle this, we use the standard "puncturing" technique. In the security proof we would not choose $\mathbf{b} = \mathbf{aR}$ as part of the public key, but we rather pick a uniformly-random "guess" $i' \in S$, and set $\mathbf{b} = \mathbf{aR} - [i'\ i'\sqrt{q}]$ as part of the public key. It's not hard to see that if the Adversary produces a solution for (9) with $i = i'$, then one again obtains the same Ring-SIS solution as in (10). If the Adversary cannot tell how $\mathbf{b}$ was constructed, even after querying for preimages $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3$, then there is exactly a $1/|S|$ chance that $i = i'$. Therefore there is a $1/|S|$ loss in the tightness of the security reduction.

For the purposes of allowing opening, the signer will also create a Ring-LWE encryption of the three polynomials comprising the vector $\mathbf{r}$ used in the commitment of i in (7) using the one-shot verifiable encryption / proof of plaintext knowledge from [LN17] combined with the proofs of knowledge for (7). The reason that we encrypt $\mathbf{r}$ rather than i is that the coefficients in $\mathbf{r}$ are small, whereas i comes from a set that is stable under some automorphism, and such sets contain elements with large coefficients. Once the opener decrypted $\mathbf{r}$, he knows from (7) that $\mathbf{a}_1 \cdot \mathbf{r} = c \cdot t_1^{(1)}$, and so he can recover c. Then using this c, he can recover i from the equality $\mathbf{a}_2 \cdot \mathbf{r} + ci = c \cdot t_2^{(1)}$.

Reducing the Commitment Size. To reduce the size of the signature, we can slightly modify the (2) so that it works over two different moduli, one for the top and another for the bottom part (call them q_1 and q_2 respectively). In our group signature scheme, the value of q_2 needs to be large due to the fact that the Ring-SIS solution in (10) is fairly large itself. The value of q_1, on the other hand, only needs to be set so that the commitments to i and $i\sqrt{q}$ in (7) are binding and hiding.

[8] Notice that we combined the $\mathbf{a}_2 \cdot \mathbf{s}_3$ term with $\mathbf{a}_2 \cdot \tilde{\mathbf{r}}$ term to obtain $\mathbf{a}_2 \cdot \mathbf{r}''$. This was the reason that we used exactly $\mathbf{a}_2$ from the commitment scheme in the key generation in (5).

Since a smaller q_1 will result in smaller sizes of $t_1^{(1)}$, $t_1^{(2)}$ in the commitment, it is sensible to set it as small as possible. We show that our proofs of automorphism stability still work if the two moduli are different.

1.3 Acknowledgements.

We thank the anonymous reviewers for many suggestions and corrections that helped improve the presentation of the paper. This research has been supported by the SNSF ERC Transfer Grant CRETP2-166734 FELICITY.

2 Preliminaries

2.1 Notation

Throughout this paper we will consider a polynomial ring $\mathcal{R}$ of the form $\mathbb{Z}[X]/(\Phi_m(X))$, with $\Phi_m(X)$ the m^{th} cyclotomic polynomial. We will denote elements of $\mathcal{R}$ by lowercase letters, vectors over $\mathcal{R}$ in bold lowercase and matrices over $\mathcal{R}$ in bold uppercase. e.g. $\mathbf{A} = \left[\mathbf{a}_1 \mid \ldots \mid \mathbf{a}_k\right] \in \mathcal{R}^{l\times k}$ with $\mathbf{a}_i = (a_{i1}, \ldots, a_{im})^T \in \mathcal{R}^l$, remark that we consider column vectors over $\mathcal{R}$. We will consider the norm of elements in $\mathcal{R}$ to be $\|a\| = |a|$ if $a \in \mathbb{Z}$, and $\|a\| = \sqrt{\sum a_i^2}$ if $a = \sum a_i X^i \in \mathbb{Z}[X]/(X^d+1)$. We extend the notation to vectors and matrices $\|\mathbf{a}\| = \sqrt{\sum \|a_i\|^2}$, $\|\mathbf{A}\| = \sqrt{\sum \|\mathbf{a}_i\|^2}$. We will also consider the quotient ring $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$ for which the norm of an element $\mathcal{R}_q$ will be the norm of its unique representative $\mathcal{R}$ with coefficient s in $\left[-\frac{q-1}{2}, \frac{q-1}{2}\right]$.

We will also consider the operator norm of matrices in $\mathcal{R}^{l\times k}$ defined as $s_1(\mathbf{A}) = \underset{\|x\|\neq 0}{max} \left(\frac{\|\mathbf{Ax}\|}{\|\mathbf{x}\|}\right)$. For $\beta \in \mathbb{R}$, we define the set S_β to be the set of all polynomials of infinity norm less than beta, i.e. $S_\beta = \{a \in \mathcal{R} \mid \|a\|_\infty \leq \beta\}$.

2.2 Invertibility of Challenges

For many of our protocols we will use the challenge set of polynomials

$$\mathcal{C} = \{c \in \mathcal{R} \mid \|c\|_1 = \kappa, \|c\|_\infty = 1\}$$

and define $\bar{C}$ as in Table 2. We will sometimes need the fact that all polynomials in $\bar{C}$ are invertible over some particular polynomial ring. The following lemma from [LS18] guarantees such a property for well chosen power-of-two cyclotomics. A similar theorem holds for general cyclotomics.

Lemma 2.1. *[LS18, Corollary 1.2] Let $d \geq k > 1$ be powers of 2 and $q \equiv 2k+1 (\mathrm{mod}\, 4k)$ be a prime. Then the polynomial X^d+1 any c in $\mathbb{Z}_q[X]/(X^d+1)$ such that $0 < \|c\| < \frac{1}{\sqrt{k}} \cdot q^{1/k}$ has an inverse in the ring.*

2.3 Norms and Gaussians

While we prove results for any cyclotomic ring $\mathcal{R} = \mathbb{Z}[X]/(\Phi_m(X))$ in section 3, we will, for simplicity, restrict the construction of section 5 to power-of-two cyclotomics, i.e. rings of the form $\mathbb{Z}[X]/(X^d+1)$ for d a power of two. This way we can use the euclidean norm rather than the

embedding norm.
Define the function $\rho_\sigma(x) = \exp\left(\frac{-x^2}{2\sigma^2}\right)$ and the discrete Gaussian distribution centered in $v \in \mathbb{R}$ over the integers, D_σ, as

$$D_\sigma(x) = \frac{\rho(x)}{\rho(\mathbb{Z})} \text{ where } \rho(\mathbb{Z}) = \sum_{v\in\mathbb{Z}} \rho(v).$$

We will write $x \leftarrow D_{\mathcal{R},\sigma}$ to mean that every coefficient of the polynomial $x \in \mathcal{R}$ is distributed according to D_σ. When clear from context we will simply write this as $x \leftarrow D_\sigma$.

Using the tail bounds for the 0-centered discrete Gaussian distribution (cf. [Ban93]), we can show that for any $\sigma > 0$, $x \leftarrow D_\sigma$ is likely to be close to σ. Namely, for any $k > 0$ it holds that

$$\Pr_{x\leftarrow D_\sigma}[|x| > k\sigma] \leq 2e^{-k^2/2}, \tag{11}$$

and when $\mathbf{x}$ is drawn from D_σ^n, we have

$$\Pr_{\mathbf{x}\leftarrow D_\sigma^n}[\|\mathbf{x}\| > \sqrt{2n}\cdot\sigma] < 2^{-n/4}. \tag{12}$$

We give an important lemma on rejection sampling which will guarantee that the responses used in our zero-knowledge protocols do not leak information.

Algorithm 1 Rej$(\mathbf{z}, \mathbf{b}, \sigma)$

$u \leftarrow [0, 1)$
if $u > \frac{1}{3}\cdot\exp\left(\frac{-2\langle\mathbf{z},\mathbf{b}\rangle+\|\mathbf{b}\|^2}{2\sigma^2}\right)$ **then**
  return 0
else
  return 1
end if

Lemma 2.2 ([Lyu12]). *Let V be a subset of $\mathcal{R}^n$ with elements of norm less than T, let h be a distribution of V. $\mathbf{b} \in \mathcal{R}^n$. Consider a procedure that samples a $\mathbf{y} \leftarrow D_\sigma^n$ and then returns the output of Rej$(\mathbf{z} := \mathbf{y} + \mathbf{b}, \mathbf{b}, \sigma)$ where $\sigma \geq 11\|\mathbf{b}\|$. The probability that this procedure outputs 1 is within 2^{-100} of $1/3$. The distribution of $\mathbf{z}$, conditioned on the output being 1, is within statistical distance 2^{-100} of D_σ^n.*

2.4 M-SIS and M-LWE

In this section we introduce the hard problems on which our schemes rely. We will be using the "Module" (or "Generalized") variants of the LWE and SIS problems, introduced in [BGV12,LS15]. These are generalizations of the usual LWE and SIS problems in the sense that, while the former are defined over the ring $\mathbb{Z}_q$, the latter are instantiated over polynomial rings $\mathcal{R}_q$. Since we will instantiate our scheme with power-of-two cyclotomic ring we only define those problems for this setting. *M-SIS* and *M-LWE* can be defined for any ring but the definitions are more cumbersome (see [LS15]). For simplicity we will consider the *M-LWE* problem in which the secret and the randomness are sampled in S_1 (i.e. uniformly from the set of elements bounded in infinity norm) this assumption is common in practical cryptographic schemes, e.g. [BDK+18,DKL+18,LN17].

$\mathcal{R}$	The cyclotomic ring $= \mathbb{Z}[X]/(X^d+1)$
d	The dimension of $\mathcal{R}$
q_1, q_2	The moduli used in our commitment
k	Width (over $\mathcal{R}$) of the commitment matrices
n	Height (over $\mathcal{R}$) of the commitment matrix $\mathbf{A}_1$
l	Dimension (over $\mathcal{R}$) of the message space
$\mathcal{C}, \kappa$	Challenge set $\mathcal{C} = \{c \in \mathcal{R} \mid \|c\|_1 = \kappa,\ \|c\|_\infty = 1\}$
$\bar{\mathcal{C}}$	The set of differences $\mathcal{C} - \mathcal{C}$ except 0
s,r	The standard deviation of the secret keys in our group signature
p	The plaintext modulus for our verifiable encryption
Q	The ciphertext modulus for our verifiable encryption

Table 2. The parameters of our commitment, zero-knowledge, group signature, and verifiable encryption schemes

Definition 2.3 (M-SIS [LS15]). *The $M\text{-}SIS_{q,n,m,\beta}$ problem (over an implicit ring $\mathcal{R}$) is defined as follows. Given $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ sampled uniformly at random, find $\mathbf{z} \in \mathcal{R}^m$ such that $\mathbf{Az} = 0$ and $0 < \|\mathbf{z}\| \leq \beta$.*

Definition 2.4 (M-LWE [BGV12,LS15]). *The decision $M\text{-}LWE_{q,m,n}$ problem (over an implicit ring $\mathcal{R}$) is defined as follows. Let $\mathbf{s} \xleftarrow{\$} S_1^n$, let $A_{q,\mathbf{s}}$ be the distribution obtained by sampling $\mathbf{a} \xleftarrow{\$} \mathcal{R}_q^n$, $e \xleftarrow{\$} S_1$, and returning $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathcal{R}_q^n \times \mathcal{R}_q$. The goal is to distinguish between m samples from either $A_{q,\mathbf{s}}$ or $\mathcal{U}\left(\mathcal{R}_q^n, \mathcal{R}_q\right)$.*

The number of samples m in the above definition of $M\text{-}LWE$ does not have any known effect on the hardness of the problem unless it is large enough (at least as large as $(nd)^3$) in order for the Arora-Ge linearization attack to apply [AG11]. In this paper, the number of samples will always be significantly lower than this (only linear in nd), and so we will omit the m and simply write $M\text{-}LWE_{q,n}$.

Definition 2.5 (NTRU). *The $NTRU_{q,r}$ problem (over an implicit ring $\mathcal{R}$) is defined as follows. The distribution A is defined by sampling ring elements $f, g \xleftarrow{\$} D_r$ and outputting $h = f/g$, if g is invertible in $\mathcal{R}_q$ (otherwise, re-sample g). The $NTRU_{q,r}$ problem is to distinguish h from a random element in $\mathcal{R}_q$.*

2.5 Commitments

We use a variant of the commitment scheme of [BDL+18] with *computationally* hiding (based on M-LWE) and binding (based on M-SIS) commitments. As per [BDL+18], this allows for the most efficient setting of parameters. We will also slightly modify the original scheme by consider the top part of the commitment equation $\mathbf{A}_1\mathbf{r} = \mathbf{t}_1$ modulo a prime q_1 while the bottom part $\mathbf{A}_2\mathbf{r} + \mathbf{m} = \mathbf{t}_2$ will be modulo a different prime q_2.

The reason for using two different moduli is that the hardness of the group signature is, in part, based on the hardness of the M-SIS problem modulo q_2. The solution to this M-SIS problem that

we're able to extract is large, and so it makes sense to set q_2 larger and make the M-SIS problem harder. Thus we're taking $q_2 > q_1$ not for the purposes of making the commitment scheme harder, but due to the fact that q_2 also comes up in the hardness of a different part of the protocol.

CKeyGen: Create the public parameters $\mathbf{A}_1 \in \mathcal{R}_{q_1}^{n \times k}$ and $\mathbf{A}_2 \in \mathcal{R}_{q_2}^{l \times k}$ such that:

$$\mathbf{A}_1 := \left[\mathbf{I}_n \ \mathbf{A}_1'\right], \text{ where } \mathbf{A}_1' \xleftarrow{\$} \mathcal{R}_{q_1}^{n \times (k-n)}$$
$$\mathbf{A}_2 := \left[\mathbf{0}^{l \times n} \ \mathbf{I}_l \ \mathbf{A}_2'\right], \text{ where } \mathbf{A}_2' \xleftarrow{\$} \mathcal{R}_{q_2}^{l \times (k-n-l)}$$

Commit: To commit to a message $\mathbf{m} \in \mathcal{R}_{q_2}^l$, sample a randomness $\mathbf{r} \xleftarrow{\$} S_1^k$ and output:

$$Com(\mathbf{m}; \mathbf{r}) := \begin{bmatrix}\mathbf{t}_1 \\ \mathbf{t}_2\end{bmatrix} = \begin{bmatrix}\mathbf{A}_1 \\ \mathbf{A}_2\end{bmatrix} \mathbf{r} + \begin{bmatrix}\mathbf{0} \\ \mathbf{m}\end{bmatrix}$$

Open: A valid opening of $\mathbf{t}_1, \mathbf{t}_2 \in \mathcal{R}_{q_1}^n \times \mathcal{R}_{q_2}^l$ consists of a message $\mathbf{m} \in \mathcal{R}_{q_2}^l$, a randomness $\mathbf{r} \in \mathcal{R}^k$, and a polynomial $c \in \mathcal{R}$ such that:

$$c \begin{bmatrix}\mathbf{t}_1 \\ \mathbf{t}_2\end{bmatrix} = \begin{bmatrix}\mathbf{A}_1 \\ \mathbf{A}_2\end{bmatrix} \mathbf{r} + c \begin{bmatrix}\mathbf{0} \\ \mathbf{m}\end{bmatrix}$$

With $\|\mathbf{r}\| \leq B_{com}$ and $c \in \bar{\mathcal{C}}$.

The parameters $n, k \in \mathbb{Z}$ must be set so that the commitment is hiding and binding. The parameter $l \in \mathbb{Z}$ dictates the size of the message space. We remark that an opening of a commitment does not simply consist of a message and a randomness but also includes a small polynomial c which multiplies the commitment. The reason is that when doing zero-knowledge proofs for commitments the knowledge extractor will not be able to extract an exact opening of $\mathbf{t} = Com(\mathbf{m}; \mathbf{r})$ but only an opening of $c\mathbf{t}$ where c will be the difference of two challenges. We prove that our commitment is hiding and binding in Section 6.1.

