SMILE: Set Membership from Ideal Lattices with Applications to Ring Signatures and Confidential Transactions

Vadim Lyubashevsky[1], Ngoc Khanh Nguyen[1,2], and Gregor Seiler[1,2]

[1] IBM Research Europe – Zurich, Switzerland
[2] ETH Zurich, Switzerland

Abstract. In a set membership proof, the public information consists of a set of elements and a commitment. The prover then produces a zero-knowledge proof showing that the commitment is indeed to some element from the set. This primitive is closely related to concepts like ring signatures and "one-out-of-many" proofs that underlie many anonymity and privacy protocols. The main result of this work is a new succinct lattice-based set membership proof whose size is logarithmic in the size of the set.

We also give a transformation of our set membership proof to a ring signature scheme. The ring signature size is also logarithmic in the size of the public key set and has size 16 KB for a set of 2^5 elements, and 22 KB for a set of size 2^{25}. At an approximately 128-bit security level, these outputs are between 1.5X and 7X smaller than the current state of the art succinct ring signatures of Beullens et al. (Asiacrypt 2020) and Esgin et al. (CCS 2019).

We then show that our ring signature, combined with a few other techniques and optimizations, can be turned into a fairly efficient Monero-like confidential transaction system based on the MatRiCT framework of Esgin et al. (CCS 2019). With our new techniques, we are able to reduce the transaction proof size by factors of about 4X - 10X over the aforementioned work. For example, a transaction with two inputs and two outputs, where each input is hidden among 2^{15} other accounts, requires approximately 30KB in our protocol.

1 Introduction

Privacy-based transaction systems are steadily gaining in popularity to the point that central banks of the US and the EU are exploring an eventual shift to digital currency. Transaction systems can be equipped with various degrees of privacy, possibilities for auditability, and permission types for joining the transaction network. The common element at the heart of most of these schemes is a zero-knowledge proof which can be adapted to endow the scheme with the desired features. The most efficient zero-knowledge proofs which allow for proving a rich set of statements are generally based on the hardness of the discrete logarithm problem over elliptic curves. This poses a problem for the eventual use of digital

currency because the timeline for widescale deployment of these transaction systems could very well coincide with the advent of a quantum computer that is able to break them. It is therefore important to begin considering schemes which are based on assumptions that are believed to be resistant to quantum attacks.

The currently most efficient, in terms of size and speed, quantum-safe basic primitives are based on the hardness of lattice problems with algebraic structure. Lattice-based constructions are therefore natural candidates for more advanced cryptographic tools like zero-knowledge proofs. Over the last few years, there has indeed been rapid progress in the field of lattice-based zero knowledge (e.g. [2, 7, 12, 25, 6, 11, 13, 1, 10, 18]). There now exist fairly practical protocols for proving knowledge of pre-images of lattice-based 1-way functions, arithmetic sums and products of committed values, as well as various primitives such as ring signatures and group signatures. In virtually all of these cases, the lattice-based solutions result in the most efficient (potentially) quantum-safe option.

As far as a relatively complete quantum-safe transaction system, the recent work of Esgin et al. [13], also based on the hardness of lattice problems, appears to be the most efficient solution. Their work adapts the RingCT protocol [22], which serves as the foundation of the digital currency Monero, and provides formal definitions upon which they construct their MatRiCT protocol. While certainly not as efficient as discrete logarithm based schemes, this work showed that a lattice-based confidential transaction system is something that may eventually be a very reasonable solution.

Our Results and Related Work. At the core of many privacy-based protocols (including the one from [13]) is a set membership proof in which the prover shows, in zero-knowledge, that a commitment is to a value from a public set. This concept is very closely related to "one-out-of-many" proofs [14] and ring signatures [24]. The main result of this work is a new set membership proof which is logarithmic in the size of the set and leads to a ring signature scheme with outputs noticeably smaller than the currently shortest schemes from [13, 4].[3] We point out that "one-out-of many" proofs [14], in which the prover shows that one of the commitments in a set is a commitment to 0, are actually equivalent to the ring signatures that we construct. This is because lattice-based public keys can be thought of as commitments to 0. We then show how to use our ring signature scheme / "one-out-of-many" proof, together with a few other optimizations of prior work, to create a more efficient confidential transaction system based upon the MatRiCT definitions.

We now give a brief overview of where the efficiency advantage comes from. The shorter proofs in our scheme are partly a result of the fact that the modulus in our underlying polynomial ring stays the same for all practical set sizes. On the

[3] One can also obtain ring signatures which are linear (rather than logarithmic) in the size of the public key set by plugging in a lattice-based signature scheme based on a trapdoor function, such as [23], into the generic framework of [24]. Even though for small set sizes (around a dozen), this may be smaller than our solution, it quickly becomes much larger (see Table 2).

other hand, if the size of the set is $n = 32^m$, then the exponent of the modulus in the ring used in [13] increases linearly in m. The reason for this difference is that [13] use "Ajtai-type" commitments which compress the input, but only allow for commitments of "short" messages. In our construction, however, we use BDLOP commitments [3] which allow commitments to arbitrary-size elements, at the expense of a slightly larger commitment size. But because the number of commitments we need is logarithmic in the size of the set, this does not pose a problem with the commitment size becoming too big.

An additional advantage of BDLOP commitments that we extensively use is that if one plans ahead by choosing a long-enough randomness vector in the beginning of the protocol, then one can adjoin a new commitment at any time and the size of the commitment only increases by the size of the committed message. In particular, the increase in size does not depend on the security parameter, which is what one would need if creating a new commitment. We use this property when combining our new techniques along with the framework for proving various relations committed to in BDLOP commitments from [1, 10, 18]. Thus our constructions essentially have just one BDLOP commitment for the entire protocol. We further reduce the transaction size by employing an amortization technique so that the proof contains just two elements whose size depends on the security parameter.

In the rest of the introduction, we give rather detailed high-level descriptions of our constructions. The reason for this level of detail is that the protocols in the body of the paper use optimizations that combine the new ideas together with prior work in a non-black box manner, which tends to somewhat obfuscate the high level picture. In the introduction, we instead give slightly less efficient constructions that try to highlight the separate parts making up the complete protocols. We would then hope that with the high-level intuition in hand, the interested reader can better follow the complete protocols in the body.