2.6 Trapdoor sampling

Recall in the group signature the manager has to sample short vectors $\mathbf{s}_1, \mathbf{s}_2$ such that

$$[\ \mathbf{a} \mid \mathbf{b} + i \cdot [\ 1 \ \lceil\sqrt{q}\rfloor]\ ] \cdot \begin{bmatrix}\mathbf{s}_1 \\ \mathbf{s}_2\end{bmatrix} = u + \mathbf{a}_2 \cdot \mathbf{s}_3.$$

By [MP12, Lemma 5.3] there exists a basis $\mathbf{S} \in \mathbb{Z}^{4d \times 4d}$ for the lattice $\Lambda^\perp = \{\mathbf{x} \in \mathcal{R}^4 \mid [\ \mathbf{a} \mid \mathbf{b} + i \cdot [\ 1 \ \lceil\sqrt{q}\rfloor]\ ] \cdot \mathbf{x} \equiv 0 \pmod q\}$ whose Gram-Schmidt orthogonalization fulfills $\|\tilde{\mathbf{S}}\| \leq (s_1(\mathbf{R})+1)\sqrt{\delta^2+1}$ with $\delta = \lceil\sqrt{q}\rfloor$. For random matrices in $\mathbb{Z}^{2d \times 2d}$, the expected value of the largest singular is $2 \cdot \sqrt{2d}$. We found experimentally that for our structured matrix $\mathbf{R}$ it is slightly larger but less than $3\sqrt{d}$.

Now a short preimage $[\ \mathbf{s}_1 \ \mathbf{s}_2\ ]^T$ of $u + \mathbf{a}_2 \cdot \mathbf{s}_3$ can be sampled by computing an arbitrary solution $[\ \mathbf{x}_1 \ \mathbf{x}_2\ ]$, expressing this solution in the basis $\mathbf{S}$ of the orthogonal lattice $\Lambda^\perp$ and decoding using the randomized nearest plane discrete Gaussian sampler from [GPV08]. This gives a solution that is distributed as a discrete Gaussian with parameter $s = 2 \cdot \|\tilde{S}\| \leq 2(3\sqrt{d}+1)\sqrt{\delta^2+1}$ and is statistically independent from the trapdoor.

In the security proof of the group signature scheme we also need that one can sample preimages of matrices $[\ 1 \ b\ ]$ of NTRU lattices with the help of a trapdoor $b = f/g$ with short f, g. Using [DLP14] we can assume that f, g lead to a basis with maximum Gram-Schmidt norm less than $1.17\sqrt{q}$. Then we can compute discrete Gaussian preimages with Gaussian parameter $r = 2 \cdot 1.17\sqrt{q}$.

3 Galois Group Structure of Cyclotomic Rings

In Section 4 we will want to construct a proof of knowledge of an opening $\mu \in \mathcal{R}_q = \mathcal{R}/q\mathcal{R}$ to a commitment with the additional guarantee that μ lies in a certain subset of $\mathcal{R}_q$. For our purposes of group signatures we need μ to be invertible and therefore want the subset to be a subfield. We do this by proving that μ is fixed by certain automorphisms that we construct from the Galois automorphisms of our cyclotomic field K. This then shows that μ is contained in $\mathcal{S}_q = \mathcal{S}/q\mathcal{S}$ where $\mathcal{S} \subset \mathcal{R}$ is the ring of integers of a subfield of K. Here we can arrange for the prime number q to stay inert in $\mathcal{S}$ so that $\mathcal{S}_q$ is a field.

3.1 Generic Cyclotomic Rings

We have the following setup. $K = \mathbb{Q}[X]/(\Phi_m(X))$ is the m-th cyclotomic number field of degree $d = \varphi(m)$ with ring of integers $\mathcal{R} = \mathbb{Z}[X]/(\Phi_m(X))$. Let $L \subset K$ be a subfield of K, not necessarily cyclotomic, with ring of integers $\mathcal{S} \subset \mathcal{R}$. We thus have the following diagram of rings and fields.

$$\begin{array}{ccc} \mathcal{R} & \subset & K \\ | & & | \\ \mathcal{S} & \subset & L \\ | & & | \\ \mathbb{Z} & \subset & \mathbb{Q} \end{array}$$

Finally suppose q is a prime number that is inert in L, i.e. such that $\mathcal{S}_q = \mathcal{S}/q\mathcal{S} \subset \mathcal{R}_q$ is a field, and unramified in K.

The automorphisms of K form a group under composition called the Galois group of K which we denote by $G = \mathrm{Gal}(K/\mathbb{Q})$. It is easy to see that all automorphisms fix the rational numbers $\mathbb{Q}$, $\sigma(x) = x$ for all $\sigma \in G$ and $x \in \mathbb{Q}$. Conversely, cyclotomic fields are special among general number fields in that they are Galois over $\mathbb{Q}$ meaning that only the elements of $\mathbb{Q}$ are fixed by all automorphisms. The Galois group of cyclotomic fields is isomorphic to $\mathbb{Z}_m^\times$ where the isomorphism

$$j \mapsto \sigma_j \colon \mathbb{Z}_m^\times \to \mathrm{Gal}(K/\mathbb{Q})$$

is defined by $\sigma_j(X) = X^j$ and $\mathbb{Q}$-linear extension. In general the degree of a Galois extension of fields is always equal to the order of the Galois group. The main theorem of Galois theory says that there is a one-to-one correspondence between the subgroups of G and the subfields of K. For example the subgroup $H < G$ corresponding to the subfield $L \subset K$ is the Galois group of K over L consisting of the automorphisms of K that fix the elements in L,

$$H = \mathrm{Gal}(K/L) = \{\sigma \in G \mid \sigma(x) = x \ \forall x \in L\}.$$

Conversely, L is the subfield of K consisting precisely of the elements that are fixed by all automorphisms in H and as such it is called the fixed field of H. Note that this implies the extension K/L is again Galois. Since G is abelian, also L is Galois over $\mathbb{Q}$ and the Galois group $\mathrm{Gal}(L/\mathbb{Q})$ is isomorphic to G/H. So the index of H in G, i.e. the order of G/H, is equal to the degree of L.

By restricting the automorphisms of K to the cyclotomic ring $\mathcal{R} \subset K$ we get ring automorphisms of $\mathcal{R}$. Moreover, since any ideal generated by a rational integer is stabilized under all automorphisms, the automorphisms factor through to automorphisms of the quotient ring $\mathcal{R}_q$.

Theorem 3.1. *Let $\mu \in \mathcal{R}_q$ be an element that is fixed modulo q by all Galois automorphisms $\sigma \in H$ of K fixing L; that is,*

$$\sigma(\mu) \equiv \mu \pmod{q\mathcal{R}} \text{ for all } \sigma \in H.$$

Then μ is contained in the subfield $\mathcal{S}_q$ of $\mathcal{R}_q$.

Proof. (Theorem 3.1) Since $\mathcal{R}$ is a Dedekind domain, the ideal $q\mathcal{R}$ of $\mathcal{R}$ can be (uniquely) written as a product of prime ideals, $q\mathcal{R} = \mathfrak{q}_1 \dots \mathfrak{q}_r$ [NS99, Theorem 3.1]. Let $\mathfrak{q}$ be one of these prime ideals, say $\mathfrak{q} = \mathfrak{q}_1$. The residue class field $\mathcal{R}/\mathfrak{q}$ is a finite extension of the finite field $\mathcal{S}_q$ since $\mathfrak{q}$ lies over the prime ideal $q\mathcal{S}$ of $\mathcal{S}$, $\mathfrak{q} \cap \mathcal{S} = q\mathcal{S}$. As such it is Galois over $\mathcal{S}_q$ with cyclic Galois group $\mathrm{Gal}((\mathcal{R}/\mathfrak{q})/\mathcal{S}_q)$ [Hun12, Proposition 5.10]. In contrast to $q\mathcal{R}$, the prime ideal $\mathfrak{q}$ is not stabilized by the whole Galois group H. So let $H_\mathfrak{q}$ be the subgroup of H that stabilizes $\mathfrak{q}$ so that $\sigma(\mathfrak{q}) = \mathfrak{q}$ for all $\sigma \in H_\mathfrak{q}$. This group is called the decomposition group of $\mathfrak{q}$ over $\mathcal{S}$ in Hilbert's ramification theory [NS99, Definition 9.2]. Then we have the canonical homomorphism $H_\mathfrak{q} \to \mathrm{Gal}((\mathcal{R}/\mathfrak{q})/\mathcal{S}_q)$, $\sigma \mapsto \tilde{\sigma}$ where $\tilde{\sigma}(x + \mathfrak{q}) = \sigma(x) + \mathfrak{q}$. It is an important fact in class field theory that this homomorphism is surjective [NS99, Proposition 9.4]. Hence, $\tilde{\sigma}(\mu + \mathfrak{q}) = \sigma(\mu) + \mathfrak{q} = \mu + \mathfrak{q}$ for all $\tilde{\sigma} \in \mathrm{Gal}((\mathcal{R}/\mathfrak{q})/\mathcal{S}_q)$ and it follows from Galois theory that $\mu + \mathfrak{q} \in \mathcal{S}_q$ and therefore $\mu \equiv x \pmod{\mathfrak{q}}$ for some $x \in S$. The Galois group H acts transitively on the prime ideals $\mathfrak{q}_i$ over $q\mathcal{S}$ [NS99, Proposition 9.1]. Therefore, for every $i = 1, \dots, r$, there is some $\sigma \in H$ such that $\sigma(\mathfrak{q}) = \mathfrak{q}_i$. It follows that $\mu \equiv \sigma(\mu) \equiv \sigma(x) \equiv x \pmod{\mathfrak{q}_i}$ and hence $\mu \equiv x \pmod{q\mathcal{R}}$ which shows that $\mu \in \mathcal{S}_q$.

We explain how we use Theorem 3.1. If we want to be able to prove knowledge of a message μ in a subfield $\mathcal{S}_q$ of size q^k for some k dividing d, we start from the Galois group G and select a subgroup H of order $\frac{d}{k}$. Then its fixed field L has degree $|G/H| = k$ and the quotient $\mathcal{S}_q$ has size q^k as we wanted. Unfortunately, primes q that are inert in L do not always exist. In fact it is necessary that G/H is cyclic. To understand this recall from the proof of Theorem 3.1 that if q is inert in L there is an isomorphism from G/H to the cyclic Galois group of the extension of finite fields $\mathcal{S}_q/\mathbb{Z}_q$ and so also G/H is cyclic. On the other hand, if G/H is cyclic, then it follows from the Chebotarev density theorem that infinitely many inert primes exist with density $\varphi(k)/k$ [NS99, Theorem 13.4]. Write $m = p_1^{\nu(p_1)} \dots p_r^{\nu(p_r)}$ for the prime decomposition of m. Then we have that $G \cong \mathbb{Z}_m^\times$ factors as $\mathbb{Z}^\times_{p_1^{\nu(p_1)}} \times \dots \times \mathbb{Z}^\times_{p_r^{\nu(p_r)}}$. All the direct factors for odd p_i are cyclic and $\mathbb{Z}^\times_{2^{\nu(2)}}$ is cyclic if $\nu(2) \leq 2$ and otherwise isomorphic to the product of the two cyclic groups $\mathbb{Z}_2$ and $\mathbb{Z}_{2^{\nu(2)-2}}$. We see that in order for the quotient G/H to be cyclic we can divide out all but one of the cyclic factors to get orders k of cyclic quotients that divide either $(p_i - 1)p_i^{\nu(p_i)-1}$ for an odd p_i or $p_i = 2$ with $\nu(2) \leq 2$, or $2^{\nu(2)-2}$ if $\nu(2) > 2$.

If we now compute a commitment to a message $\mu \in \mathcal{S}_q$, then, by Theorem 3.1, it is enough to prove that μ is fixed by the Galois automorphisms in H in order to establish that μ lies in $\mathcal{S}_q$ of order q^k. In fact, it clearly suffices to only prove that μ is fixed by a set of generators for H, which are usually only one or two.

For using Theorem 3.1 in practice we need to be able to compute elements in $\mathcal{S}_q$. For this we give a $\mathbb{Z}_q$-basis of the field $\mathcal{S}_q$. By the primitive element theorem there exists a single generator α of $L = \mathbb{Q}[\alpha]$ that lies in $\mathcal{S}$ [Hun12, Proposition 6.15]. Its powers form a $\mathbb{Q}$-basis for L but in general, since L is not necessarily cyclotomic, they do not form an integral basis so they are not a $\mathbb{Z}$-basis for $\mathcal{S}$. In fact, $\mathcal{O} = \mathbb{Z}[\alpha]$ is only a so-called order in L which can be strictly smaller than the ring of integers $\mathcal{S}$. Fortunately this does not pose a problem for us as we are only interested in $\mathcal{S}_q$. $\mathcal{O}_q = \mathcal{O}/q\mathcal{O}$ is a subring of $\mathcal{S}_q$ but since it also has q^k elements it must be equal to $\mathcal{S}_q$. This

shows there is an element of $\mathcal{O}$ in every coset of $\mathcal{S}$ modulo $q\mathcal{S}$ and we can use the powers of α as a $\mathbb{Z}_q$-basis for $\mathcal{S}_q$. More precisely, $1, \alpha, \alpha^2, \ldots, \alpha^{k-1}$ is a $\mathbb{Z}_q$-basis for $\mathcal{S}_q$ and we have

$$\mathcal{S}_q = \{c_0 + c_1\alpha + \cdots + c_{k-1}\alpha^{k-1} \in \mathcal{R}_q \mid c_i \in \mathbb{Z}_q\}.$$

As we are mainly interested in the power-of-two case we only give concrete generators α for all cyclic subfields of power-of-two cyclotomic fields in the next section but it is easy to compute generators for other examples.

What remains is to decide whether a given prime q is inert in L. A general approach is to compute the minimal polynomial of a generator α of L. Then the Dedekind-Kummer theorem says that q is inert in L if and only if the minimal polynomial is irreducible modulo q [Coh00, Theorem 4.8.13]. For the power-of-two case we give a much more direct answer to this question.

3.2 Power-of-Two Cyclotomic Rings

We consider the case where $K = \mathbb{Q}[X]/(X^d + 1)$ is a power-of-two cyclotomic field. The Galois group G, being isomorphic to $\mathbb{Z}_{2d}^\times$, has structure $\mathbb{Z}_2 \times \mathbb{Z}_{d/2}$ where the cyclic subgroups $\mathbb{Z}_2$ and $\mathbb{Z}_{d/2}$ are generated by σ_{-1} and σ_5, respectively [LS18, Lemma 2.4].

In the simplest case when we choose the subgroup H to be the full Galois group G, then the fixed field L is $\mathbb{Q}$ and $\mathcal{S}_q = \mathbb{Z}_q$ is a field for every prime number q. Theorem 3.1 gives that if some element $\mu \in \mathcal{R}_q$ is fixed by σ_{-1} and σ_5, then $\mu \in \mathbb{Z}_q$.

For subfields of degree $k|d$ with $k < d$, take as H the subgroup $\langle \sigma_{-1}, \sigma_5^k \rangle = \mathbb{Z}_k$.

Theorem 3.2. *Let $d > k \geq 1$ be powers of 2. The subgroup $H = \langle \sigma_{-1}, \sigma_5^k \rangle$ of the Galois group $G = \mathrm{Gal}(K/\mathbb{Q})$ has index k. Its fixed field L is generated by $\alpha = X^{d-\frac{d}{2k}} - X^{\frac{d}{2k}}$ over $\mathbb{Q}$ inside K, $L = \mathbb{Q}[\alpha] \subset K$.*

Proof. (Theorem 3.2) For $k = 1$ the subgroup H is equal to the whole Galois group G so its fixed field is $\mathbb{Q}$ which indeed is generated by $\alpha = 0$. For $k > 1$, H has order $2 \cdot \frac{d}{2k} = \frac{d}{k}$ and thus G/H has order k. Now observe $\sigma_{-1}(\alpha) = X^{\frac{d}{2k}-d} - X^{-\frac{d}{2k}} = -X^{\frac{d}{2k}} + X^{d-\frac{d}{2k}} = \alpha$ since $X^d = -1$. By repeated squaring one finds $5^k \equiv 1+4k \pmod{8k}$ and hence $\frac{d}{2k}(5^k-1) \equiv 0 \pmod{2d}$. Consequently $\sigma_5^k(\alpha) = X^{(d-\frac{d}{2k})5^k} - X^{\frac{d}{2k}5^k} = X^{d-\frac{d}{2k}} - X^{\frac{d}{k}} = \alpha$. So, α lies in L. Now consider the subfield $L' \subset K$ fixed by the Galois group $\langle \sigma_{-1}, \sigma_5^{k/2} \rangle$. It has degree $k/2$ and is contained in L since its Galois group contains H, which implies L is of degree 2 over L'. α is not fixed by $\sigma_5^{k/2}$ which means it does not lie in the subfield L' of degree $k/2$ and therefore α generates L.

For selecting primes it is helpful to compute the minimal polynomial of the generator α of L. The roots of the minimal polynomial are the conjugates of α under the action of G/H so it is given by

$$\phi(Y) = \prod_{\sigma \in G/H} (Y - \sigma(\alpha))$$

and the coefficients are the symmetric polynomials in the conjugates of α. We give an example for the case where the subfield has degree $k = 4$. A system of representatives for the Galois group G/H is given by $\{1, \sigma_5, \sigma_5^2, \sigma_5^3\}$. The conjugates of $\alpha = X^{d-\frac{d}{8}} - X^{\frac{d}{8}}$ are $\sigma_5(\alpha) = X^{d-5\frac{d}{8}} - X^{5\frac{d}{8}}$, $\sigma_5^2(\alpha) = X^{d-25\frac{d}{8}} - X^{25\frac{d}{8}} = X^{-\frac{d}{8}} - X^{d+\frac{d}{8}} = X^{\frac{d}{8}} - X^{d-\frac{d}{8}} = -\alpha$ and $\sigma_5^3(\alpha) = -\sigma_5(\alpha)$. Now it is easy to compute $\phi(Y) = Y^4 - 4Y^2 + 2$.

Theorem 3.3. *The prime numbers that are inert in the fixed field L of $\langle \sigma_{-1}, \sigma_5^k \rangle$ for some power of two $1 < k < d$ are precisely the primes that are congruent to 3 or 5 modulo 8. They split into two prime ideals in K.*

Proof. (Theorem 3.3) First consider the case $q \equiv 5 \pmod 8$. The Legendre symbol $(-1/q)$ is equal to 1 so there is a square root r of -1 modulo q. We get that $X^d + 1 \equiv (X^{d/2} - r)(X^{d/2} + r) \pmod q$. It follows from [LN86, Theorem 3.35] that the two factors of degree $d/2$ are irreducible. Hence q splits into two prime ideals in K. They are fixed by σ_5 and mapped to each other by σ_{-1}. Therefore q splits over the course of the subextension K/L where L is the fixed field of $\langle \sigma_{-1} \rangle$ and stays inert in L and in all subfields of L. Next we handle the case $q \equiv 3 \pmod 8$. In this case $(-2/q) = 1$ so that there exist an $r \in \mathbb{Z}_q$ with $r^2 \equiv -2 \pmod q$. This allows us to write $X^d + 1 \equiv (X^{d/2} + rX^{d/2} - 1)(X^{d/2} - rX^{d/4} - 1)$. These factors are again irreducible and q splits only into two prime ideals that are stabilized by the subgroup $\langle \sigma_{-1}\sigma_5 \rangle$ of order $d/2$, which therefore is the decomposition group of q. Since σ_{-1} is not contained in this group we see that q again splits over the course of the extension K/L and consequently stays inert in all subfields of L.

We end this section by giving and alternate subfield of degree 2 where we take $H = \mathbb{Z}_{d/2} = \langle \sigma_5 \rangle$. Then $\sigma_5(X^{d/2}) = X^{5d/2} = X^{d/2}$ and hence $X^{d/2}$ lies in the fixed field L of H. But since L is only of degree 2, it is already generated by $\alpha = X^{d/2}$. In fact, α has minimal polynomial $Y^2 + 1$ and L is the field of Gaussian integers inside K. An odd prime number q is inert in this imaginary quadratic field if and only if the Legendre symbol $\left(\frac{-1}{q}\right)$ is equal to -1 [NS99, Proposition 8.5], which is the case if and only if $q \equiv 3 \pmod 4$. This subfield of degree 2 has the advantage that only proving stability under one automorphism σ_5 is necessary. Also in this case there exist prime numbers q that split into more than two primes in K which allows for faster multiplication using the Fast Fourier Transform. We have

$$\mathcal{S}_q = \{c_0 + c_{d/2}X^{d/2} \in \mathcal{R}_q \mid c_0, c_{d/2} \in \mathbb{Z}_q\}.$$

Table 3 summarizes the subfields of degree $k \leq 8$ in K for $d = 4096$. We give an example how it can be used. Suppose we want to commit to messages in a subfield of size q^2 in $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^{4096} + 1)$ and give a proof that they are really contained in this subfield. Then we can compute messages $\mu = c_0 + c_1\alpha = c_0 + c_1(X^{3072} - X^{1024}) \in \mathcal{R}_q$ with arbitrary $c_0, c_1 \in \mathbb{Z}_q$. They are fixed by the two automorphisms σ_{-1} and σ_5^2 and if we prove that they indeed are, then it follows the messages are of this form. Moreover, these messages form a subfield of $\mathcal{R}_q$ and are therefore invertible.

4 Proofs of Knowledge

The goal of this section will be to construct a proof that a commitment opens to a message that is invariant under certain automorphisms. We will first introduce a more generic protocol that allows one to prove that the openings of a set of commitments verify a given linear relation. This protocol can be used even if the aforementioned commitments use different public keys and have different dimensions. We present this protocol as a standalone result as we think it can be of independent interest. We then show how, using the results of Section 3.1, we can adapt this generic protocol to obtain a proof of automorphism stability.

Degree	Galois group H	Generator α of L	Minimal polynomial of α
1	$\langle \sigma_{-1}, \sigma_5 \rangle$	1	$Y - 1$
2	$\langle \sigma_{-1}, \sigma_5^2 \rangle$	$X^{3072} - X^{1024}$	$Y^2 - 2$
2	$\langle \sigma_5 \rangle$	X^{2048}	$Y^2 + 1$
4	$\langle \sigma_{-1}, \sigma_5^4 \rangle$	$X^{3584} - X^{512}$	$Y^4 - 4Y^2 + 2$
8	$\langle \sigma_{-1}, \sigma_5^8 \rangle$	$X^{3849} - X^{256}$	$Y^8 - 8Y^6 + 20Y^4 - 16Y^2 + 2$

Table 3. Subfields of the power-of-two cyclotomic field $K = \mathbb{Q}[X]/(X^{4096} + 1)$ of degree at most 8 with generators for the corresponding Galois group H, generators of the subfields, and their minimal polynomials over $\mathbb{Q}$.

4.1 Generic Proof for Linear Relations

In this section we will present a novel proof of knowledge that a set of commitments $\mathbf{t}_1, \ldots, \mathbf{t}_\tau$ are such that their openings $\mathbf{m}_1, \ldots, \mathbf{m}_\tau$ verify

$$\sum_1^\tau \mathbf{B}_j \mathbf{m}_j = \mathbf{0}$$

for any fixed $\mathbf{B}_1, \ldots, \mathbf{B}_\tau$. An interesting property of this protocol will be that the commitment matrices $\mathbf{A}_1, \ldots, \mathbf{A}_\tau$ do not have to be identical, in fact they can even have different dimensions as long as all the matrices $\mathbf{B}_1, \ldots, \mathbf{B}_\tau$ have the same number of rows.

Concretely, for $j \in [\tau]$ let $\mathbf{A}_j := \begin{bmatrix} \mathbf{A}_{j,1} \\ \mathbf{A}_{j,2} \end{bmatrix}$ with $\mathbf{A}_{j,1} \in \mathcal{R}_{q_1}^{n_j \times k_j}$ and $\mathbf{A}_{j,2} \in R_{q_2}^{l_j \times k_j}$, let $\mathbf{B}_j \in \mathcal{R}_{q_2}^{x \times l_j}$, and let

$$\mathbf{t}_j := \begin{bmatrix} \mathbf{t}_{j,1} \\ \mathbf{t}_{j,2} \end{bmatrix} = \mathbf{A}_j \mathbf{r}_j + \begin{bmatrix} \mathbf{0} \\ \mathbf{m}_j \end{bmatrix}, \text{ with } \mathbf{r}_j \in \mathcal{R}^{k_j},\ \mathbf{m}_j \in \mathcal{R}_{q_2}^{l_j}$$

be such that

$$\sum_1^\tau \mathbf{B}_j \mathbf{m}_j = \mathbf{0}. \tag{13}$$

We prove in Lemma 4.1 that the protocol of Figure 1, in which the challenge space is

$$\mathcal{C} = \{c \in \mathcal{R} :\ \|c\|_1 = \kappa,\ \|c\|_\infty = 1\},$$

is a proof of knowledge that the commitments $\mathbf{t}_1, \ldots, \mathbf{t}_\tau$ are well formed and that their openings verify (13)

Lemma 4.1. *Let $\mathbf{r}_j \xleftarrow{\$} S_1^{k_j}$. Let $\xi \geq 11\kappa\sqrt{d \sum k_j}$ and $B'_j \geq \sqrt{2dk_j}\xi$. Also, let $\bar{C}$ (defined as in Table 2) be such that all elements in it are invertible over $\mathcal{R}_{q_2}$. Then the protocol Π_{lin} of Figure 1 achieves the following properties:*

- Correctness: *The prover aborts with probability at most $\frac{2}{3} + 2^{-100}$, and if he does not abort the verifier accepts with overwhelming probability.*
- Honest-Verifier Zero-Knowledge: *Non-aborting transcripts with an honest verifier can be simulated with statistically indistinguishable distribution.*

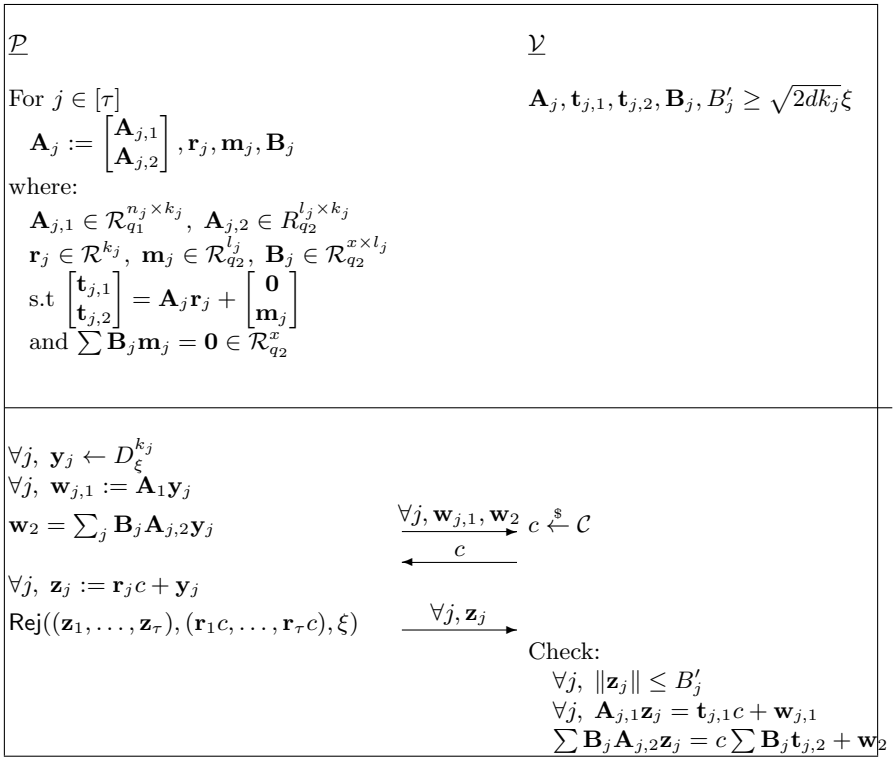

Fig. 1. Proof that a set of commitments $\mathbf{t}_1, \ldots, \mathbf{t}_\tau$ open to messages $\mathbf{m}_1, \ldots, \mathbf{m}_\tau$ such that $\sum \mathbf{B}_j\mathbf{m}_j = 0 \mod q_2$.

– Special Soundness: *Given two accepting transcripts one can extract valid openings* $\bar{\mathbf{z}}_j$, $\bar{\mathbf{m}}_j$, $\bar{c}$ *of* $\mathbf{t}_j$ *for* $j \in [\tau]$ *such that* $\bar{c} \in \bar{C}$, $\|\bar{\mathbf{z}}_j\| \leq 2B'_j$, *and* $\sum \mathbf{B}_j\bar{\mathbf{m}}_j = \mathbf{0}$.

Proof.

– *Correctness:* If $\mathcal{P}$ and $\mathcal{V}$ are honest then $\|(\mathbf{r}_1 c, \ldots, \mathbf{r}_\tau c)\| \leq \kappa\sqrt{d\sum k_j}$, the probability of abort is exponentially close to 2/3 by definition of rejection sampling and lemma 2.2. Since each coefficient of $\mathbf{z}_j$ is statistically close to D_ξ, then according to (12) we have $\|\mathbf{z}_j\| \leq \sqrt{2k_j d}\xi$ with overwhelming probability.
– *Honest-Verifier Zero-Knowledge:* We only show that the protocol is zero-knowledge when the prover does not abort prior to sending $\mathbf{z}_j$. The reason that this is enough for practical purposes is that HVZK Σ-protocols are first converted into non-interactive proofs via the Fiat-Shamir transform, in which case the verifier never sees the aborting transcripts. One can also, using the standard technique of sending commitments of $\mathbf{w}_{j,1}, \mathbf{w}'_2$ and only opening them in case a non-abort occurs, make the interactive protocol zero-knowledge.
Let $\mathcal{S}(\mathbf{A}, \mathbf{t}_1, \ldots, \mathbf{t}_\tau, \mathbf{B}_1, \ldots, \mathbf{B}_j)$ be the following PPT algorithm:
  - Sample $\mathbf{c} \leftarrow \mathcal{C}$
  - Sample $\mathbf{z}_j \leftarrow D_\xi^{k_j}$
  - Set $\mathbf{w}_{j,1} = \mathbf{A}_{j,1}\mathbf{z}_j - \mathbf{t}_{j,1}c$
  - Set $\mathbf{w}_2 = \sum \mathbf{B}_j\mathbf{A}_{j,2}\mathbf{z}_j - c\sum \mathbf{B}_j\mathbf{t}_{j,2}$
  - Output $(\mathbf{w}_{1,1}, \ldots, \mathbf{w}_{\tau,1}, \mathbf{w}_2, c, \mathbf{z}_1, \ldots, \mathbf{z}_\tau)$

It is clear that $\mathbf{z}_1, \ldots, \mathbf{z}_\tau$ verifies with overwhelming probability. We know that in the real protocol when no abort occurs the distribution of $\mathbf{z}_j$ is within statistical distance 2^{-100} of $D_\xi^{k_j}$. As $\mathbf{w}_{1,1}, \ldots, \mathbf{w}_{\tau,1}, \mathbf{w}_2$ are completely determined by $\mathbf{A}_1, \ldots, \mathbf{A}_\tau, \mathbf{B}_1, \mathbf{B}_\tau, \mathbf{t}_1, \ldots, \mathbf{t}_\tau, \mathbf{z}_1, \ldots, , \ldots, \mathbf{z}_\tau, c$, the distribution of $(\mathbf{w}_{1,1}, \ldots, \mathbf{w}_{\tau,1}, \mathbf{w}_2, c, \mathbf{z}_1, \ldots, \mathbf{z}_\tau)$ output by $\mathcal{S}$ is within 2^{-100} of the distribution of these variables in the actual protocol.

– *Special Soundness:* Let $(\mathbf{w}_{1,1}, \ldots, \mathbf{w}_{\tau,1}, \mathbf{w}_2, c, \mathbf{z}_1, \ldots, \mathbf{z}_\tau)$ and $(\mathbf{w}_{1,1}, \ldots, \mathbf{w}_{\tau,1}, \mathbf{w}_2, c', \mathbf{z}'_1, \ldots, \mathbf{z}'_\tau)$ be two accepting transcripts with $c \neq c'$. We will prove that there exists messages $\bar{\mathbf{m}}_j,\ j \in [\tau]$ such that $(\mathbf{z}_j - \mathbf{z}'_j, \bar{\mathbf{m}}_j, c - c')$ is a valid opening of $\mathbf{t}_j$, and $\sum \mathbf{B}_j \bar{\mathbf{m}}_j = \mathbf{0}$. Let $\bar{\mathbf{z}}_j = \mathbf{z}_j - \mathbf{z}'_j$, and $\bar{c} = c - c'$. By computing the difference of the verification equations for both transcripts we obtain:

$$\mathbf{A}_{j,1}\bar{\mathbf{z}}_j = \mathbf{t}_1 \bar{c} \tag{14}$$

$$\sum \mathbf{B}_j \mathbf{A}_{j,2} \bar{\mathbf{z}}_j = \bar{c} \sum \mathbf{B}_j \mathbf{t}_{j,2} \tag{15}$$

Since $\bar{c}$ has an inverse in $\mathcal{R}_{q_2}$ we can define $\bar{\mathbf{m}}_j \in \mathcal{R}_{q_2}^{l_j}$ such that $\bar{c}\mathbf{t}_{j,2} = \mathbf{A}_{j,2}\bar{\mathbf{z}}_j + \bar{c}\bar{\mathbf{m}}_j$. By replacing $\bar{c}\mathbf{t}_{j,2}$ in equation 15 we have:

$$\sum \mathbf{B}_j \bar{\mathbf{m}}_j = \mathbf{0}$$

In conclusion we have extracted $\bar{\mathbf{z}}_1, \ldots, \bar{\mathbf{z}}_\tau$, $\bar{\mathbf{m}}_1, \ldots, \bar{\mathbf{m}}_\tau$, and $\bar{c}$ such that:

$$\mathbf{A}_j \bar{\mathbf{z}}_j + \bar{c} \begin{bmatrix} 0 \\ \bar{\mathbf{m}}_j \end{bmatrix} = \bar{c}\mathbf{t}_j$$

with $\|\bar{\mathbf{z}}_j\| \leq 2B'_j$, and $\sum \mathbf{B}_j \bar{\mathbf{m}}_j = \mathbf{0}$.