1.1 The Polynomial Ring and BDLOP Commitments

Throughout this paper, we will be working over the polynomial ring $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^{128}+1)$ where q is set such that $X^{128}+1 = \prod_{i=1}^{32}(X^4 - r_i)$ and $X^4 - r_i$ are irreducible modulo q (c.f. [21] for how to set q to obtain such a factorization). We will be exclusively using BDLOP commitments [3], where a commitment to a polynomial vector $\vec{\boldsymbol{m}} \in \mathcal{R}_q^k$ is of the form

$$\begin{bmatrix} \boldsymbol{B}_0 \\ \boldsymbol{B}_1 \end{bmatrix} \vec{\boldsymbol{r}} + \begin{bmatrix} \vec{\boldsymbol{0}} \\ \vec{\boldsymbol{m}} \end{bmatrix} = \begin{bmatrix} \vec{\boldsymbol{t}}_0 \\ \vec{\boldsymbol{t}}_1 \end{bmatrix}, \tag{1}$$

where $\boldsymbol{B}_i$ are uniform[4] public random matrices and $\vec{\boldsymbol{r}}$ is a random low-norm vector which serves as the commitment randomness. To open the commitment

[4] For efficiency, a large portion of $\boldsymbol{B}_i$ can be the identity matrix (c.f. [3]), but we ignore the form of the public randomness in this paper, as it does not affect any output sizes.

without revealing it, one would ideally want to give a zero-knowledge proof of a low-norm $\vec{r}$ satisfying $\boldsymbol{B}_0\vec{r} = \vec{t}_0$. Unfortunately, there is no particularly efficient zero-knowledge proof for this statement, and so a relaxed opening is defined which consists of a vector $\vec{v}$ and a polynomial $\boldsymbol{c}$ satisfying $\boldsymbol{B}_0\vec{v} = \vec{t}_0$ such that $\|\boldsymbol{c}\|$ and $\|\boldsymbol{c}\vec{v}\|$ are small (but $\vec{v}$ is not necessarily small itself). The committed message is then implicitly

$$\vec{m} = \vec{t}_1 - \boldsymbol{B}_1\vec{v}. \tag{2}$$

An efficient zero-knowledge proof for the above opening was given in [3]. That work also showed how to prove linear (over $\mathcal{R}_q$) relations of $\vec{m}$ without increasing the proof size. For this, it's in fact enough to just be able to prove that the commitment is to $\vec{0}$. The reason is that a commitment of $\vec{m}$ can be easily converted to a commitment of $\vec{m} + \vec{m}'$ by adding $\vec{m}'$ to $\vec{t}_1$. Similarly, for any matrix $\boldsymbol{L}$ over $\mathcal{R}_q$, one can convert a commitment of $\vec{m}$ to one of $\boldsymbol{L}\vec{m}$ by multiplying the bottom part by $\boldsymbol{L}$ to obtain $\begin{bmatrix} \boldsymbol{B}_0 \\ \boldsymbol{L}\boldsymbol{B}_1 \end{bmatrix} \cdot \vec{r} + \begin{bmatrix} \vec{0} \\ \boldsymbol{L}\vec{m} \end{bmatrix} = \begin{bmatrix} \vec{t}_0 \\ \boldsymbol{L}\vec{t}_1 \end{bmatrix}$. Thus proving that the message $\vec{m}$ in (1) satisfies $\boldsymbol{L}\vec{m} = \vec{u}$, involves proving that the commitment $\begin{bmatrix} \vec{t}_0 \\ \boldsymbol{L}\vec{t}_1 - \vec{u} \end{bmatrix}$ with public key $\begin{bmatrix} \boldsymbol{B}_0 \\ \boldsymbol{L}\boldsymbol{B}_1 \end{bmatrix}$ is a commitment to $\vec{0}$.

Later works (e.g. [1, 10, 18]) showed how to prove more complicated relations between the committed messages in BDLOP commitments. These include proving multiplicative relations among the polynomials comprising $\vec{m}$ and proving linear relations over $\mathbb{Z}_q$ (rather than $\mathcal{R}_q$) of the integer coefficients comprising $\vec{m}$. An important feature of these aforementioned proofs is that the proof size does not grow with the number of relations that one needs to prove about one commitment. So the cost, in terms of proof size, of proving multiple relations about one commitment is the cost of proving the most expensive one.

1.2 The New Set Membership Proof

In this work we extend the toolbox of what can be proved about $\vec{m}$ in BDLOP commitments by showing how to do set membership proofs. Given a collection of polynomial vectors $\vec{p}_i$, and a commitment to one on them, we would like to prove that the committed $\vec{w}$ is indeed one of the $\vec{p}_i$.

More specifically, the public information consists of $\boldsymbol{P} = [\vec{p}_1 \mid \ldots \mid \vec{p}_n]$, where $n = l^m = 32^m$, and a commitment ω. The prover gives a zero knowledge proof that a commitment ω opens to $(\vec{v}_1, \ldots, \vec{v}_m, \vec{w})$ where

$$\boldsymbol{P} \cdot (\vec{v}_1 \otimes \ldots \otimes \vec{v}_m) = \vec{w} \tag{3}$$

$$\forall i, \vec{v}_i \in \{0,1\}^l \text{ and } \|\vec{v}_i\|_1 = 1. \tag{4}$$

Notice that by definition of the $\vec{v}_i$, their tensor product will be a vector of length n consisting of all zeros and one 1 (this decomposition observation was originally used in [14]). If each vector $\vec{v}_i$ will be committed as a polynomial $\boldsymbol{m}_i$ in

the BDLOP commitment,[5] then (4) can already be proved using the aforementioned techniques from [1, 10]. Our main result in this work is an efficient proof of (3) whose size is linear in m, and thus logarithmic in the number of elements in $\boldsymbol{P}$. We also prove a more generic k-dimensional version of this problem. In this version, there are k public lists

$$\boldsymbol{P}^{(1)} = \left[\vec{\boldsymbol{p}}_1^{(1)} \mid \ldots \mid \vec{\boldsymbol{p}}_n^{(1)}\right], \ldots, \boldsymbol{P}^{(k)} = \left[\vec{\boldsymbol{p}}_1^{(k)} \mid \ldots \mid \vec{\boldsymbol{p}}_n^{(k)}\right]$$

and $\vec{w}$ is a sum of k elements, one taken from each set. The prover gives a zero knowledge proof that the commitment ω opens to

$$(\vec{v}_1^{(1)}, \ldots, \vec{v}_m^{(1)}, \ldots, \vec{v}_1^{(k)}, \ldots, \vec{v}_m^{(k)}, \vec{\boldsymbol{w}})$$

where

$$\sum_{j=1}^{k} \boldsymbol{P}^{(j)} \cdot (\vec{v}_1^{(j)} \otimes \ldots \otimes \vec{v}_m^{(j)}) = \vec{\boldsymbol{w}} \tag{5}$$

$$\forall i, j,\ \vec{v}_i^{(j)} \in \{0,1\}^l \text{ and } \|\vec{v}_i^{(j)}\|_1 = 1 \tag{6}$$

This proof is of size $O(mk)$, so there is no amortization happening. But being able to prove the above will allow us to amortize away many of the other parts of the anonymous transaction protocol.