4.2 Proof of automorphism stability

We present in this section a proof of knowledge that a commitment opens to a message $\mathbf{m} \in \mathcal{R}_{q_2}^l$ that is invariant under a certain set of automorphisms $(\sigma_j)_{j \in S \subset \mathbb{Z}_m^*}$, where the ring we consider is $\mathcal{R} = \mathbb{Z}[X]/\Phi_m$. As shown in section 3.1 as a special case we can show that $\mathbf{m} \in \mathbb{Z}_{q_2}^l$ by proving that it is invariant under a well chosen automorphism when m is prime or $m = 2^b p^j$ for $b \in \{0, 1\}, j > 0$ and p an odd prime, or by proving that it is invariant under two automorphisms (specifically σ_{-1} and σ_5) when m is a power of two. We can also prove that $\mathbf{m}$ belongs to certain sets of size q^{li} for specific integers i by proving it is invariant under a well-chosen set of automorphism, section 3.2 shows for which i this is possible and which set $S \subset \mathbb{Z}_m^*$ should be used. We now show how to use the previous proof to prove that for a set of automorphisms $\sigma_j, j \in S$ a commitment $\mathbf{t}$ opens to a message $\mathbf{m}$ such that $\forall j \in S, \sigma_j(\mathbf{m}) = \mathbf{m}$. For ease of presentation we will rewrite S as the set $\{1, \ldots, |S|\}$ (while the previous section used S as the set of the powers corresponding to the Galois automorphisms).
We consider a commitment

$$\mathbf{t} := \mathbf{A}\mathbf{r} + \begin{bmatrix} \mathbf{0} \\ \mathbf{m} \end{bmatrix}$$

with $\mathbf{t} \in \mathcal{R}_{q_1}^n \times \mathcal{R}_{q_2}^l$, $\mathbf{A} \in \mathcal{R}_{q_1}^{n \times k} \times \mathcal{R}_{q_2}^{l \times k}$, and $\mathbf{m} \in \mathcal{R}_{q_2}^l$. We will use the proof of Figure 1 with the following parameters:

- $\tau := |S| + 1$
- $x := |S|$
- $\mathbf{A}_1 = \mathbf{A}$
- For $j \in S$, $\mathbf{A}_{j+1} := \sigma_j^{-1}(\mathbf{A})$
- $\mathbf{t}_1 := \mathbf{t}$
- For $j \in S$, $\mathbf{t}_{j+1} := \sigma_j^{-1}(\mathbf{t})$
- $\mathbf{B}_1 := \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \in \mathcal{R}_{q_2}^{|S| \times l}$
- For $j \in S$, $\mathbf{B}_{j+1} := \begin{bmatrix} \mathbf{0}^{j-1} \\ 1 \\ \mathbf{0}^{|S|-j} \end{bmatrix} \in \mathcal{R}_{q_2}^{|S| \times l}$
- For $j \in S$, $B'_j = B_{aut}$

Corollary 4.2. *Let $\mathbf{r} \xleftarrow{\$} S_1^k$. Let S be a set of automorphisms of size $|S|$. Let $\mathbf{t} = Com(\mathbf{m}; \mathbf{r})$ with $\sigma_j(\mathbf{m}) = \mathbf{m}$ for all j in S. Let $\xi \geq 11\kappa\sqrt{k\,|S|\,d}$ and $B_{aut} \geq \sqrt{2dk}\xi$. If $B_{com} \geq 2B_{aut}$, then the protocol of Figure 1 instantiated with the parameters set as above achieves the following properties:*

- Correctness: *The prover aborts with probability at most $2/3 + 2^{-100}$, and if he does not abort the verifier accepts with overwhelming probability.*
- Honest-Verifier Zero-Knowledge: *Non-aborting transcripts with an honest verifier can be simulated with statistically indistinguishable distribution.*
- Special Soundness: *Given two accepting transcripts one can extract a valid opening $\bar{\mathbf{z}}$, $\bar{\mathbf{m}}$, $\bar{c}$ of $\mathbf{t}$ such that $\|\bar{\mathbf{z}}\| \leq 2B_{aut}$, and $\forall j \in S, \bar{\mathbf{m}} = \sigma_j(\bar{\mathbf{m}})$.*

Proof. Correctness and zero-knowledge follow directly from Lemma 4.1.
By the special soundness of Lemma 4.1 we have that given two accepting transcripts one can extract valid openings $\bar{\mathbf{z}}_j$, $\bar{\mathbf{m}}_j$, $\bar{c}$ of $\mathbf{t}_j$ for $j \in [\tau]$ such that $\|\bar{\mathbf{z}}_j\| \leq 2B'_j$, and $\sum \mathbf{B}_j \bar{\mathbf{m}}_j = \mathbf{0}$. Using how the matrices $\mathbf{A}_j, j \in [\tau]$ are defined we obtain

$$\bar{c}\mathbf{t} = \mathbf{A}\bar{\mathbf{z}}_1 + \bar{c} \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}}_1 \end{bmatrix}$$
$$\bar{c}\sigma_j^{-1}(\mathbf{t}) = \sigma_j^{-1}(\mathbf{A})\bar{\mathbf{z}}_{j+1} + \bar{c} \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}}_{j+1} \end{bmatrix}, \forall j \in S$$

By applying σ_j to the second equation we can rewrite it as

$$\sigma_j(\bar{c})\mathbf{t} = \mathbf{A}\sigma_j(\bar{\mathbf{z}}_{j+1}) + \sigma_j(\bar{c}) \begin{bmatrix} \mathbf{0} \\ \sigma_j(\bar{\mathbf{m}}_{j+1}) \end{bmatrix}, \forall j \in S$$

Since $\|\sigma_j(\bar{\mathbf{z}}_{j+1})\| \leq 2B_{aut} \leq B_{Com}$, we have by the binding property of the commitment scheme that all these openings open to the same message, i.e.

$$\bar{\mathbf{m}}_1 = \sigma_j(\bar{\mathbf{m}}_{j+1}), \forall j \in S.$$

	Full signature	Commitment	Ciphertext	Proof	Secret key
Parameters I	581 KB	113 KB	123 KB	345 KB	146 KB
Parameters II	1173 KB	204 KB	254 KB	715 KB	292 KB

Table 4. Size of a signature and user secret key

Furthermore we know that $\sum \mathbf{B}_j \bar{\mathbf{m}}_j = \mathbf{0}$ which by construction of the $\mathbf{B}_j$ corresponds to

$$\bar{\mathbf{m}}_1 = \bar{\mathbf{m}}_{j+1}, \forall j \in S$$

from these two equality we have

$$\bar{\mathbf{m}}_1 = \sigma_j(\bar{\mathbf{m}}_1), \forall j \in S.$$

In conclusion we have extracted $\bar{\mathbf{z}}_1$, $\bar{\mathbf{m}}_1$, and $\bar{c}$ such that:

$$\bar{c}\mathbf{t} = \mathbf{A}\bar{\mathbf{z}}_1 + \bar{c} \begin{bmatrix} 0 \\ \bar{\mathbf{m}}_1 \end{bmatrix}$$

with $\|\bar{\mathbf{z}}_1\| \leq 2B_{aut}$ and $\bar{\mathbf{m}}_1 = \sigma_j(\bar{\mathbf{m}}_1), \forall j \in S$.

5 Group Signatures

We first recall the definitions and security model of group signatures. A group signature scheme consists of a tuple of four algorithms $(\mathbf{GSetup}, \mathbf{Sign}, \mathbf{Verify}, \mathbf{Open})$:

- $\mathbf{GSetup}(1^\lambda, 1^N)$:Takes as input the security parameter λ as well as the maximum number of identities N. Outputs the group public key gpk, the group manager secret key $gmsk$, and the secret keys of each identity $sk_1, \ldots, sk_N$.
- $\mathbf{Sign}(sk_i, M)$: Takes as input a user secret key sk_i and a message $M \in \{0,1\}^*$. Outputs a signature z of M.
- $\mathbf{Verify}(gpk, M, z)$: Takes as input the group public key gpk, a message M, and a signature z. Outputs 1 if z is a valid signature of M and 0 otherwise.
- $\mathbf{Open}(gmsk, M, z)$: Takes as input the group manager secret key $gmsk$, a message M, and a valid signature z of M. Outputs an identity $id \in [N]$ or $\perp$.

For correctness, we want that for any $(gpk, gmsk, sk_1, \ldots, sk_n) \leftarrow \mathbf{GSetup}(1^\lambda)$, any $j \in [N]$, $M \in \{0,1\}^*$, and $z \leftarrow \mathbf{Sign}(gpk, sk_j, M)$, with overwhelming probability:

$$\mathbf{Verify}(gpk, M, z) = 1, \text{ and } \mathbf{Open}(gpk, gmsk, M, z) = j$$

The security of the group signature is captured by two notions: anonymity and traceability, we describe these notions informally, for a more formal approach see [BMW03]. For anonymity we consider a PPT adversary $\mathcal{A}$ who has access to all the signing keys $sk_1, \ldots, sk_N$ but not the manager secret key $gmsk$. $\mathcal{A}$ chooses a message M and two identities i_0 and i_1, his goal is to distinguish between signatures of M under these identities. There are multiple flavors of anonymity depending on whether $\mathcal{A}$ can access an opening oracle (full anonymity) or not (weak anonymity), intuitively full anonymity will be achieved when the PKE used in the opening is CCA-secure while

Parameter	Notation	Params I	Params II
Ring dimension	d	4096	8192
Commitment modulus ("Top")	q_1	$\sim 2^{30}$	$\sim 2^{20}$
Commitment modulus ("Bottom")	q_2	$\sim 2^{80}$	$\sim 2^{80}$
Commitment row dimension ("Top")	n	1	1
Commitment message dimension	l	1	1
Verifiable encryption plaintext module	p	$\sim 2^{27}$	$\sim 2^{27}$
Verifiable encryption ciphertext module	Q	$\sim 2^{60}$	$\sim 2^{62}$
Bound on the challenge space	κ	26	24
Standard deviation of the GPV trapdoor	s	$\sim 2^{49}$	$\sim 2^{50}$
Standard deviation of the NTRU trapdoor	r	$\sim 2^{42}$	$\sim 2^{42}$
Root-Hermite Factor	δ_0	1.0036	1.002
Security in space for sieving (bits)		93	207
Post-Quantum Security in time for enumeration (bits)		242	1084

Table 5. Concrete parameters for our Group signature

weak anonymity corresponds to a CPA-secure encryption scheme. In this paper, for simplicity, we present a weakly anonymous group signature but mention that the verifiable encryption scheme we use [LN17] can achieve CCA security, which would end up somewhat increasing the signature size (by $\approx 20\%$ in our case). For the full-traceability notion, the adversary $\mathcal{A}$ has access to the signing keys $(sk_i)_{i\in S}$ for any arbitrary set $S \subset [N]$ (possible $S = [N]$) as well as the manager secret key $gmsk$, his goal is to produce a valid signature z of some message M (i.e. which passes verification) such that either $\text{Open}(gpk, gmsk, M, z) = j \notin S$ or $\text{Open}(gpk, gmsk, M, z) = \perp$. Full-traceability captures the notion that all signatures, even when computed by a collusion of users and the group manager, should trace to a member of the forging coalition. Note that full-traceability implies unforgeability, since the forgery game is a special case of the full-traceability game with $S = \emptyset$. Here the condition $\text{Open}(gpk, gmsk, M, z) = j \notin S = \emptyset$ or $\text{Open}(gpk, gmsk, M, z) = \perp$ is vacuous.

5.1 The Scheme

The group signature we present in this section will be for fixed parameters as per Table 5, for which the signatures will be of size 581 KB, as described in Section 8. In particular we consider the power-of-two cyclotomic ring $\mathcal{R} = \mathbb{Z}[X]/X^{4096} + 1$, and identity set $[N] = \mathbb{Z}_{q_2}$. This entails that user identities are exactly the elements $x \in \mathcal{R}_{q_2}$ that are left invariant under the automorphisms $\sigma_{-1} : X \to X^{-1} = -X^{d-1}$ and $\sigma_5 : X \to X^5$. We also use commitments that rely on R-LWE

and R-SIS (which can be seen as specific instances of the module variant of the corresponding problems for modules of dimension 1). Using other cyclotomic rings can result in smaller signatures (especially since for some of them only one automorphism is needed to prove that elements belong to $\mathbb{Z}_{q_2}$), and using higher dimension commitments that rely on the module variants of LWE and SIS would allow for more fine-tuned parameters. We have chosen the parameters in this section as such for easier presentation and because they allow for simpler implementations.
We will first present in this section a group signature scheme without opening, and show in section 5.2 how to add an opening.
Let $\delta = \lceil \sqrt{q_2} \rceil$, and $\mathbf{g}^T$ be the gadget matrix $\begin{bmatrix} 1 & \delta \end{bmatrix} \in \mathcal{R}_{q_2}^{1\times 2}$, we will consider the set of identities $\mathcal{I}d = \mathbb{Z}_{q_2}$.

GSetup(1^λ):

- Sample $\mathbf{A} := \begin{bmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \end{bmatrix} \leftarrow$ **CSetup**(1^λ), with $\mathbf{a}_1 \in \mathcal{R}_{q_1}^3$, and $\mathbf{a}_2 \in \mathcal{R}_{q_2}^3$.
- Sample $\mathbf{a} \xleftarrow{\$} \mathcal{R}_{q_2}^2$.
- Sample $\mathbf{R} \xleftarrow{\$} S_1^{2\times 2}$ and set $\mathbf{b}^T = \mathbf{a}^T\mathbf{R} \in \mathcal{R}_{q_2}^{1\times 2}$.
- Sample $(\mathbf{s}_{0_1}, \mathbf{s}_{0_2}, \mathbf{s}_{0_3}) \leftarrow D_s^2 \times D_s^2 \times D_r^3$
- Set $u := \begin{bmatrix} \mathbf{a}^T \mid \mathbf{b}^T \mid \mathbf{a}_2^T \end{bmatrix} \begin{bmatrix} \mathbf{s}_{0_1} \\ \mathbf{s}_{0_2} \\ \mathbf{s}_{0_3} \end{bmatrix}$
- Set $gpk := (\mathbf{A}, \mathbf{a}, \mathbf{b}, u)$
- For $i \in \mathbb{Z}_{q_2}^*$, sample $\mathbf{s}_{i_3} \leftarrow D_r^3$
- For $i \in \mathbb{Z}_{q_2}^*$, sample $(\mathbf{s}_{i_1}, \mathbf{s}_{i_2}) \in \mathcal{R}^4$ s.t
  $\begin{bmatrix} \mathbf{a}^T \mid \mathbf{b}^T + i\mathbf{g}^T \end{bmatrix} \begin{bmatrix} \mathbf{s}_{i_1} \\ \mathbf{s}_{i_2} \end{bmatrix} = u - \mathbf{a}_2^T\mathbf{s}_{i_3}$, and $(\mathbf{s}_{i_1}, \mathbf{s}_{i_2}) \leftarrow D_s^4$
- For $i \in \mathbb{Z}_{q_2}$, set $sk_i := \mathbf{s}_i := (\mathbf{s}_{i_1}, \mathbf{s}_{i_2}, \mathbf{s}_{i_3})$

Intuitively user i could sign a message $M \in \{0,1\}^*$ by doing a non-interactive proof that he knows a small $\mathbf{s} \in \mathcal{R}^7$ such that $\begin{bmatrix} \mathbf{a}^T \mid \mathbf{b}^T + i\mathbf{g}^T \mid \mathbf{a}_2^T \end{bmatrix} \mathbf{s} = u$ in which the message is part of the hash that generates the challenge. However doing so would reveal his identity – a verifier would need to know the matrix $\begin{bmatrix} \mathbf{a}^T \mid \mathbf{b}^T + i\mathbf{g}^T \mid \mathbf{a}_2^T \end{bmatrix}$ to verify the signature and since $\mathbf{a}$, $\mathbf{a}_2$, $\mathbf{b}$, and $\mathbf{g}$ are public, he could recover the identity of the signer. As explained in section 1.2, we circumvent this issue by committing to the part of the matrix that depends on i (that is $i\mathbf{g}^T$) and proving knowledge of a solution to a related equation.

Sign$(M, \mathbf{s}_i)$:

- Set $\mathbf{t} := \begin{bmatrix} t_1 \\ t_2 \end{bmatrix} = Com(i, \mathbf{r}) \in \mathcal{R}_{q_1} \times \mathcal{R}_{q_2}$, where $\mathbf{r} \leftarrow S_1^3$.
- Set $\mathbf{t}' := \begin{bmatrix} t_1' \\ t_2' \end{bmatrix} = Com(i\delta, \mathbf{r}')$, where $\mathbf{r}' \leftarrow S_1^3$.
- Set $\mathbf{v}^T := \begin{bmatrix} \mathbf{a}^T \mid \mathbf{b}^T + \begin{bmatrix} t_2 & t_2' \end{bmatrix} \mid \mathbf{a}_2^T \end{bmatrix} \in \mathcal{R}_{q_2}^{1\times 7}$, and $\mathbf{s}' = \begin{bmatrix} \mathbf{s}_{i_1} \\ \mathbf{s}_{i_2} \\ \mathbf{s}_{i_3} - \begin{bmatrix} \mathbf{r} & \mathbf{r}' \end{bmatrix} \mathbf{s}_{i_2} \end{bmatrix} \in \mathcal{R}^7$, observe that $\mathbf{v}^T\mathbf{s}' = u$
- In parallel (see below for explanation):

- Compute a proof Π_1 that $\mathbf{t}, \mathbf{t}'$ open to messages m, m' such that $m' = \delta m$
- Compute a proof Π_2 that $\mathbf{t}$ opens to a message m such that $m = \sigma_{-1}(m) = \sigma_{-5}(m)$
- Compute a proof Π_3 of knowledge of $\mathbf{s}'$ such that $\mathbf{v}^T\bar{\mathbf{s}} = u$

– Output the signature $z = (\mathbf{t}, \mathbf{t}', \Pi_1, \Pi_2, \Pi_3)$

The proofs Π_1, Π_2, Π_3 will use the Fiat-Shamir heuristic to transform the interactive proofs of Section 4 into non-interactive proofs in the random oracle model, we will also include the message M in the random oracle call to obtain a signature. For extraction, we will need all of these proofs to be executed with the same challenge (or output of the hash function in the non-interactive version). This is done in the standard way with the signer running all three proofs in parallel and, in the non-interactive version, computing a common challenge for all three as a hash of all the relevant information. We describe the full non-interactive proof, including the opening, in more details in Section 7.
To verify a signature one simply verifies the proofs Π_1, Π_2, Π_3.

Verify$(\mathbf{t}, \mathbf{t}', \Pi_1, \Pi_2, \Pi_3)$:

– Let $\begin{bmatrix} t_1 \\ t_2 \end{bmatrix} := \mathbf{t}$
– Let $\begin{bmatrix} t'_1 \\ t'_2 \end{bmatrix} := \mathbf{t}'$
– Let $\mathbf{v}^T = \left[\mathbf{a}^T \mid \mathbf{b}^T + \begin{bmatrix} t_2 & t'_2 \end{bmatrix} \mid \mathbf{a}_2^T\right]$
– Verify Π_1 using $\mathbf{t}, \mathbf{t}', \delta$
– Verify Π_2 using $\mathbf{t}, \sigma_{-1}, \sigma_5$
– Verify Π_3 using $\mathbf{v}$

5.2 Adding the Opening

To be able to open the group signature scheme of Section 5.1 we will add a verifiable encryption to the signature. In essence we want the signer to encrypt his identity, using a public key associated to a decryption key that the group manager possesses, and prove that this encryption is indeed of his identity. To do so we will encrypt the randomness $\mathbf{r}$ of $\mathbf{t} = Com(id; \mathbf{r})$ and prove that $\mathbf{a}_1^T\mathbf{r} = t_1$, note that encrypting id directly would result in a smaller ciphertext but a very large proof since id itself is not small. We use the verifiable encryption of [LN17] which consists in a *R-LWE* encryption and a proof of knowledge. We let p be the modulus of the plaintext space of our encryption scheme (which we only need large enough to accommodate the decryption slack, see [LN17]) and Q the modulus of the ciphertext.

PKESetup(1^λ):

– Sample $a \xleftarrow{\$} \mathcal{R}_Q$
– Sample $\mathbf{s}, \mathbf{e} \leftarrow S_1^3$
– Set $\mathbf{b} := a\mathbf{s} + \mathbf{e} \in \mathcal{R}_Q^3$
– Output $(\mathbf{s}, (a, \mathbf{b}))$

Encryption will consist in creating a standard *R-LWE* encryption and a proof that the message $\mathbf{r}$ encrypted is the randomness in $\mathbf{t} = Com(id; \mathbf{r})$.

Enc$((a, \mathbf{b}), \mathbf{r}, t_1)$:

– Sample $r, e_1 \leftarrow S_1$
– Sample $\mathbf{e}_2 \leftarrow S_1^3$
– Set $u := p(ar + e_1)$
– Set $\mathbf{v} := p(\mathbf{b}r + \mathbf{e}_2) + \mathbf{m}$
– Set $\mathbf{B}_1 := \begin{bmatrix} pa & p & 0 & 0 & 0 & 0 & 0 & 0 \\ pb_1 & 0 & p & 0 & 0 & 1 & 0 & 0 \\ pb_2 & 0 & 0 & p & 0 & 0 & 1 & 0 \\ pb_3 & 0 & 0 & 0 & p & 0 & 0 & 1 \end{bmatrix} \in \mathcal{R}_Q^{4\times 8}$
– Set $\mathbf{B}_2 := \begin{bmatrix} \mathbf{0}^{1\times 5} & \mathbf{a}_1^T \end{bmatrix} \in \mathcal{R}_{q_1}^{1\times 8}$
– Set $\mathbf{B} := \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \end{bmatrix}$
– Set $\mathbf{x} := \begin{bmatrix} r \\ e_1 \\ \mathbf{e}_2 \\ \mathbf{r} \end{bmatrix} \in \mathcal{R}^8$
– Set $\mathbf{y} := \begin{bmatrix} u \\ \mathbf{v} \\ t_1 \end{bmatrix} \in \mathcal{R}_Q^4 \times \mathcal{R}_{q_1}$
– Compute a proof Π of knowledge of $\mathbf{x}$ such that $\mathbf{B}\mathbf{x} = \mathbf{y}$.
– Output $(u, \mathbf{v}, \Pi)$

To verify an encryption one simply verifies the proof Π.

Verify$((u, \mathbf{v}, \Pi), t_1)$:

– Set $\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}, \mathbf{y}$ as in **Enc** above.
– Output **Verify**$(\Pi, \mathbf{B}, \mathbf{y})$

Decryption is not as simple as standard R-LWE decryption. By completeness we know that honestly generated ciphertexts can be decrypted but soundness should guarantee that as long as the proof verifies one should be able to decrypt. This is not clear since the proof Π does not imply that $(u, \mathbf{v})$ is a valid ciphertext but that there exists some $\bar{c} \in \bar{\mathcal{C}}$ such that $(\bar{c}u, \bar{c}\mathbf{v})$ is a valid ciphertext and we do not know which one. In [LN17] the authors show that in fact trying random $\bar{c}$ is a valid approach and the expected number of attempts is the same as the expected number of oracle calls that are needed to generate the proof (in particular only one attempt is necessary if the prover is honest). This will be sufficient for our scheme.

Dec$((u, \mathbf{v}, \Pi), \mathbf{s})$:

– If **Verify**$((u, \mathbf{v}, \Pi), t_1) = 1$, Let c be the challenge used in Π
– Loop:
– $\quad c' \leftarrow \mathcal{C}$
– $\quad \bar{c} := c - c'$
– $\quad \bar{\mathbf{r}} := (\mathbf{v} - u\mathbf{s})\bar{c} \mod Q$
– $\quad$ If $\|\bar{\mathbf{r}}\|_\infty \leq Q/8\kappa$ then:
– $\quad\quad \bar{\mathbf{r}} := \bar{\mathbf{r}} \mod q$
– $\quad\quad$ return $(\bar{\mathbf{r}}, \bar{c})$

The following lemma shows that if decryption succeeds then the decrypted value $(\bar{\mathbf{r}}, \bar{c})$ will essentially be a preimage for the zero-knowledge proof.

Lemma 5.1 ([LN17] Lemma 3.1). *Let $sk = \mathbf{s}$, and $\mathbf{e}$ be the error in $\mathbf{b} = a\mathbf{s} + \mathbf{e}$. If for given*

$$(u, \mathbf{v}, t_1) \in \mathcal{R}_Q^4 \times \mathcal{R}_{q_1}$$

there exists $\bar{\mathbf{r}}_B := (\bar{r}, \bar{e}_1, \bar{\mathbf{e}}_2, \bar{\mathbf{r}}) \in \mathcal{R}^8$, and $\bar{c} \in \mathcal{R}$ such that :

$$\mathbf{Br}_B = \begin{bmatrix} u & \bmod Q \\ \mathbf{v} & \bmod Q \\ t_1 & \bmod q_1 \end{bmatrix}$$

and

$$\|p(\bar{u}\mathbf{e} + \bar{\mathbf{e}}_2 - \bar{e}\mathbf{s}) + \bar{\mathbf{r}}\|_\infty \leq Q/4\kappa \tag{16}$$

Then for $(\bar{\mathbf{r}}', \bar{c}') = \mathbf{Dec}(u, \mathbf{v}, \Pi, t_1)$, we have:

$$\frac{\bar{\mathbf{r}}}{\bar{c}} \mod p = \frac{\bar{\mathbf{r}}'}{\bar{c}'} \mod p$$

Once we have verifiable encryption, adding traceability to our group signature is straightforward. During key generation we will create $(pk, sk) \leftarrow \mathbf{PKESetup}(1^\lambda)$, add pk to the group public key and set $gmsk = sk$. When signing a user will compute an encryption v of his randomness $\mathbf{r}$, which is such that $\mathbf{a}_1^T\mathbf{r} = t_1 \mod q_1$, and add v to the signature. For verification one only needs to check the extra proof Π. We consider how to open a signature, this is not completely straightforward because soundness only guarantees that a verifying signature will open to $\bar{c}\mathbf{r}$ for some $\bar{c} \in \bar{\mathcal{C}}$.
Open(msk, z) :

- Parse z as $(\mathbf{t}, \mathbf{t}', \Pi_1, \Pi_2, \Pi_3, v)$
- Let $(\bar{\mathbf{r}}, \bar{c})$ =Dec(msk, t_1, z)
- Set $id := \bar{c}^{-1}(t_2 - \mathbf{a}_2^T\bar{\mathbf{r}}) \in \mathcal{R}_{q_2}$
- If $id \in \mathbb{Z}_{q_2}$ then output id, otherwise output $\perp$

Note that if decryption succeeds then the proof Π verifies, which entails that there exists $\bar{\mathbf{r}}', \bar{c}'$ such that $\mathbf{a}_1^T\bar{\mathbf{r}}' = \bar{c}'t_1 \mod q_1$ and by lemma 5.1 we know that:

$$\frac{\bar{\mathbf{r}}}{\bar{c}} \mod p = \frac{\bar{\mathbf{r}}'}{\bar{c}'} \mod p$$

if we multiply this equation by $\bar{c}$ and $\bar{c}'$ we have that $\bar{\mathbf{r}}'\bar{c} = \bar{\mathbf{r}}\bar{c}' \mod p$, and since both sides are smaller than p this equation will be true over the integer. From which we get:

$$\mathbf{a}_1^T\bar{\mathbf{r}} = \bar{c}t_1 \mod q_1$$

which entails that if $\mathbf{t} = (t_1, t_2)$ is a well formed commitment the identity returned by the Open algorithm will be its message.

6 Security Proofs

6.1 Security of The Commitment

Lemma 6.1 (Binding). *Let $\kappa \geq \max_{c \in \mathcal{C}} (\|c\|_1)$. If there is an adversary $\mathcal{A}$ who can output a commitment $\mathbf{t}$ with two valid openings $(\mathbf{m}, \mathbf{r}, c)$ and $(\mathbf{m}', \mathbf{r}', c')$ such that $\mathbf{m} \neq \mathbf{m}'$ with probability ε, then there is an algorithm $\mathcal{A}'$ who can break M-SIS$_{q_1,n,m,4\kappa B_{Com}}$ in the same time and with advantage ε.*

We prove the hiding property for a slightly modified variant of our commitment scheme in which the error is sample according to $D^n_{\xi'} \times D^{k-n}_{\xi}$, where $\xi' = \sqrt{\frac{q_2}{q_1}\xi + 1 + 2d(k-n-l)\xi^2}$. This difference is mostly an artefact of the modulus switching used in the proof. We use the distribution S^k_1 in our paper as it makes for easier analysis and implementation and does not entail better attacks.

Lemma 6.2 (Hiding). *For any $\mathbf{m}, \mathbf{m}' \in \mathcal{R}^l_{q_2}$, if there is an adversary $\mathcal{A}$ who can distinguish between $Com(\mathbf{m})$ and $Com(\mathbf{m}')$ with advantage ε, then there exists an algorithm $\mathcal{A}'$ who runs in the same time and breaks M-LWE$_{q_2,m-n-l,\xi}$ with probability $\varepsilon/2$.*

Proof. Given an instance $(\mathbf{B}, \mathbf{y}) \in \mathcal{R}^{(n+l)\times(m-n-l)}_{q_2} \times \mathcal{R}^{n+l}_{q_2}$ of M-LWE$_{q_2,m-n-l,\xi}$, parse $\mathbf{B}$ and $\mathbf{y}$ as $\begin{bmatrix}\mathbf{B}_1\\ \mathbf{B}_2\end{bmatrix}$ and $\begin{bmatrix}\mathbf{y}_1\\ \mathbf{y}_2\end{bmatrix}$. Let $\rho : \mathbb{R} \to \mathbb{Z}$ be a randomized rounding function which maps $x \in \mathbb{R}$ to $\rho(x) \leftarrow \lfloor x \rfloor + B_{x-\lfloor x \rfloor}$, where $B_{x-\lfloor x \rfloor}$ is a Bernouilli variable which outputs 1 with probability $x - \lfloor x \rfloor$. Remark that for $q_1 \leq q_2$, $\rho\left(\mathcal{U}(\frac{q_1}{q_2}\mathbb{Z}_{q_2})\right) = \mathcal{U}(\mathbb{Z}_{q_1})$. Let $\mathbf{B}'_1 := \rho\left(\frac{q_1}{q_2}\mathbf{B}_1\right)$ and $\mathbf{y}'_1 := \rho(\frac{q_1}{q_2}\mathbf{y}_1)$. $\mathcal{A}'$ samples $\mathbf{R} \xleftarrow{\$} \mathcal{R}^{n\times l}_{q_1}$ and sets:

$$\mathbf{A} := \begin{bmatrix}\mathbf{I}_n & \mathbf{R}\\ \mathbf{0}^{l\times n} & \mathbf{I}_l\end{bmatrix} \cdot \begin{bmatrix}\mathbf{I}_n & \mathbf{0}^{n\times l} & \mathbf{B}'_1\\ \mathbf{0}^{l\times n} & \mathbf{I}_l & \mathbf{B}_2\end{bmatrix}$$

where the products are done over the integers and then taken modulo q_1 for the top part and modulo q_2 for the bottom part. $\mathcal{A}'$ sends $\mathbf{A}$ to the adversary $\mathcal{A}$ and receives messages $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{R}^l_{q_2}$ such that $\mathbf{m}_0 \neq \mathbf{m}_1$. $\mathcal{A}'$ samples $b \xleftarrow{\$} \{0,1\}$, computes:

$$\mathbf{t} = \begin{bmatrix}\mathbf{I}_n & \mathbf{R}\\ \mathbf{0}^{l\times n} & \mathbf{I}_l\end{bmatrix} \begin{bmatrix}\mathbf{y}'_1\\ \mathbf{y}_2\end{bmatrix} + \begin{bmatrix}\mathbf{0}\\ \mathbf{m}_b\end{bmatrix}$$

where the products are done over the integers and then taken modulo q_1 for the top part and modulo q_2 for the bottom part, and sends $\mathbf{t}$ to $\mathcal{A}$. When $\mathcal{A}$ returns b', $\mathcal{A}'$ returns 1 if $b' = b$ and 0 otherwise.
We first show that the public commitment matrix $\mathbf{A}$ is taken according to the correct distribution. We have:

$$\mathbf{A} = \begin{bmatrix}\mathbf{I}_n & \mathbf{R}\,\mathbf{B}'_1 + \mathbf{R}\mathbf{B}_2\\ \mathbf{0}^{l\times n} & \mathbf{I}_l \quad\quad \mathbf{B}_2\end{bmatrix}$$