1.3 Set Membership Proof Sketch

We now give a sketch of how to prove (3) and (4). Let us first define the set $\mathcal{M}_q = \mathbb{Z}_q + \mathbb{Z}_q X + \mathbb{Z}_q X^2 + \mathbb{Z}_q X^3$. Because of the way we defined $\mathcal{R}_q$, the NTT and inverse NTT functions are bijective functions $\mathsf{NTT}(\boldsymbol{w}) : \mathcal{R}_q \to \mathcal{M}_q^{32}$ and $\mathsf{NTT}^{-1}(\vec{w}) : \mathcal{M}_q^{32} \to \mathcal{R}_q$ where

$$\mathsf{NTT}(\boldsymbol{w}) = (\boldsymbol{w} \bmod X^4 - r_1, \ldots, \boldsymbol{w} \bmod X^4 - r_{32}).$$

These functions extend to polynomial vectors in the natural way by being applied to each polynomial separately.

We will also need to overload the inner product operator. For a polynomial $\boldsymbol{w}$ such that $\mathsf{NTT}(\boldsymbol{w}) = \vec{w} = (w_1, \ldots, w_{32}) \in \mathcal{M}_q^{32}$, define the function $g(\boldsymbol{w}) = \sum_{i=1}^{32} w_i$. In other words, it's just the sum of the NTT coefficients as polynomials in $\mathcal{M}_q$. For two vectors $\vec{w}, \vec{w}' \in \mathcal{M}_q^{32}$, we define $\langle \vec{w}, \vec{w}' \rangle = g(\mathsf{NTT}^{-1}(w)\,\mathsf{NTT}^{-1}(w'))$. It resembles an inner product because we can equivalently write it as

$$\langle \vec{w}, \vec{w}' \rangle = \sum_{i=1}^{32} w_i w_i' \bmod (X^4 - r_i).$$

[5] Actually the inverse NTT of the vector $\vec{v}_i$, which is an element of $\mathcal{R}_q$, will be committed – see Section 1.3.

The multiplication is performed modulo different polynomials, amd so this function is not an inner product. But it is commutative and satisfies $\langle \vec{w} + \vec{w}', \vec{w}'' \rangle = \langle \vec{w}, \vec{w}'' \rangle + \langle \vec{w}', \vec{w}'' \rangle$. Similarly, for $\vec{w} = (\vec{w}_1, \ldots, \vec{w}_k), \vec{w}' = (\vec{w}'_1, \ldots, \vec{w}'_k)$, where each $\vec{w}_i, \vec{w}'_i \in \mathcal{M}_q^{32}$, one defines $\langle \vec{w}, \vec{w}' \rangle = \sum_{i=1}^{k} \langle \vec{w}_i, \vec{w}'_i \rangle$.

For convenience, we will now rewrite the set membership problem to be over $\mathcal{M}_q$. In particular, the public information consists of vectors $P = [\vec{p}_1 \mid \ldots \mid \vec{p}_n]$ where each $\vec{p}_i \in \mathcal{M}_q^{32k}$, for some arbitrary k. And we also have a commitment to a vector $\vec{w} \in \mathcal{M}_q^{32k}$ such that $\vec{w} = \vec{p}_i$ for some i. Notice that the $\vec{p}_i$ and $\vec{w}$ are the NTT of the $\vec{\boldsymbol{p}}_i, \vec{\boldsymbol{w}}$ from (3). To commit to the vector $\vec{w}$, we define the polynomial vector $\vec{\boldsymbol{w}} = \mathsf{NTT}^{-1}(\vec{w}) \in \mathcal{R}_q^k$ and then use the BDLOP commitment from (1) to commit to $\vec{\boldsymbol{w}}$. Later rows of this BDLOP commitment will also include commitments to the vectors $\vec{v}_1, \ldots, \vec{v}_m \in \mathcal{M}_q^{32}$ (defined as in (4)). We will define the polynomials $\boldsymbol{v}_j = \mathsf{NTT}^{-1}(\vec{v}_j)$ and commit to them in the BDLOP commitment. Note that we can already prove (4) using the techniques from [1, 10] by proving that $\vec{\boldsymbol{v}} \cdot (\vec{\boldsymbol{1}} - \vec{\boldsymbol{v}}) = \vec{\boldsymbol{0}}$ and that the NTT coefficients of each polynomial in $\vec{\boldsymbol{v}}$ sum to 1.

We now describe how to prove (3) – in other words, that $P \cdot (\vec{v}_1 \otimes \ldots \otimes \vec{v}_m) - \vec{w} = \vec{0}$. We will prove this by showing that for a random challenge $\vec{\gamma} \in \mathcal{M}_q^{32k}$, the "inner product" $\langle P \cdot (\vec{v}_1 \otimes \ldots \otimes \vec{v}_m) - \vec{w}, \vec{\gamma} \rangle = 0$. Because $\mathbb{Z}_q[X]/(X^4 - r_i)$ are fields and of size q^4, it's not hard to see that if the left term in the inner product is not $\vec{0}$, then the probability of the inner product being 0 is exactly q^{-4}. Because we will be working with a $q \approx 2^{32}$, this probability is approximately 2^{-128}, so no repetitions are required.

We now get to the main technical part of the protocol. Let's break up P into 32 parts as $P = [\ P_1 \mid \ldots, \mid P_{32}\ ]$ and define $P' := \begin{bmatrix} \gamma^T P_1 \\ \vdots \\ \gamma^T P_{32} \end{bmatrix} \in \mathcal{M}_q^{32 \times 32^{m-1}}$.

Then using the property that $\vec{v}_i$ are vectors over $\mathcal{M}_q$ with just constant coefficients,[6] with some algebraic manipulation (see (18)), it can be shown that

$$\langle P(\vec{v}_1 \otimes \ldots \otimes \vec{v}_m) - \vec{w}, \vec{\gamma} \rangle = \langle \vec{v}_1, P'(\vec{v}_2 \otimes \ldots \otimes \vec{v}_m) \rangle - \langle \vec{w}, \vec{\gamma} \rangle. \tag{7}$$

To prove that the left-hand side is 0, it is therefore equivalent to prove that the right-hand side is 0. The crucial part is that the right-hand side contains an expression which selects one element from a set P' – but this set is 32 times smaller than P. If we define $\vec{x} = P'(\vec{v}_2 \otimes \ldots \otimes \vec{v}_m)$ and send a commitment to $\vec{x}$, then proving the original set membership involves proving a new set membership proof in which the set is 32 times smaller, as well as the equation $\langle \vec{v}_1, \vec{x} \rangle = \langle \vec{w}, \vec{\gamma} \rangle$.

[6] Intuitively, if the coefficients of $\vec{v}_i$ were polynomials of degree > 0, then the term $\langle \vec{v}_1, P'(\vec{v}_2 \otimes \ldots \otimes \vec{v}_m) \rangle$ in (7) would make very little algebraic sense because there is a multiplication on one side of P' which involves reduction modulo $X^4 - r_j$, and then there would be a multiplication on the other side which would get reduced modulo different $X^4 - r_{j'}$. But since vectors $\vec{v}_i$ only have constant terms, the "inner product" with $\vec{v}_i$ does not involve any modular reduction.

If this latter equation can be proved with a constant number of commitments (in our case, it will essentially be one), then continuing the proof recursively would mean that the whole proof requires approximately $2m$ commitments for sets P containing $n = 32^m$ elements.