Since $\mathbf{B}'_1$ is uniform modulo q_1 and independent from $\mathbf{R}$ and $\mathbf{B}_2$, $\mathbf{B}'_1 + \mathbf{R}\mathbf{B}_2 \mod q_1$ is uniform modulo q_1. Since $\mathbf{R}$ and $\mathbf{B}_2$ are also uniform, the distribution of $\mathbf{A}$ is identical to the one output by **CSetup**.
If $\mathbf{y}$ is uniform in $\mathcal{R}^{n+l}_{q_2}$ then $\mathbf{y}'_1$ is uniform in $\mathcal{R}^n_{q_1}$ and $\mathbf{t}$ is uniform in $\mathcal{R}^{n+l}_{q_1}$ and $b' = b$ with

probability exactly 1/2. However if $\mathbf{y} = \left[\mathbf{I}_{n+l}\ \mathbf{B}\right]\mathbf{r}$, write $\mathbf{r}$ as $\begin{bmatrix}\mathbf{r}_1\\ \mathbf{r}_2\\ \mathbf{r}_3\end{bmatrix}$, with $\mathbf{r}_1 \in \mathcal{R}^n$, $\mathbf{r}_2 \in \mathcal{R}^l$, and $\mathbf{r}_3 \in \mathcal{R}^{m-n-l}$. Applying ρ component-wise to $\frac{q_1}{q_2}(\mathbf{B}_1, \mathbf{B}_1\mathbf{r}_3 + \mathbf{r}_1)$ we get:

$$\begin{aligned}\left(\frac{q_1}{q_2}\mathbf{B}_1 + \Delta, \frac{q_1}{q_2}B_1\mathbf{r}_3 + \frac{q_1}{q_2}\mathbf{r}_1 + \delta\right) &= \left(\frac{q_1}{q_2}\mathbf{B}_1 + \Delta, \left(\frac{q_1}{q_2}B_1 + \Delta\right)\mathbf{r}_3\right.\\ &\qquad \left. + \frac{q_1}{q_2}\mathbf{r}_1 + \delta - \Delta\mathbf{r}_3\right)\\ &= \left(\mathbf{B}'_1, \mathbf{B}'_1\mathbf{r}_3 + \frac{q_1}{q_2}\mathbf{r}_1 + \delta - \Delta\mathbf{r}_3\right)\\ &= \left(\mathbf{B}'_1, \mathbf{B}'_1\mathbf{r}_3 + \mathbf{r}'_1\right)\end{aligned}$$

where $\mathbf{r}'_1$ is subgaussian with parameter $\sqrt{\frac{q_1}{q_2}\alpha + 1 + \|\mathbf{r}_3\|^2} \leq \xi'$. Setting $\mathbf{r}' = \begin{bmatrix}\mathbf{r}'_1\\ \mathbf{r}_2\\ \mathbf{r}_3\end{bmatrix}$ we have that

$$\mathbf{t} = \mathbf{A}\mathbf{r}' + \begin{bmatrix}\mathbf{0}\\ \mathbf{m}_b\end{bmatrix}$$

is distributed according to $Com(\mathbf{m}_b)$, and $\mathcal{A}$ will output $b' = b$ with probability $1/2+\varepsilon$. $\mathcal{A}'$ therefore has advantage $\varepsilon/2$ in the M-LWE$_{q_2,m-n-l,\xi}$ problem.

6.2 Security of the Scheme

Lemma 6.3 (Anonymity). *Let $\mathcal{A}$ be a PPT adversary. Let $\mathsf{Adv}^{Hid}_{\mathcal{A}}(\lambda)$ be the advantage of $\mathcal{A}$ over the Hiding property of the commitment scheme. Let $\mathsf{Adv}^{ind-cpa}_{\mathcal{A}}(\lambda)$ be the advantage of $\mathcal{A}$ over the IND-CPA property of the encryption scheme. The advantage of $\mathcal{A}$ against the CPA-anonymity of our group signature is at most:*

$$\mathsf{Adv}^{anon}_{\mathcal{A}}(\lambda) \leq 2\mathsf{Adv}^{Hid}_{\mathcal{A}}(\lambda) + \mathsf{Adv}^{ind-cpa}_{\mathcal{A}}(\lambda) + 2^{-\lambda}$$

Proof. We use a succession of games.
Game $\mathbf{G}_0$: In this game the challenger runs **GSetup** honestly and gives $(gpk, sk_1, \ldots, sk_N)$ to $\mathcal{A}$. $\mathcal{A}$ outputs a message M^* and two identities $i_0, i_1 \in [N]$. The challenger chooses a bit $b \xleftarrow{\$} \{0,1\}$ and computes $z^* := (\mathbf{t}, \mathbf{t}', e, \pi) \leftarrow \mathbf{Sign}(M^*, sk_{i_b})$
Game $\mathbf{G}_1$: In this the challenger uses the simulator of the proof Π_{Sign} when queried for $\mathbf{Sign}(M^*, sk_{i_b})$. This game is statistically indistinguishable from the previous by the zero-knowledge of Π_{Sign}.

$$\left|\mathsf{Adv}^{\mathbf{G}_1}_{\mathcal{A}} - \mathsf{Adv}^{\mathbf{G}_0}_{\mathcal{A}}\right| \leq 2^{-\lambda}$$

Game $\mathbf{G}_2$: In this the challenger replaces the commitment $\mathbf{t}$ by a commitment of 0 when answering the query $\mathbf{Sign}(M^*, sk_{i_b})$. The proof Π_{Sign} can still be used since it uses the simulator, and this game is indistinguishable from the previous one by the hiding property of the commitment.

$$\left|\mathsf{Adv}^{\mathbf{G}_2}_{\mathcal{A}} - \mathsf{Adv}^{\mathbf{G}_1}_{\mathcal{A}}\right| \leq \mathsf{Adv}^{Hid}_{\mathcal{A}}(\lambda)$$

Game $\mathbf{G}_3$: In this the challenger replaces the commitment $\mathbf{t}'$ by a commitment of 0 when answering the query $\mathbf{Sign}(M^*, sk_{i_b})$. This game is indistinguishable from the previous one by the hiding property of the commitment.

$$\left|\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_3} - \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_2}\right| \leq \mathsf{Adv}_{\mathcal{A}}^{Hid}(\lambda)$$

Game $\mathbf{G}_4$: In this the challenger replaces the commitment ciphertext e with an encryption of 0. Since the proof Π_{Sign} uses the simulator it is independent of the decryption of e. This game is indistinguishable from the previous one by the IND-CPA property of the encryption scheme.

$$\left|\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_4} - \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_3}\right| \leq \mathsf{Adv}_{\mathcal{A}}^{ind-cpa}(\lambda)$$

The signature $(\mathbf{t}, \mathbf{t}', e, \pi)$ output in **Game $\mathbf{G}_3$** is independent of i_b and the adversary has thus probability $1/2$ of outputting $b' = b$. We obtain the desired result by summing the advantages.

We will prove traceability in two steps. We will first prove that an adversary $\mathcal{A}$ cannot distinguish between the regular traceability game and the traceability game in which the setup algorithm has been replaced by $\mathbf{GSetup}^*$ which we define below. We will then prove that a challenger $\mathcal{B}$ can extract an M-SIS solution from an adversary who succeeds in producing a forgery in the traceability game with $\mathbf{GSetup}^*$.

$\mathbf{GSetup}^*(1^\lambda)$:

- Sample $i^* \xleftarrow{\$} \mathbb{Z}_{q_2}$
- Sample $\mathbf{A} := \begin{bmatrix}\mathbf{a}_1^T \\ \mathbf{a}_2^T\end{bmatrix} \leftarrow \mathbf{CSetup}(1^\lambda)$, with $\mathbf{a}_1 \in \mathcal{R}_{q_1}^3$, and $\mathbf{a}_2 \in \mathcal{R}_{q_2}^3$.
- Sample $\mathbf{a} \xleftarrow{\$} \mathcal{R}_{q_2}^2$.
- Sample $\mathbf{R} \xleftarrow{\$} S_1^{2\times 2}$ and set $\mathbf{b}^T = \mathbf{a}^T\mathbf{R} \in \mathcal{R}_{q_2}^{1\times 2}$.
- Sample $(\mathbf{s}_{i_1^*}, \mathbf{s}_{i_2^*}, \mathbf{s}_{i_3^*}) \leftarrow D_s \times D_s \times D_r$
- Set $u := \begin{bmatrix}\mathbf{a}^T \mid \mathbf{b}^T \mid \mathbf{a}_2^T\end{bmatrix} \begin{bmatrix}\mathbf{s}_{i_1^*} \\ \mathbf{s}_{i_2^*} \\ \mathbf{s}_{i_3^*}\end{bmatrix}$
- Set $gpk := \left(\mathbf{A}, \mathbf{a}, \mathbf{b} - i^*\mathbf{g}^T, u\right)$
- For $i \in \mathbb{Z}_{q_2} \setminus \{i^*\}$, sample $\mathbf{s}_{i_3} \leftarrow D_r^3$
- For $i \in \mathbb{Z}_{q_2} \setminus \{i^*\}$, sample $(\mathbf{s}_{i_1}, \mathbf{s}_{i_2}) \in \mathcal{R}^4$ s.t
  $\begin{bmatrix}\mathbf{a}^T \mid \mathbf{b}^T + (i - i^*)\mathbf{g}^T\end{bmatrix} \begin{bmatrix}\mathbf{s}_{i_1} \\ \mathbf{s}_{i_2}\end{bmatrix} = u - \mathbf{a}_2^T\mathbf{s}_{i_3}$, and $(\mathbf{s}_{i_1}, \mathbf{s}_{i_2}) \leftarrow D_s^4$
- For $i \in \mathbb{Z}_{q_2}$, set $sk_i := \mathbf{s}_i := (\mathbf{s}_{i_1}, \mathbf{s}_{i_2}, \mathbf{s}_{i_3})$

We consider the following advantages for an adversary $\mathcal{A}$

- $\mathsf{Adv}_{\mathcal{A}}^{trace}(\lambda)$ the advantage of $\mathcal{A}$ in the traceability game.
- $\mathsf{Adv}_{\mathcal{A}}^{trace*}(\lambda)$ the advantage of $\mathcal{A}$ in the traceability game where **GSetup** is replaced with $\mathbf{GSetup}^*$.
- $\mathsf{Adv}_{\mathcal{A}}^{NTRU}(\lambda)$ the advantage of $\mathcal{A}$ in solving the $\mathrm{NTRU}_{q,r}$ problem.
- $\mathsf{Adv}_{\mathcal{A}}^{MLWE}(\lambda)$ the advantage of $\mathcal{A}$ in solving the $\mathrm{M\text{-}LWE}_{q,1,s}$ problem.

Lemma 6.4. *The advantage of any PPT adversary $\mathcal{A}$ against the traceability game of the group signature is at most:*

$$\mathsf{Adv}_{\mathcal{A}}^{trace}(\lambda) \leq 2(\mathsf{Adv}_{\mathcal{A}}^{NTRU}(\lambda) + \mathsf{Adv}_{\mathcal{A}}^{MLWE}(\lambda)) + \mathsf{Adv}_{\mathcal{A}}^{trace*}(\lambda)$$

Proof. We use a succession of games.
Game $\mathbf{G}_0$: The challenger $\mathcal{B}$ runs the Group signature protocol honestly. He gives $(sk_i)_{i\in S}$ as well as $gmsk$ to $\mathcal{A}$ who has advantage ε in the traceability game.

$$\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_0} = \mathsf{Adv}_{\mathcal{A}}^{trace}$$

Game $\mathbf{G}_1$: $\mathcal{B}$ samples $\mathbf{a}_2^T$ as $[0 \mid 1 \mid f/g]$ where $f, g \in\leftarrow D_r$ are taken as in Section 2.6. $\mathbf{G}_2$ is indistinguishable from $\mathbf{G}_1$ under the $\text{NTRU}_{q,r}$ assumption.

$$\left|\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_1} - \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_0}\right| \leq \mathsf{Adv}_{\mathcal{A}}^{NTRU}$$

Game $\mathbf{G}_2$: $\mathcal{B}$ sets $\mathbf{b}^T \xleftarrow{\$} \mathcal{R}_{q_2}^{1\times 2}$. Note that if $\mathbf{b}^T \neq \mathbf{a}^T\mathbf{R}$, $\mathcal{B}$ can no longer use the GPV trapdoor of $\left[\mathbf{a}^T \mid \mathbf{b}^T + i\mathbf{g}^T\right]$ to sample secret keys for user i. To generate keys for i he will instead sample $\mathbf{s}_{i_1}, \mathbf{s}_{i_2} \leftarrow D_s^2$ and use his NTRU trapdoor on $\mathbf{a}_2$ to sample $\mathbf{s}_{i_3}$. This game will be indistinguishable from the previous one by the hardness of $\text{M-LWE}_{q,1,s}$ (since $\mathbf{a}^T\mathbf{R}$ is two $\text{M-LWE}_{q,1,s}$ samples).

$$\left|\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_2} - \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_1}\right| \leq \mathsf{Adv}_{\mathcal{A}}^{MLWE}$$

Game $\mathbf{G}_3$: $\mathcal{B}$ replaces $\mathbf{b}^T$ with $\mathbf{b}^{*T} := \mathbf{b}^T - i^*\mathbf{g}^T$. Since $\mathbf{b}^T$ is uniform this game is identical to the previous one.

$$\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_3} = \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_2}$$

Game $\mathbf{G}_4$: $\mathcal{B}$ sets $\mathbf{b}^{*T} := \mathbf{a}^T\mathbf{R} - i^*\mathbf{g}^T$. This game is indistinguishable from the previous one under $\text{M-LWE}_{q,1,s}$.

$$\left|\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_4} - \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_3}\right| \leq \mathsf{Adv}_{\mathcal{A}}^{MLWE}$$

Game $\mathbf{G}_5$: $\mathcal{B}$ sets $\mathbf{a}_2^T$ as $[0 \mid 1 \mid a_2]$, with $a_2 \xleftarrow{\$} \mathcal{R}_{q_2}$ and uses the GPV trapdoor of $\left[\mathbf{a}^T \mid \mathbf{b}^T + (i - i^*)\mathbf{g}^T\right]$ to sample secret keys for user i. This game is indistinguishable from the previous one under the $\text{NTRU}_{q,r}$ assumption.

$$\left|\mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_5} - \mathsf{Adv}_{\mathcal{A}}^{\mathbf{G}_4}\right| \leq \mathsf{Adv}_{\mathcal{A}}^{NTRU}$$

Note that **Game $\mathbf{G}_5$** is the traceability game that uses $\mathbf{GSetup}^*$ (simply by renaming $\mathbf{s}_0$ to $\mathbf{s}_{i^*}$), the result follows.

Lemma 6.5. *Let $\mathcal{A}$ be a PPT algorithm with advantage ε in the traceability game with $\mathbf{GSetup}^*$. Let h be a bound on the number of hash queries made by $\mathcal{A}$. Let $B_S \geq 4\kappa B_1 + 12\kappa\sqrt{d}B_1 + 2\kappa B_2 + 6\sqrt{d}\xi B_1 + \kappa^2(1+3\sqrt{d})2\sqrt{d}s + \kappa^2\sqrt{6d}s$. There exists $\mathcal{B}$ a challenger for the $\text{M-SIS}_{q_2,1,4,B_S}$ such that:*

$$\mathsf{Adv}_{\mathcal{B}}^{MSIS}(\lambda) \geq \frac{\varepsilon}{q_2}\left(\frac{\varepsilon}{h} - 2^{-\lambda}\right)$$

Proof. Formally $\mathcal{B}$ is given a matrix $\mathbf{x}^T := \left[x_1, x_2, x_3, x_4\right] \in \mathcal{R}_{q_2}^4$ and must output $\mathbf{y}$ s.t $\mathbf{x}^T\mathbf{y} = 0 \mod q_2$ and $\|\mathbf{y}\| \leq B_S$, w.l.o.g we consider $\mathbf{x} = \left[x_1, x_2, x_3, 1\right]$ instead since with high probability one of the x_i will have an inverse.
$\mathcal{B}$ will set $\mathbf{a} := (x_1, x_2)$ and $\mathbf{a}_2^T := (0, 1, x_3)$ during setup, since x_1, x_2, x_3 are uniform in $\mathcal{R}_{q_2}$ this does not change the distribution of $\mathbf{GSetup}^*$. When asked signing queries, $\mathcal{B}$ runs the signing algorithm honestly, when asked corrupt queries $\mathcal{B}$ outputs the corresponding secret key. Suppose

the adversary $\mathcal{A}$ outputs a forgery $z := (\mathbf{t}, \mathbf{t}', \Pi, (u, \mathbf{v}))$ by programming the random oracle with two different challenges $\mathcal{B}$ will be able to extract $\bar{\mathbf{z}} \in \mathcal{R}^3$, $\bar{id} \in \mathbb{Z}_{q_2}$, $\bar{\mathbf{z}}' \in \mathcal{R}^3$, $\bar{\mathbf{z}}_s \in \mathcal{R}^7$, $\bar{\mathbf{z}}_B \in \mathcal{R}^8$, $\bar{c} \in \bar{\mathcal{C}}$ such that:

$$\begin{aligned}
&\bar{c}\mathbf{t} = Com(\bar{c}\bar{id}; \bar{\mathbf{z}}) \\
&\bar{c}\mathbf{t}' = Com(\bar{c}\bar{id}\delta; \bar{\mathbf{z}}') \\
&\bar{c}\begin{bmatrix} u \\ \mathbf{v} \\ t_1 \end{bmatrix} = \mathbf{B}\bar{\mathbf{z}}_B \\
&\bar{c}u = \mathbf{v}^T\bar{\mathbf{z}}_s
\end{aligned}$$

such that $\|(\bar{\mathbf{z}}, \bar{\mathbf{z}}', \bar{\mathbf{z}}_B)\| \leq 2B \wedge \|\bar{\mathbf{z}}_1, \bar{\mathbf{z}}_2\| \leq 2B_1 \wedge \|\bar{\mathbf{z}}_3\| \leq 2B_2$, with $(\bar{\mathbf{z}}_1, \bar{\mathbf{z}}_2, \bar{\mathbf{z}}_3) := \bar{\mathbf{z}}$. Using the forking lemma of [BN06], $\mathcal{B}$ will be able to do this with probability at least $\varepsilon\left(\frac{\varepsilon}{h} - 2^{-\lambda}\right)$. Let $(\tilde{\mathbf{r}}, \tilde{c})$:=Dec(u, v), the parameters set in section 8 are such that, by soundness of the verifiable encryption scheme, with overwhelming probability $\tilde{\mathbf{r}}\bar{c} = \bar{\mathbf{z}}\tilde{c}$ over the integers, which implies that $\mathbf{Open}(z) \in \mathbb{Z}_{q_2}$ i.e. the forgery opens to an identity in $\mathbb{Z}_{q_2}$ and not $\perp$.
Since i^* is taken uniformly at random in $\mathbf{GSetup}^*$, z will open to this identity with probability $1/q_2$. Suppose that z opens to i^*. Then

$$\begin{aligned}
&\bar{c}t_2 = \mathbf{a}_2^T\bar{\mathbf{z}} + \bar{c}i^* \\
&\bar{c}t_2' = \mathbf{a}_2^T\bar{\mathbf{z}}' + \bar{c}i^*\delta \\
&\left[\mathbf{a}^T \mid \mathbf{b}^T + [t_2 \mid t_2'] - i^*\mathbf{g}^T \mid \mathbf{a}_2^T\right]\bar{\mathbf{z}}_s = \bar{c}u
\end{aligned}$$

If we multiply the third equation by $\bar{c}$ and replace $\bar{c}[t_2 \mid t_2']$ we get:

$$\left[\bar{c}\mathbf{a}^T \mid \bar{c}\mathbf{b}^T + [\mathbf{a}_2^T\bar{\mathbf{z}} \mid \mathbf{a}_2^T\bar{\mathbf{z}}'] \mid \bar{c}\bar{\mathbf{a}}_2^T\right]\bar{\mathbf{z}}_s = \bar{c}^2u$$

Let

$$\tilde{\mathbf{z}} = \begin{bmatrix} \bar{c}\bar{\mathbf{z}}_1 + \mathbf{R}\bar{c}\bar{\mathbf{z}}_2 \\ \bar{c}\bar{\mathbf{z}}_3 - \left[\bar{\mathbf{z}}\ \bar{\mathbf{z}}'\right]\bar{\mathbf{z}}_2 \end{bmatrix}$$

Then

$$\left[\mathbf{a}^T \mid \bar{\mathbf{a}}_2^T\right]\tilde{\mathbf{z}} = \bar{c}^2u$$

Since $\mathcal{A}$ has to output a valid forgery this means that he has never obtained the key sk_{i^*}, we can thus consider that $\mathbf{s}_{i^*}$ was sampled after receiving the forgery, conditioned on $\left[\mathbf{a}^T \mid \mathbf{b} \mid \mathbf{a}_2^T\right]\mathbf{s}_{i^*} = u$. Let $\mathbf{s}^* := \left[\mathbf{s}_{i_1^*} + \mathbf{R}\mathbf{s}_{i_2^*} \mid \mathbf{s}_{i_3^*}\right]$, the probability that $\bar{c}\mathbf{s}^* = \tilde{\mathbf{z}}$ is negligible. Finally we have a solution $\tilde{\mathbf{z}} - \bar{c}\mathbf{s}^*$ to the M-SIS problem defined by $\left[\mathbf{a}^T \mid \mathbf{a}_2^T\right]$. Using the bounds on the extracted values and the distribution of $\mathbf{s}^*$ we have the following bound on the norm of the solution:

$$\begin{aligned}
\|\tilde{\mathbf{z}} - \bar{c}\mathbf{s}^*\| &\leq \|\tilde{\mathbf{z}}\| + 2\kappa\|\mathbf{s}^*\| \\
&\leq 2\kappa\|\mathbf{z}_1\| + 6\kappa\sqrt{d}\|\bar{\mathbf{z}}_2\| + \kappa\|\bar{\mathbf{z}}_3\| + 3\sqrt{d}\xi\|\bar{\mathbf{z}}_2\| \\
&\quad + \kappa^2(1 + 3\sqrt{d})(2\sqrt{d}s) + \kappa^2\sqrt{6d}s \\
&\leq 4\kappa B_1 + 12\kappa\sqrt{d}B_1 + 2\kappa B_2 + 6\sqrt{d}\xi B_1 \\
&\quad + \kappa^2(1 + 3\sqrt{d})2\sqrt{d}s + \kappa^2\sqrt{6d}s
\end{aligned}$$

The largest terms in this solution are by far $2\kappa B_2$ and $6\sqrt{d}\xi B_1$ which we will consider when setting the parameters in section 8.

7 The Full Non-Interactive Proof

We give the full non-interactive zero-knowledge proof that the signer will output. We only consider the parameter choice made in section 8. The user $i \in \mathbb{Z}_{q_2}$ will use the following elements for his proof:

$$\mathbf{t} := \begin{bmatrix} t_1 \\ t_2 \end{bmatrix} = \mathbf{A}\mathbf{r} + \begin{bmatrix} 0 \\ i \end{bmatrix} \in \mathcal{R}_{q_1} \times \mathcal{R}_{q_2}$$
$$\mathbf{t}' := \begin{bmatrix} t'_1 \\ t'_2 \end{bmatrix} = \mathbf{A}\mathbf{r}' + \begin{bmatrix} 0 \\ i\delta \end{bmatrix} \in \mathcal{R}_{q_1} \times \mathcal{R}_{q_2}$$
$$\mathbf{v}^T = \left[\mathbf{a}^T \mid \mathbf{b}^T + \left[t_2\ t'_2\right] \mid \mathbf{a}_2^T\right] \in \mathcal{R}_{q_2}^{1\times 7}$$
$$\mathbf{s}' = \begin{bmatrix} \mathbf{s}_{i_1} \\ \mathbf{s}_{i_2} \\ \mathbf{s}_{i_3} - \left[\mathbf{r}\ \mathbf{r}'\right]\mathbf{s}_{i_2} \end{bmatrix} \in \mathcal{R}^7$$

We first note that since $\mathbf{a}_2^T = \left[0\ 1\ a'_2\right]$, we can ignore the 5^{th} coefficient of $\mathbf{v}^T$ (corresponding to 0) in $\mathbf{v}^T\mathbf{s}' = u$ and thus consider $\mathbf{v}^T \in \mathcal{R}_{q_2}^{1\times 6}$ and $\mathbf{s}' \in \mathcal{R}^6$ such that $\mathbf{v}^T\mathbf{s}' = u$. The gain in proof size obtained by discarding one element of this equation may seems negligible at first but it is in fact rather important because the last three coefficients of $\mathbf{s}'$ will be much larger than the other four. We also recall the matrices needed for the proof of verifiable encryption:

$$\mathbf{B}_1 = \begin{bmatrix} pa & p & 0 & 0 & 0 & 0 & 0 & 0 \\ pb_1 & 0 & p & 0 & 0 & 1 & 0 & 0 \\ pb_2 & 0 & 0 & p & 0 & 0 & 1 & 0 \\ pb_3 & 0 & 0 & 0 & p & 0 & 0 & 1 \end{bmatrix} \in \mathcal{R}_Q^{4\times 8}$$
$$\mathbf{B}_2 = \left[\mathbf{0}^{1\times 5}\ \mathbf{a}_1^T\right] \in \mathcal{R}_{q_1}^{1\times 8}$$
$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \end{bmatrix}$$

Which are such that :

$$\mathbf{B}\mathbf{r}_B = \begin{bmatrix} u & \bmod Q \\ v_1 & \bmod Q \\ v_2 & \bmod Q \\ v_3 & \bmod Q \\ t_1 & \bmod q_1 \end{bmatrix}, \text{ for } \mathbf{r}_B = \begin{bmatrix} r \\ e_1 \\ \mathbf{e}_2 \\ \mathbf{r} \end{bmatrix}$$

An important point for proof size will be rejection sampling. After doing rejection sampling $\mathrm{Rej}(\mathbf{z}, \mathbf{a}, \xi)$ on a vector $\mathbf{z}$ we know by lemma 2.2 that all of its coefficients will be statistically close to D_ξ with $\xi \geq 11\,\|\mathbf{a}\|$, meaning that for very imbalanced vectors it would be worthwhile to do rejection sampling multiple times. For example if $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2)$ with $\|\mathbf{a}_2\| >> \|\mathbf{a}_1\|$ then by doing two rejection samplings $\mathrm{Rej}(\mathbf{z}_1, \mathbf{a}_1, \xi_1)$ and $\mathrm{Rej}(\mathbf{z}_2, \mathbf{a}_2, \xi_2)$ one obtains a smaller vector $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2)$ at the cost of having acceptation probability $1/9$, since the proof is non-interactive aborts have a minimal impact and this approach can help reduce the proof size significantly. We will use two rejection samplings for $\mathbf{s}' \in \mathcal{R}^6$ in which the last two coefficients will be much larger than the other four (because they correspond to a product of $\mathbf{s}_i$ and $\mathbf{r}, \mathbf{r}'$). We will thus write $\mathbf{s}'$ as $\mathbf{s}' = (\mathbf{s}'_1, \mathbf{s}'_2) \in \mathcal{R}^4 \times \mathcal{R}^2$. We can now write the full zero-knowledge proof of the verifier.

Algorithm 2 Π_{Sign}

Require: Message $M \in \{0,1\}^*$.
Public information: $\mathbf{t}, \mathbf{t}', \mathbf{v}^T, \mathbf{B}, \delta = \lfloor \sqrt{q} \rceil, \sigma_{-1}, \sigma_5$.
Private information: $\mathbf{r}, \mathbf{r}', i, \mathbf{s}', \mathbf{r}_B$

$\mathbf{y}, \mathbf{y}', \mathbf{y}_{-1}, \mathbf{y}_5 \leftarrow D_\xi^3$
$\mathbf{y}_B \leftarrow D_\xi^8$
$\mathbf{y}_{s_1} \leftarrow D_{\xi_1}^4$
$\mathbf{y}_{s_2} \leftarrow D_{\xi_2}^2$
$\mathbf{y}_s = (\mathbf{y}_{s_1}, \mathbf{y}_{s_2})$
$\mathbf{w}_1 := \mathbf{a}_1^T \mathbf{y}$
$\mathbf{w}_1' := \mathbf{a}_1^T \mathbf{y}'$
$\mathbf{w}_{1,-1} := \sigma_{-1}(\mathbf{a}_1^T)\mathbf{y}_{-1}$
$\mathbf{w}_{1,5} := \sigma_5(\mathbf{a}_1^T)\mathbf{y}_5$
$\mathbf{w}_2 := \delta \mathbf{a}_2^T \mathbf{y} - \mathbf{a}_2^T \mathbf{y}'$
$\mathbf{w}_{2,-1} := \mathbf{a}_2^T \mathbf{y} - \sigma_{-1}(\mathbf{a}_2)\mathbf{y}_{-1}$
$\mathbf{w}_{2,5} := \mathbf{a}_2^T \mathbf{y} - \sigma_5(\mathbf{a}_2)\mathbf{y}_5$
$\mathbf{w}_s := \mathbf{v}^T \mathbf{y}_s$
$\mathbf{w}_B := \mathbf{B}\mathbf{y}_B$
$c := H(\mathbf{t}, \mathbf{t}', \mathbf{v}, \mathbf{B}, \delta, \sigma_{-1}, \sigma_5, \mathbf{w}_1, \mathbf{w}_1', \mathbf{w}_{1,-1}, \mathbf{w}_{1,5}, \mathbf{w}_2, \mathbf{w}_{2,-1}, \mathbf{w}_{2,5}, \mathbf{w}_s, \mathbf{w}_B, M)$
$\mathbf{z} := \mathbf{r}c + \mathbf{y}$
$\mathbf{z}' := \mathbf{r}'c + \mathbf{y}'$
$\mathbf{z}_{-1} := \sigma_{-1}(\mathbf{r})c + \mathbf{y}_{-1}$
$\mathbf{z}_5 := \sigma_5(\mathbf{r})c + \mathbf{y}_5$
$\mathbf{z}_{s_1} := \mathbf{s}_1' c + \mathbf{y}_{s_1}$
$\mathbf{z}_{s_2} := \mathbf{s}_2' c + \mathbf{y}_{s_2}$
$\mathbf{z}_B := \mathbf{r}_B c + \mathbf{y}_B$
if $\mathrm{Rej}((\mathbf{z}, \mathbf{z}', \mathbf{z}_{-1}, \mathbf{z}_5, \mathbf{z}_B), (\mathbf{r}c, \mathbf{r}'c, \sigma_{-1}(\mathbf{r})c, \sigma_5(\mathbf{r})c, \mathbf{r}_B c), \xi) \wedge \mathrm{Rej}(\mathbf{z}_{s_1}, \mathbf{s}_1' c, \xi_1) \wedge \mathrm{Rej}(\mathbf{z}_{s_2}, \mathbf{s}_2' c, \xi_2)$ **then**
    Output $z = (\mathbf{z}, \mathbf{z}', \mathbf{z}_{-1}, \mathbf{z}_5, \mathbf{z}_{s_1}, \mathbf{z}_{s_2}, \mathbf{z}_B, c)$
else
    Restart
end if

Algorithm 3 Verify

Require: Message $M \in \{0,1\}^*$.
Signature $\Pi = (\mathbf{z}, \mathbf{z}', \mathbf{z}_{-1}, \mathbf{z}_5, \mathbf{z}_{s_1}, \mathbf{z}_{s_2}, \mathbf{z}_B)$.
Public information: $\mathbf{t}, \mathbf{t}', \mathbf{v}^T, \mathbf{B}, \delta = \lfloor \sqrt{q} \rceil, \sigma_{-1}, \sigma_5$.