Both $\vec{w}$ and $\vec{\gamma}$ are vectors in $\mathcal{M}_q^{32k}$, so let us write them as $\vec{w} = (\vec{w}_1, \ldots, \vec{w}_k)$ and $\vec{\gamma} = (\vec{\gamma}_1, \ldots, \vec{\gamma}_k)$ where $\vec{w}_i, \vec{\gamma}_i \in \mathcal{M}_q^{32}$. Then

$$\langle \vec{v}_1, \vec{x} \rangle = \langle \vec{w}, \vec{\gamma} \rangle \Leftrightarrow g(\boldsymbol{v}_1 \boldsymbol{x}) = g\left(\sum_{i=1}^{k} \boldsymbol{w}_i \boldsymbol{\gamma}_i \right),$$

where the bold letters correspond to the inverse NTTs and the function g is the sum of the NTT's of the polynomial. Because we have BDLOP commitments to $\boldsymbol{x}$ and $\boldsymbol{w}_i$, we can compute a commitment to $\boldsymbol{y} = \boldsymbol{v}_1 \boldsymbol{x} - \sum_{i=1}^{k} \boldsymbol{w}_i \boldsymbol{\gamma}_i$, and then we just have to prove that the sum of the NTT coefficients of this polynomial is 0. For this, we employ a lemma used in [10], which states that for the ring $\mathcal{R}_q$ as defined in this section and a polynomial $\boldsymbol{y} \in \mathcal{R}_q = \sum_{i=0}^{127} y_i X^i$, we have $g(\boldsymbol{y}) = 32(y_0 + y_1 X + y_2 X^2 + y_3 X^3)$. In other words, the sum of the NTT coefficients is 0 if and only if the first four coefficients of the polynomial representation are 0. To prove this in zero knowledge, we can first commit to a masking polynomial $\boldsymbol{z}$ whose first 4 coefficients are 0 and the rest uniform in $\mathbb{Z}_q$, and then output $\boldsymbol{y} + \boldsymbol{z}$ and prove that this is indeed the right sum. The verifier can then check that the first four coefficients are 0. We don't need to multiply $\boldsymbol{y}$ by a challenge because in our case, it already contains a challenge $\vec{\gamma}$. In the body of the paper, we present an efficient way to do this proof which does not require committing to $\boldsymbol{y}$ and so we just need an extra commitment to $\vec{x} \in \mathcal{M}_q^{32}$ at each level of the recursion.

1.4 From Set Membership to Ring Signatures

A ring signature scheme allows a signer to sign in a way that hides the public key that he is using. More specifically, the signer creates a set comprised of his public key and other public keys for which he may not know the secret key. He then creates a signature with the property that the verifier can check that the message was signed by an entity who knows the secret key to one of the public keys in the list. We now sketch how one can convert a "Schnorr-like" lattice-based signature scheme into a ring signature by using a set membership proof.

The basic signature scheme underlying the ring signature follows the usual "Fiat-Shamir with Aborts" approach for constructing lattice-based digital signatures (e.g. [16, 17, 9]). In particular, the secret key is a low-norm vector $\vec{\boldsymbol{s}}$, while the public key consists of a random matrix $\boldsymbol{A}$ and a vector $\vec{\boldsymbol{t}} = \boldsymbol{A}\vec{\boldsymbol{s}}$. The signature is then a "relaxed" zero-knowledge proof of knowledge (made non-interactive using the Fiat-Shamir transform) of a vector $\vec{\boldsymbol{s}}'$ and a polynomial $\boldsymbol{c}'$, both with small norms, satisfying $\boldsymbol{c}'\vec{\boldsymbol{t}} = \boldsymbol{A}\vec{\boldsymbol{s}}'$.

The ring signature public information consists of the matrix $\boldsymbol{A}$ and vectors $\vec{\boldsymbol{t}}_1, \ldots, \vec{\boldsymbol{t}}_n$. A signer who knows an $\vec{\boldsymbol{s}}_i$ satisfying $\boldsymbol{A}\vec{\boldsymbol{s}}_i = \vec{\boldsymbol{t}}_i$ will want to give a zero-knowledge proof knowledge of $\vec{\boldsymbol{s}}', \boldsymbol{c}'$, and $i \in [0, n)$ satisfying $\boldsymbol{c}'\vec{\boldsymbol{t}}_i = \boldsymbol{A}\vec{\boldsymbol{s}}'$. An interactive version of this proof is presented in Figure 1 and it is then made non-interactive using the Fiat-Shamir transform and inserting the message to be signed into the random oracle which is used to produce the challenge.

Private information: $\vec{v}_1, \ldots, \vec{v}_m \in \{0,1\}^l$ as in (4), and $\vec{\boldsymbol{s}}$ with a small norm
Public information: $\boldsymbol{A}$, $\boldsymbol{T} = [\ \vec{\boldsymbol{t}}_1 \mid \ \ldots \ \mid \vec{\boldsymbol{t}}_n]$, where $n = l^m$ s.t. $\boldsymbol{T} \cdot (\vec{v}_1 \otimes \ldots \otimes \vec{v}_m) = \boldsymbol{A}\vec{\boldsymbol{s}}$

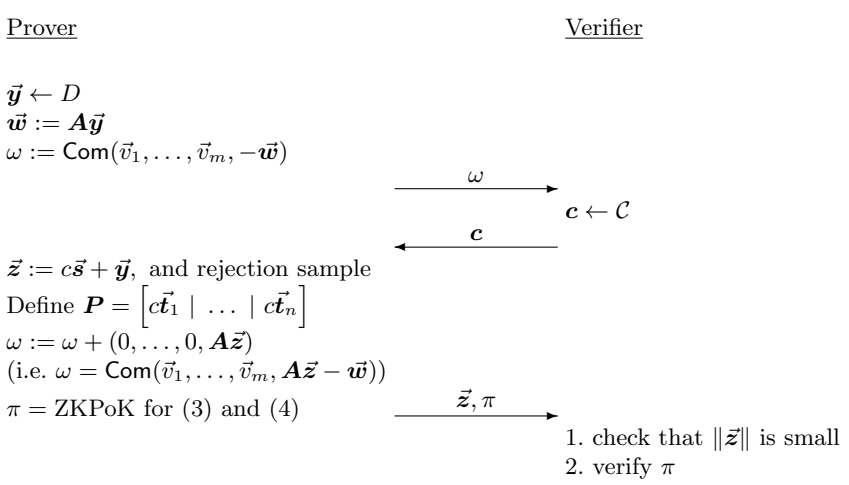

Fig. 1. A lattice-based ring signature using the set membership proof. Com is a BDLOP commitment, while D is a distribution that outputs polynomial vectors with small coefficients. As in Section 1.3, a BDLOP commitment to $\vec{v}_i$ is a commitment to the polynomial $\mathsf{NTT}^{-1}(\vec{v}_i) \in \mathcal{R}_q$.

To see that this proof is complete (assuming that all the norm-checks pass), notice that $\boldsymbol{A}\vec{\boldsymbol{z}} - \boldsymbol{c}\vec{\boldsymbol{t}}_i = \boldsymbol{A}\vec{\boldsymbol{y}} = \vec{\boldsymbol{w}}$. And this is exactly what π proves. The zero-knowledge property follows from the fact that π is a zero-knowledge proof and that $\vec{\boldsymbol{z}}$ is independent of $\vec{\boldsymbol{s}}$ and $\boldsymbol{c}$ due to the employed rejection sampling. To see that the protocol is a proof of knowledge, note that verifying π implies that $\boldsymbol{A}\vec{\boldsymbol{z}} - c\vec{\boldsymbol{t}}_i = \vec{\boldsymbol{w}}$. Because the $\vec{v}_i$ and $\vec{\boldsymbol{w}}$ in the commitment are fixed, if we rewind the prover with a different challenge $\boldsymbol{c}'$, we will obtain $\boldsymbol{A}\vec{\boldsymbol{z}}' - \boldsymbol{c}'\vec{\boldsymbol{t}}_i = \vec{\boldsymbol{w}}$. Eliminating $\vec{\boldsymbol{w}}$ by subtracting the two equations results in the statement that we would like to extract.

1.5 Bimodal Gaussians (almost) for Free

The goal of the rejection sampling in the signing algorithm is to remove the dependence of the secret key $\vec{\boldsymbol{s}}$ from the output $\vec{\boldsymbol{z}}$. If the distribution D in

Ring Size	2^3	2^5	2^6	2^{10}	2^{12}	2^{15}	2^{21}	2^{25}
Falafl [4]	30		32		35		39	
Esgin et al. [13]	19		31		59		148	
Raptor [15] / [24]+[23]	10		81		5161			
This Work		**16**		**18**		**19**		**22**

Fig. 2. Sizes, in KB, of the different lattice-based ring signature schemes with approximately 128 bits of security. The sizes for [4, 13, 15] are taken from [4, Table 1].

Figure 1 is a zero-centered discrete Gaussian, then the distribution of $\vec{\boldsymbol{z}} = \boldsymbol{c}\vec{\boldsymbol{s}} + \vec{\boldsymbol{y}}$ before rejection sampling is performed is a discrete Gaussian centered at $\boldsymbol{c}\vec{\boldsymbol{s}}$. In order for the rejection probability to not be too large (e.g. $< 1 - 1/e$), one needs the standard deviation of the $\vec{\boldsymbol{z}}$ after the rejection sampling to be around $12 \cdot \|\boldsymbol{c}\vec{\boldsymbol{s}}\|$ [17]. In [8], it was shown that if one can get the distribution of $\vec{\boldsymbol{z}}$ before rejection sampling to follow a *bimodal* Gaussian distribution with the two centers being $\pm\boldsymbol{c}\vec{\boldsymbol{s}}$, then one only needs the standard deviation of the $\vec{\boldsymbol{z}}$ after rejection sampling to be $\|\boldsymbol{c}\vec{\boldsymbol{s}}\|/\sqrt{2}$ for the same repetition rate. Such a reduction has a direct consequence on reducing the output length and increasing the SIS-hardness of the underlying problem.

The way to create a bimodal gaussian with the two centers being $\pm\boldsymbol{c}\vec{\boldsymbol{s}}$ is for the prover to choose a $y \leftarrow D$ and also a $b \leftarrow \{-1, 1\}$ and then create $\vec{\boldsymbol{z}} = b\boldsymbol{c}\vec{\boldsymbol{s}} + \vec{\boldsymbol{y}}$. It is crucial for security that b remains hidden and so the verifier is not allowed to know b or use it during verification. This could be an issue in regular signature schemes because the verifier would need to directly check that

$$\boldsymbol{A}\vec{\boldsymbol{z}} = \boldsymbol{c}\vec{\boldsymbol{t}} + \vec{\boldsymbol{w}}. \tag{8}$$

Since $\boldsymbol{A}\vec{\boldsymbol{z}} = \boldsymbol{A}(b\boldsymbol{c}\vec{\boldsymbol{s}} + \vec{\boldsymbol{y}})$, we would need $\boldsymbol{A}\vec{\boldsymbol{s}} = -\boldsymbol{A}\vec{\boldsymbol{s}}$ to always hold. In our case, this does not hold, but it will not pose a problem because the verifier does not directly verify (8) because, for privacy, the prover cannot send $\vec{\boldsymbol{w}}$ in the clear anyway. Instead, the verifier gets $\mathsf{Com}(\vec{\boldsymbol{w}})$ and a ZK proof that this commitment opens to a $\vec{\boldsymbol{w}}$ satisfying (8). Since the prover already sends a commitment to $\vec{\boldsymbol{w}}$ along with the ones for $\vec{v}_i$ (and eventually all the "garbage terms" required in π), he can just increase the commitment size by one (128-degree) polynomial and also commit to b. Then the proof π would need to be modified to prove that

$$[b\boldsymbol{c}\vec{\boldsymbol{t}}_1 \mid \ldots \mid b\boldsymbol{c}\vec{\boldsymbol{t}}_n] \cdot (\vec{v}_1 \otimes \ldots \otimes \vec{v}_m) = \vec{\boldsymbol{w}} - \boldsymbol{A}\vec{\boldsymbol{z}}.$$

Notice that because $b \in \{-1, 1\}$ and all the $\vec{v}_i$ consist of all 0's and one 1, this can be rewritten as

$$[\boldsymbol{c}\vec{\boldsymbol{t}}_1 \mid \ldots \mid \boldsymbol{c}\vec{\boldsymbol{t}}_n] \cdot (b\vec{v}_1 \otimes \vec{v}_2 \otimes \ldots \otimes \vec{v}_m) = \vec{\boldsymbol{w}} - \boldsymbol{A}\vec{\boldsymbol{z}},$$

and so the only thing that changes is that instead of committing to $\vec{v}_1$, the prover commits to $b\vec{v}_1$. He then just has to show that the coefficients of $b\vec{v}_1$ are in $\{0, b\}$ rather than $\{0, 1\}$ – but this proof is exactly the same if we already have a commitment to b (which we proved to be in $\{-1, 1\}$).