$\mathbf{w}_1 := \mathbf{a}_1^T \mathbf{z} - t_1 c$
$\mathbf{w}_1' := \mathbf{a}_1^T \mathbf{z}' - t_1' c$
$\mathbf{w}_{1,-1} := \sigma_{-1}(\mathbf{a}_1^T)\mathbf{z}_{-1} - \sigma_{-1}(t_1)c$
$\mathbf{w}_{1,5} := \sigma_5(\mathbf{a}_1^T)\mathbf{z}_5 - \sigma_5(t_1)c$
$\mathbf{w}_2 := \delta \mathbf{a}_2^T \mathbf{z} - \mathbf{a}_2^T \mathbf{z}' - (\delta t_2 - t_2')\, c$
$\mathbf{w}_{2,-1} := \mathbf{a}_2^T \mathbf{z} - \sigma_{-1}(\mathbf{a}_2^T)\mathbf{z}_{-1} - (t_2 - \sigma_{-1}(t_2))\, c$
$\mathbf{w}_{2,5} := \mathbf{a}_2^T \mathbf{z} - \sigma_5(\mathbf{a}_2^T)\mathbf{z}_5 - (t_2 - \sigma_5(t_2))\, c$
$\mathbf{w}_s := \mathbf{v}^T \mathbf{z}_s - uc$
$\mathbf{w}_B := \mathbf{B}\mathbf{z}_B - (v, v_1, v_2, v_3, t_1)c \in \mathcal{R}_Q^4 \times \mathcal{R}_{q_1}$
if $\|(\mathbf{z}, \mathbf{z}', \mathbf{z}_{-1}, \mathbf{z}_5, \mathbf{z}_B)\| \le B \wedge \|\mathbf{z}_{s_1}\| \le B_1 \wedge \|\mathbf{z}_{s_2}\| \le B_2$
and $c = H(\mathbf{t}, \mathbf{t}', \mathbf{v}, \mathbf{B}, \delta, \sigma_{-1}, \sigma_5, \mathbf{w}_1, \mathbf{w}_1', \mathbf{w}_{1,-1}, \mathbf{w}_{1,5}, \mathbf{w}_2, \mathbf{w}_{2,-1}, \mathbf{w}_{2,5}, \mathbf{w}_s, \mathbf{w}_B, M)$ **then**
    Output 1
else
    Output 0
end if

Lemma 7.1. *Let* $\mathbf{r}, \mathbf{r}' \leftarrow S_1^3$, *let* $\mathbf{s}_{i_1}, \mathbf{s}_{i_2} \leftarrow D_s^2$, *let* $\mathbf{s}_{i_3} \leftarrow D_r^3$, *let* $u, v_1, v_2, v_3 \leftarrow S_1$. *Let* $\mathbf{t}, \mathbf{t}', \mathbf{v}, \mathbf{s}', \mathbf{B}, \mathbf{r}_B$ *be defined as previously. Let* $\xi \geq 11\kappa\sqrt{20d}$, $\xi_1 \geq 11\kappa\sqrt{8ds}$, $\xi_2 \geq 11\kappa(d\sqrt{24}s + \sqrt{2d}r)$ *and* $B \geq 2\sqrt{10d}\xi$, $B_1 \geq 2\sqrt{2d}\xi_1$, $B_2 \geq 2\sqrt{d}\xi_2$. *If* $B_{com} \geq 2B$, *then the algorithm* Π_{sign} *achieves the following properties:*

- Correctness: *The prover restarts with probability at most* $1/27 + 2^{-100}$, *and if he does not abort the verifier accepts with overwhelming probability.*
- Honest-Verifier Zero-Knowledge: *Signatures can be simulated with statistically indistinguishable distribution.*
- Special Soundness: *Given two accepting transcripts one can extract* $\bar{\mathbf{z}} \in \mathcal{R}^3$, $\bar{id} \in \mathbb{Z}_{q_2}$, $\bar{\mathbf{z}}' \in \mathcal{R}^3$, $\bar{\mathbf{z}}_s \in \mathcal{R}^7$, $\bar{\mathbf{z}}_B \in \mathcal{R}^8$, $\bar{c} \in \bar{\mathcal{C}}$ *such that:*

$$\begin{aligned}
&\bar{c}\mathbf{t} = Com(\bar{c}\bar{id}; \bar{\mathbf{z}}) \\
&\bar{c}\mathbf{t}' = Com(\bar{c}\bar{id}\delta; \bar{\mathbf{z}}') \\
&\bar{c}\begin{bmatrix} u \\ v_1 \\ v_2 \\ v_3 \\ t_1 \end{bmatrix} = \mathbf{B}\bar{\mathbf{z}}_B \\
&\bar{c}u = \mathbf{v}^T\bar{\mathbf{z}}_s
\end{aligned}$$

such that $\|(\bar{\mathbf{z}}, \bar{\mathbf{z}}', \bar{\mathbf{z}}_B)\| \leq 2B \wedge \|\bar{\mathbf{z}}_{s_1}\| \leq 2B_1 \wedge \|\bar{\mathbf{z}}_{s_2}\| \leq 2B_2$.

Proof. The proof is simply a combination of the proofs for Lemma 4.1 and Corollary 4.2.

8 Parameters and Implementation

8.1 Fixing the Parameters

We will set the parameters as per Table 5. In this section we discuss the bounds that must be verified by these parameters and the resulting security guarantees. We will consider the security of our scheme in terms of "root-hermite factor" δ_0 which is a parameter often used when assessing the security of lattice-based schemes. In this section we will aim for a root-hermite factor $\delta_0 = 1.0036$. Such a factor implies at least 93-bit (post-quantum) space hardness or 242-bit time hardness (and significantly smaller space-hardness) depending on which lattice reduction strategy one uses.

First we fix the dimension to $d = 4096$, we use this dimension as anything smaller does not allow the existence of parameters that make our scheme secure. For this dimension the challenge set $\{c \in \mathcal{R} \ : \ \|c\|_1 = \kappa, \|c\|_\infty = 1\}$ will be of size greater than 2^{256} if we fix $\kappa = 26$.

The standard deviations s and r are fixed by the quality of our trapdoors to $s = 6\sqrt{dq_2}$ and $r = 2 \cdot 1.17\sqrt{q_2}$.

To fix q_1 we consider the requirements on the binding property of our commitment. From Section 5 we have that the M-SIS$_{q_1,1,3,4\kappa B}$ problem has to be hard for vectors of norm $4\kappa B = 88 \cdot \kappa^2 \cdot d\sqrt{200}$. For $q_1 \approx 2^{30}$, as in table 5, we obtain $\delta_0 = 1.0036$.
To fix q_2 we will need the M-SIS$_{q_2,1,4,B_S}$ for B_S as per lemma 6.5 to be hard. We can compute that $B_S = \|\tilde{\mathbf{z}} - \mathbf{s}^*\| \approx 180224 \cdot \sqrt{2} \cdot d^2 \cdot \sqrt{q_2}$ for $q_2 \approx 2^{80}$ as in table 5 the root-hermite factor of this problem will be $\delta_0 = 1.0036$.

The value of q_2 also affects the hardness of the hiding property of our commitment scheme. For this, we need that R-LWE is hard for dimension d, errors sampled in S_1 and both modulus q_1 and q_2. In practice the best attack will be to either solve R-LWE modulo q_1 or modulo q_2. Since we will have $q_2 > q_1$ and the hardness of R-LWE decreases as the modulus increases, we will only consider q_2 as being relevant here. For the parameters in table 5 we have $\delta_0 = 1.0036$ for R-LWE with modulo q_2, which is set intentionally to be the same as the root-hermite factor for the M-SIS problem above.

The only constraint we have on p, the plaintext modulus of our verifiable encryption scheme, is that if $\bar{\mathbf{r}}\tilde{c} = \tilde{\mathbf{r}}\bar{c} \mod p$ for some $\bar{\mathbf{r}}, \bar{c}, \tilde{\mathbf{r}}, \tilde{c}$ extracted in Π_{Sign} then this equation should hold over the integers, i.e. $\|\bar{\mathbf{z}}\tilde{c} - \tilde{\mathbf{z}}\bar{c}\|_\infty \leq p/2$. Since the vector $\mathbf{z}$ output in Π_{Sign} will have coefficients distributed according to D_ξ, we will have with overwhelming probability that $\|\mathbf{z}\|_\infty \leq 12\xi$ (we can add this as an explicit check in the verification algorithm), in which case we will require $p \geq 4 \cdot \kappa \cdot 12\xi \geq 2^{26.5}$. The ciphertext modulus Q will be fixed by equation 16 which gives:

$$Q \geq 264\sqrt{34}\kappa p d^{3/2} \geq 2^{59.5}$$

We consider the proof size that results from this parameter choice. The secret key consists of 4 polynomials of standard deviation s and two polynomials of standard deviation r resulting in a size of $4d\log(12s) + 2d\log(12r) = 154KB$. The signature itself will consist of two commitments, one ciphertext and one zero-knowledge proof, which are respectively of size:

$$\begin{aligned} 2d\log q_1 + 2d\log q_2 &= 113KB \\ 4d\log Q &= 123KB \\ 13d\log(12\xi) + 4d\log(12\xi_1) + 2d\log(12\xi_2) &= 345KB \end{aligned}$$

8.2 Accounting For Complexity Leveraging

The proof for full-traceability of Section 6.2 reduces the security of our group signature to that of a selectively secure signature. When guessing the identity of the forgery, we thus lose a factor of q_2 in the success probability of the attacker. The hardness of SIS is usually evaluated by considering exponential time/exponential space algorithms (sieving algorithms) since such algorithms have the best asymptotic complexity. While complexity leveraging implies that the running time of a successful SIS challenger is multiplied by q_2 it should not be affected in terms of space complexity, it is thus reasonable to consider the space complexity of the adversary as a lower bound on the security of the scheme. To account for the loss in success probability we will also consider polynomial space/exponential time algorithms (enumeration algorithms). For a root-hermite factor $\delta_0 = 1.0036$ enumeration estimates (e.g. [ACD+18,BCLvV]) give a post-quantum time complexity of 322 bits, which when taking into account the loss of a factor q_2 results in a security of 242 bits. For the sake of completeness, we also give a very conservative second set of parameters in which the post-quantum time and space complexity of sieving algorithms is above 128 bits even with complexity leveraging. The dimension $d = 4096$ is no longer enough to reach such a security and we will thus be forced to set $d = 8192$. For this dimension the M-SIS$_{q_2,1,3,B_S}$ has a root hermite factor of $\delta_0 = 1.002$, corresponding to a security of 207 bits in space and 262 bits in time for post-quantum security using sieving (the post-quantum time security when considering enumeration is of 1084 bits with leveraging). The R-LWE problem in dimension $d = 8192$ has a root-hermite factor of less than 1.0019 resulting in more than 300 bits of security, similarly the M-SIS$_{q_1,1,3,4\kappa B}$ becomes harder

with a higher dimension and q_1 can be reduced accordingly. The rest of the parameters will be changed to the values given in Table 5, resulting in a signature size of (c.f. Table 4).

8.3 Implementation

We have implemented the group signature scheme in the C programming language. For the preimage sampling during key generation we use the Fast Fourier version [DP16] of the randomized nearest plane algorithm [GPV08] adapted to cyclotomic rings. This was done before in the Falcon signature scheme [PFH+18]. We also use the FFT-based algorithm from [DP16] adapted to cyclotomic rings to compute the compact LDL^* decomposition of the trapdoor basis. Contrary to Falcon, double floating point precision does not suffice in our case. For the necessary multiprecision complex arithmetic we use the library MPC [EGTZ18] based on MPFR [FHL+07] and GMP [Gt16]. We compute everything with 256 bits of precision. In the complex FFT we use Cooley-Tukey butterflies in the forward transform, Gentleman-Sande butterflies in the inverse transform and no reordering. The signing algorithm does not need multiprecision floating point arithmetic and is thus suitable for small devices. It mainly requires multiprecision integer polynomial arithmetic which we have implemented using the GMP integer multiprecision library. It is possible to optimize this part of our code. For example in the zero-knowledge proofs of the commitment scheme we only need that differences of challenges are invertible modulo the second prime q_2. The first prime q_1 could thus be chosen to be fully splitting and polynomial multiplications be computed with an NTT-based algorithm. We obtain all randomness from the SHAKE-256 expandable output function [BDPVA13].

Table 1 lists the running times of the implementation. The time for key generation contains the generation of the group public key and one member secret key. For further improvements, the tree representing the L matrix of the compact LDL^* decomposition of the trapdoor basis can be precomputed, which would significantly reduce the key generation time for each individual member.

References

[ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115, 2010.

[ACD+18] Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the {LWE, NTRU} schemes! *IACR Cryptology ePrint Archive*, 2018:331, 2018.

[AG11] Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In *ICALP (1)*, pages 403–415, 2011.

[Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993.

[BCLvV] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. NTRU prime. technical report, national institute of standards and technology, 2017.

[BCN17] Cecilia Boschini, Jan Camenisch, and Gregory Neven. Relaxed lattice-based signatures with short zero-knowledge proofs. Cryptology ePrint Archive, Report 2017/1123, 2017.

[BCN18] Cecilia Boschini, Jan Camenisch, and Gregory Neven. Floppy-sized group signatures from lattices. In *ACNS*, pages 163–182, 2018.

[BDK+18] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In *EuroS&P*, pages 353–367, 2018.

[BDL+18] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In *SCN*, pages 478–498, 2018.

[BDPVA13] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Keccak. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, pages 313–314, 2013.

[BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325, 2012.

[BMW03] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *EUROCRYPT*, pages 614–629, 2003.

[BN06] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In *ACM Conference on Computer and Communications Security*, pages 390–399, 2006.

[Boy10] Xavier Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Public Key Cryptography*, pages 499–517, 2010.

[CL02] Jan Camenisch and Anna Lysyanskaya. A signature scheme with efficient protocols. In *SCN*, pages 268–289, 2002.

[Coh00] H. Cohen. *A Course in Computational Algebraic Number Theory*. Graduate Texts in Mathematics. Springer Berlin Heidelberg, 2000.

[CvH91] David Chaum and Eugène van Heyst. Group signatures. In *EUROCRYPT*, pages 257–265, 1991.

[DKL+18] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.

[DLP14] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In *ASIACRYPT*, pages 22–41, 2014.

[DM15] Léo Ducas and Daniele Micciancio. FHEW: bootstrapping homomorphic encryption in less than a second. In *EUROCRYPT*, pages 617–640, 2015.

[DP16] Léo Ducas and Thomas Prest. Fast fourier orthogonalization. In *ISSAC*, pages 191–198, 2016.

[EGTZ18] Andreas Enge, Mickaël Gastineau, Philippe Théveny, and Paul Zimmermann. *mpc — A library for multiprecision complex arithmetic with exact rounding*. INRIA, 1.1.0 edition, January 2018. http://mpc.multiprecision.org/.

[FHL+07] Laurent Fousse, Guillaume Hanrot, Vincent Lefèvre, Patrick Pélissier, and Paul Zimmermann. Mpfr: A multiple-precision binary floating-point library with correct rounding. *ACM Trans. Math. Softw.*, 33(2), June 2007.

[GKV10] S. Dov Gordon, Jonathan Katz, and Vinod Vaikuntanathan. A group signature scheme from lattice assumptions. In *ASIACRYPT*, pages 395–412, 2010.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.

[Gt16] Torbjrn Granlund and the GMP development team. *GNU MP: The GNU Multiple Precision Arithmetic Library*, 6.1.2 edition, 2016. http://gmplib.org/.

[Hun12] T.W. Hungerford. *Algebra*. Graduate Texts in Mathematics. Springer New York, 2012.

[KKW18] Jonathan Katz, Vladimir Kolesnikov, and Xiao Wang. Improved non-interactive zero knowledge with applications to post-quantum signatures. *IACR Cryptology ePrint Archive*, 2018:475, 2018.

[KY16] Shuichi Katsumata and Shota Yamada. Partitioning via non-linear polynomial functions: More compact ibes from ideal lattices and bilinear maps. In *ASIACRYPT*, pages 682–712, 2016.

[LLNW16] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for lattice-based accumulators: Logarithmic-size ring signatures and group signatures without trapdoors. In *EUROCRYPT*, pages 1–31, 2016.

[LN86] Rudolph Lidl and Harald Niederreiter. *Introduction to Finite Fields and their Applications*. Cambridge University Press, 1986.

[LN17] Vadim Lyubashevsky and Gregory Neven. One-shot verifiable encryption from lattices. In *EUROCRYPT*, pages 293–323, 2017.

[LNWX18] San Ling, Khoa Nguyen, Huaxiong Wang, and Yanhong Xu. Constant-size group signatures from lattices. In *PKC*, pages 58–88, 2018.

[LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.

[LS18] Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In *EUROCRYPT*, pages 204–224, 2018.

[Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.

[Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.

[NS99] J. Neukirch and N. Schappacher. *Algebraic Number Theory*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1999.

[NZZ15] Phong Q. Nguyen, Jiang Zhang, and Zhenfeng Zhang. Simpler efficient group signatures from lattices. In *PKC*, pages 401–426, 2015.

[PFH+18] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Fast-fourier lattice-based compact signatures over ntru, 2018. submitted to the NIST PQC standardization process.