1.6 Application to Confidential Transactions

We now show how to construct a confidential transaction system in the model of [13]. The setup is the following: at any given moment, the state (which is managed by the blockchain, and is outside the scope of this work) consists of a set of accounts $\mathsf{act} = (\mathsf{pk}, \mathsf{cn})$, each of which contains a public key and a coin. The state also contains a set of serial numbers which implicitly correspond to the accounts that were already spent (to prevent double-spending). The secret account key associated to each account is $\mathsf{ask} = (\mathsf{sk}, \mathsf{ck}, \mathsf{amt})$, which consists of the secret key corresponding to pk and the commitment key ck, which is the randomness used to create the BDLOP commitment cn to the amount amt in the account. As in [13], we will assume that amt takes values between 0 and $2^{64} - 1$. Since we are working over rings with 32 NTT slots, we will represent the values in base 4. The basic operation has the sender choosing M input accounts for which he knows the secret keys associated to $\mathsf{pk}^{(1)}, \ldots, \mathsf{pk}^{(M)}$, and then creating S new output accounts with given public keys for which he does not need to know the associated secret keys. There are three correctness constraints. The first is that the spender knows the associated secret keys for the M input accounts. The second is that the sum of the values of the input coins (i.e. the sum of the amt) equals to the sum of the values of the output coins. And the third is that none of the M input accounts were used as inputs in any previous transaction.

In addition to correctness, there are also secrecy and anonymity requirements. The secrecy requirement states that nothing about the amounts amt is known except that the sum of the input and output coins is equal. The spender's anonymity is defined by hiding the spenders account among N other accounts. In particular, rather than stating which M accounts the spender is using, he will instead choose M sets of N accounts each, and then choose one account from each set in a way that hides which of the N accounts has been chosen. How the spender chooses the $N - 1$ other accounts is a policy issue that is outside the scope of this work.

The public information for the system consists of a polynomial matrix $\boldsymbol{B}$ which forms the "top part" of the BDLOP commitment. The polynomial vectors $\vec{\boldsymbol{b}}_c$ (which will be used to commit to amt) and $\vec{\boldsymbol{b}}_s$ (which will be used to "commit" to zero, with the commitment being the serial number) form the "bottom part" of the commitments. In particular, sk is a low-norm vector $\vec{\boldsymbol{s}}$ where

$$\begin{bmatrix} \boldsymbol{B} \\ \vec{\boldsymbol{b}}_s \end{bmatrix} \vec{\boldsymbol{s}} = \begin{bmatrix} \mathsf{pk} \\ \mathsf{sn} \end{bmatrix}. \tag{9}$$

And ck is another low-norm vector $\vec{\boldsymbol{r}}$ such that

$$\begin{bmatrix} \boldsymbol{B} \\ \vec{\boldsymbol{b}}_c \end{bmatrix} \vec{\boldsymbol{r}} + \begin{bmatrix} 0 \\ \mathsf{amt} \end{bmatrix} = \mathsf{cn}. \tag{10}$$

Correctness. Let's ignore anonymity for a moment, and just briefly discuss how the correctness of the protocol could be handled. If the spender wants to spend

accounts $\mathsf{act}^{(1)}, \ldots, \mathsf{act}^{(M)}$, then he outputs the values $\mathsf{sn}^{(j)}, \vec{\boldsymbol{s}}^{(j)}, \vec{\boldsymbol{r}}^{(j)}, \mathsf{amt}^{(j)}$ for the input accounts, and the verifier can check that (9) and (10) are satisfied. Furthermore, the verifier checks that none of the $\mathsf{sn}^{(j)}$ are in the set of used serial numbers, and adds these $\mathsf{sn}^{(j)}$ to the set. Note that because the value of $\vec{\boldsymbol{s}}^{(j)}$ is uniquely determined by $\boldsymbol{B}$ and pk (unless SIS is easy), the value of sn is uniquely tied to pk; and so it is not possible to spend a coin more than once. The spender then creates valid output tokens with the values of pk that he is given and creates the output coins with by picking small vectors $\vec{\boldsymbol{r}}$ and using them to create BDLOP commitments to amt as in (10). He then outputs these $\vec{\boldsymbol{r}}$ and amt so that everyone can check that the sum of the input amounts is equal to the sum of the output amounts.

Anonymity and Secrecy. We now sketch how anonymity and secrecy is achieved in our confidential transactions protocol. The spender chooses the M accounts $\mathsf{act}^{(j)} = (\mathsf{pk}^{(j)}, \mathsf{cn}^{(j)})$ that he wants to spend. He puts each of the right hand sides of (10) (i.e. the coin commitments) from these accounts into M lists $\boldsymbol{T}^{(j)}$, one coin per list. The rest of the lists are filled with N coins from accounts among which the spender wants to hide his. He then creates S output accounts $\mathsf{act}^{(j)} = (\mathsf{pk}^{(j)}, \mathsf{cn}^{(j)})$ using the given public keys. He does not need to hide these accounts and so he just creates S lists of size 1 for the output coins. He then wants to create one BDLOP commitment that includes all the coin values (i.e. the amt) from the input and output tokens. This protocol is described in Figure 6. Once the spender has one BDLOP commitment, he can prove that the sum of the input and output tokens matches, which can be done using techniques similar to those in [13, 18].

The prover also needs to show that he knows $\vec{\boldsymbol{s}}$ that satisfy (9) for the input accounts. He does this by creating M lists $\boldsymbol{U}^{(j)}$ that are derived from $\boldsymbol{T}^{(j)}$. If the spender's coin is in position i in the list $\boldsymbol{T}^{(j)}$, then he puts $\begin{bmatrix} \mathsf{pk}_i^{(j)} \\ \mathsf{sn}^{(j)} \end{bmatrix}$ into position i. He then fills the list with the public keys from the accounts corresponding to the coins in $\boldsymbol{T}^{(j)}$. For the serial numbers, he attaches the same one (i.e. the one corresponding to his public key) to all the public keys. In particular, if the spender wants to hide the j^{th} account that he will be using in position i among $N-1$ other accounts $\mathsf{act}_1, \ldots, \mathsf{act}_{i-1}, \mathsf{act}_{i+1}, \ldots, \mathsf{act}_N$, then the lists $\boldsymbol{T}^{(j)}$ and $\boldsymbol{U}^{(j)}$ are

$$\boldsymbol{T}^{(j)} = \left[\mathsf{cn}_1^{(j)}, \ldots, \mathsf{cn}_N^{(j)}\right]$$

$$\boldsymbol{U}^{(j)} = \left[\begin{bmatrix} \mathsf{pk}_1^{(j)} \\ \mathsf{sn}^{(j)} \end{bmatrix}, \cdots, \begin{bmatrix} \mathsf{pk}_{i-1}^{(j)} \\ \mathsf{sn}^{(j)} \end{bmatrix}, \begin{bmatrix} \mathsf{pk}_i^{(j)} \\ \mathsf{sn}^{(j)} \end{bmatrix}, \begin{bmatrix} \mathsf{pk}_{i+1}^{(j)} \\ \mathsf{sn}^{(j)} \end{bmatrix}, \ldots, \begin{bmatrix} \mathsf{pk}_N^{(j)} \\ \mathsf{sn}^{(j)} \end{bmatrix}\right]$$

For the lists $\boldsymbol{U}^{(j)}$, the spender simply wants to prove that he knows the secret keys $\vec{\boldsymbol{s}}^{(j)}$ for the elements in the same position as those in $\boldsymbol{T}^{(j)}$. Since the positions are already committed to, the proof of knowledge of the $\vec{\boldsymbol{s}}^{(j)}$ does not require any extra BDLOP commitments and the proof of knowledge of the $\vec{\boldsymbol{s}}^{(j)}$ can be amortized into the output vector $\vec{\boldsymbol{z}}$ in Figure 6. The verifier will need to check that the serial numbers $\mathsf{sn}^{(j)}$ have never been used (i.e. don't appear in

(M, S)	ring size N 2^5	2^{10}	2^{15}	2^{20}	2^{25}
$(1, 2)$ This Work	22 KB	24 KB	25 KB	27 KB	28 KB
$(1, 2)$ Esgin et al. [13]	100 KB	160 KB	250 KB	375 KB	520 KB
$(2, 2)$ This Work	24 KB	27 KB	30 KB	33 KB	36 KB
$(2, 2)$ Esgin et al. [13]	110 KB	190 KB	300 KB	440 KB	660 KB

Fig. 3. Transaction proof sizes depending on ring size (anonymity set size) N, number M of input accounts, and number S of output accounts. The sizes for [13] are taken from [13, Figure 1].

M	25	50	75	100
size (This Work $N = 1024$)	100 KB	180 KB	262 KB	345 KB
size (Esgin et al. [13] $N = 100$)	370 KB	610 KB	900 KB	1170 KB

Fig. 4. Transaction proof sizes with M input accounts and $S = 2$ output accounts. The anonymity set N is 100 in [13] and $32^2 = 1024$ in our work. The sizes for [13] are taken from [13, Figure 2].

the "used" pile) and that the lists $\boldsymbol{T}^{(j)}, \boldsymbol{m}^{(j)}$ are valid (i.e. the positions $\mathsf{pk}_i^{(j)}$ in list $\boldsymbol{T}^{(j)}$ and $\mathsf{cn}_i^{(j)}$ in list $\boldsymbol{U}^{(j)}$ correspond to some account $\mathsf{act} = \left(\mathsf{pk}_i^{(j)}, \mathsf{cn}_i^{(j)}\right)$). The verifier also has to verify the proof from Figure 6 and the addition proof confirming that the amounts in the input and output accounts match.

The protocol in Figure 6, which is at the center of the confidential transaction protocol, creates a new BDLOP commitment and proves that it is committing to the same values as the M input and S output accounts. It additionally proves that the spender knows the secret keys of the M input accounts. This involves using the protocol for the k-dimensional version of the set membership problem as well as an amortization technique which will allow us to only send one "masked value" for all the randomness used in the $M+S$ accounts.

Aggregating BDLOP Commitments. Before describing the protocol in Figure 6, we ignore the part where each of the M input accounts are hidden among N others, and give a simpler protocol in Figure 5 that takes k BDLOP commitments with distinct randomnesses, and creates one BDLOP commitment to the same messages. The improvement in this protocol over the trivial one is in the fact that only one output $\vec{\boldsymbol{z}}$ is enough to prove knowledge that all k commitments are valid. The norm of this vector $\vec{\boldsymbol{z}}$ is larger by a factor of k (or $\sqrt{k}$ in the asymptote), so its representation grows only logarithmically in k.

The protocol in Figure 5 takes as input k BDLOP commitments under randomness $\vec{\boldsymbol{s}}_i$ and produces one BDLOP commitment ω under randomness $\vec{\boldsymbol{r}}$. The commitment includes all the $\boldsymbol{m}_i$ and one additional "garbage polynomial" $\tilde{\boldsymbol{w}}$. When the prover computes and outputs $\vec{\boldsymbol{z}}$, he proves that all the k commitments

Private information: For $1 \leq i \leq k$, polynomials $\boldsymbol{m}_i$, low-norm vectors $\vec{\boldsymbol{s}}_i$
Public information: Uniformly random $\boldsymbol{B}, \vec{\boldsymbol{b}}, \boldsymbol{A}, \vec{\boldsymbol{a}}_w, \vec{\boldsymbol{a}}_i, \begin{bmatrix} \vec{\boldsymbol{t}}_i \\ \boldsymbol{u}_i \end{bmatrix} = \begin{bmatrix} \boldsymbol{B} \\ \vec{\boldsymbol{b}} \end{bmatrix} \vec{\boldsymbol{s}}_i + \begin{bmatrix} \vec{\boldsymbol{0}} \\ \boldsymbol{m}_i \end{bmatrix}$

Prover		Verifier
$(\vec{\boldsymbol{y}}, \vec{\boldsymbol{r}}) \leftarrow D_y \times D_r$ $\vec{\boldsymbol{w}} := \boldsymbol{B}\vec{\boldsymbol{y}};\ \tilde{\boldsymbol{w}} := \vec{\boldsymbol{b}} \cdot \vec{\boldsymbol{y}}$ $\begin{bmatrix} \boldsymbol{A} \\ \vec{\boldsymbol{a}}_1 \\ \dots \\ \vec{\boldsymbol{a}}_k \\ \vec{\boldsymbol{a}}_w \end{bmatrix} \vec{\boldsymbol{r}} + \begin{bmatrix} \vec{\boldsymbol{0}} \\ \boldsymbol{m}_1 \\ \dots \\ \boldsymbol{m}_k \\ \tilde{\boldsymbol{w}} \end{bmatrix} = \begin{bmatrix} \vec{\boldsymbol{f}} \\ \boldsymbol{g}_1 \\ \dots \\ \boldsymbol{g}_k \\ \tilde{\boldsymbol{g}}_w \end{bmatrix} = \omega$	$\xrightarrow{\vec{\boldsymbol{w}}, \omega}$	
		$\boldsymbol{c}_1, \dots, \boldsymbol{c}_k \leftarrow \mathcal{C}$
	$\xleftarrow{\boldsymbol{c}_1, \dots, \boldsymbol{c}_k}$	
$\vec{\boldsymbol{z}} := \vec{\boldsymbol{y}} + \sum \boldsymbol{c}_i \vec{\boldsymbol{s}}_i$, and rejection sample $\vec{\boldsymbol{a}}^* := \sum_{i=1}^{k} \boldsymbol{c}_i \vec{\boldsymbol{a}}_i - \vec{\boldsymbol{a}}_w$ $\boldsymbol{g}^* := \sum_{i=1}^{k} \boldsymbol{c}_i \boldsymbol{g}_i - \vec{\boldsymbol{g}}_w + \vec{\boldsymbol{b}} \cdot \vec{\boldsymbol{z}}$ $\pi =$ ZKPoK that $\begin{bmatrix} \vec{\boldsymbol{f}} \\ \boldsymbol{g}^* \end{bmatrix}$ under public key $\begin{bmatrix} \boldsymbol{A} \\ \vec{\boldsymbol{a}}^* \end{bmatrix}$ is a commitment to $\sum_{i=1}^{k} \boldsymbol{c}_i \boldsymbol{u}_i$	$\xrightarrow{\vec{\boldsymbol{z}}, \pi}$	
		1. check that $\|\vec{\boldsymbol{z}}\|$ is small 2. check that $\sum_{i=1}^{k} \boldsymbol{c}_i \vec{\boldsymbol{t}}_i = \boldsymbol{B}\vec{\boldsymbol{z}} - \vec{\boldsymbol{w}}$ 3. Compute $\vec{\boldsymbol{a}}^*, \boldsymbol{g}^*$ and verify π

Fig. 5. A protocol which takes commitments $\begin{bmatrix} \vec{\boldsymbol{t}}_i \\ \boldsymbol{u}_i \end{bmatrix} = \begin{bmatrix} \boldsymbol{B} \\ \vec{\boldsymbol{b}} \end{bmatrix} \vec{\boldsymbol{s}}_i + \begin{bmatrix} \vec{\boldsymbol{0}} \\ \boldsymbol{m}_i \end{bmatrix}$ to $\boldsymbol{m}_i$ under distinct randomnesses $\vec{\boldsymbol{s}}_i$, and outputs one BDLOP commitment ω to all the $\boldsymbol{m}_i$ (and some auxiliary garage term(s)) under one common randomness $\vec{\boldsymbol{r}}$. Along with outputting the commitment, the protocol also proves that $\begin{bmatrix} \vec{\boldsymbol{t}}_i \\ \boldsymbol{u}_i \end{bmatrix}$ are valid commitments and that the new commitment is to the same $\boldsymbol{m}_i$.

under $\vec{s}_i$ are valid. The rest of the steps are needed to show that the commitment under $\vec{r}$ is to the same $\boldsymbol{m}_i$. We discuss this in more detail below.

The proof that the k commitments are valid follows from the ideas in [2] where one does rewinding by keeping most of the challenge fixed. As long as the new challenge still has κ bits of entropy conditioned on the prior challenge, the soundness error will still be $\approx 2^{-\kappa}$. Without loss of generality, suppose that we would like to prove that the new commitment is a commitment to $\boldsymbol{m}_1$ (in the row that contains $\boldsymbol{g}_1$). Let $(\vec{\boldsymbol{w}}, \omega, \boldsymbol{c}_1, \boldsymbol{c}_2 \ldots, \boldsymbol{c}_k, \vec{\boldsymbol{z}}, \pi)$ be the transcript of one run and $(\vec{\boldsymbol{w}}, \omega, \boldsymbol{c}'_1, \boldsymbol{c}_2 \ldots, \boldsymbol{c}_k, \vec{\boldsymbol{z}}', \pi')$ be the view of the second run when we rewind while keeping all the challenges, except for $\boldsymbol{c}_1$ fixed.

Rewinding on the second verification equation, we obtain $(\boldsymbol{c}_1 - \boldsymbol{c}'_1)\vec{\boldsymbol{t}}_1 = \boldsymbol{A}(\vec{\boldsymbol{z}} - \vec{\boldsymbol{z}}')$. By (2), this implies that the message $\boldsymbol{m}_i$ committed to by $\begin{bmatrix} \vec{\boldsymbol{t}}_1 \\ \boldsymbol{u}_1 \end{bmatrix}$ satisfies

$$(\boldsymbol{c}_1 - \boldsymbol{c}'_1)\boldsymbol{m}_1 = (\boldsymbol{c} - \boldsymbol{c}')\boldsymbol{u}_1 - \vec{\boldsymbol{b}} \cdot (\vec{\boldsymbol{z}} - \vec{\boldsymbol{z}}'). \tag{11}$$

Notice that repeating this for all i, we can prove that all the commitments $\begin{bmatrix} \vec{\boldsymbol{t}}_i \\ \boldsymbol{u}_i \end{bmatrix}$ are valid. The intuition for proving that ω is a commitment to the same messages is to prove that the messages in the commitment of ω (call them $\bar{\boldsymbol{m}}_i$ and $\bar{\boldsymbol{w}}$) satisfy the linear equation

$$\sum_i \boldsymbol{c}_i \bar{\boldsymbol{m}}_i = \sum_i \boldsymbol{c}_i \boldsymbol{u}_i + \bar{\boldsymbol{w}} - \vec{\boldsymbol{b}} \cdot \vec{\boldsymbol{z}}. \tag{12}$$

Rewinding in the same way as above, we would obtain

$$(\boldsymbol{c}_1 - \boldsymbol{c}'_1)\bar{\boldsymbol{m}}_1 = (\boldsymbol{c}_1 - \boldsymbol{c}'_1)\boldsymbol{u}_1 - \vec{\boldsymbol{b}} \cdot (\vec{\boldsymbol{z}} - \vec{\boldsymbol{z}}').$$

Substituting $\vec{\boldsymbol{b}} \cdot (\vec{\boldsymbol{z}} - \vec{\boldsymbol{z}}')$ from (11), we get $(\boldsymbol{c}_1 - \boldsymbol{c}'_1)\bar{\boldsymbol{m}}_1 = (\boldsymbol{c}_1 - \boldsymbol{c}'_1)\boldsymbol{m}_1$. And since $\boldsymbol{c}_1 - \boldsymbol{c}'_1$ is invertible, we have $\boldsymbol{m}_1 = \bar{\boldsymbol{m}}_1$ as desired.

We now observe that we exactly prove (12). The proof π proves that ω is a valid commitment and therefore there is a unique $\vec{\boldsymbol{v}}$ (and a short polynomial d s.t. $d\vec{\boldsymbol{v}}$ has small norm) satisfying $\boldsymbol{g}_i - \vec{\boldsymbol{a}}_i \cdot \vec{\boldsymbol{v}} = \bar{\boldsymbol{m}}_i$ and $\boldsymbol{g}_w - \vec{\boldsymbol{a}}_w \cdot \vec{\boldsymbol{v}} = \bar{w}$. Because we also prove that $\sum \boldsymbol{c}_i \boldsymbol{u}_i$ is a valid commitment, it implies that $\langle \vec{\boldsymbol{a}}^*, \vec{\boldsymbol{v}} \rangle + \sum \boldsymbol{c}_i \boldsymbol{v}_i = \boldsymbol{g}^*$. If we expand out the definitions of $\vec{\boldsymbol{a}}^*$ and $\boldsymbol{g}^*$, and then plug it in, along with the expressions for $(\boldsymbol{c}_i - \boldsymbol{c}'_i)\boldsymbol{g}_i$ and $(\boldsymbol{c}_i - \boldsymbol{c}'_i)\boldsymbol{g}_w$, into the previous equation, we will exactly end up with (12).

We now sketch the zero-knowledge proof. By assumption, π can be simulated and $\vec{\boldsymbol{z}}$ is independent of $\vec{\boldsymbol{s}}_i$ and $\boldsymbol{c}_i$ by rejection sampling. The BDLOP commitment ω is indistinguishable from uniform by the LWE assumption, and $\vec{\boldsymbol{w}}$ is unique once $\vec{\boldsymbol{z}}$ and $\boldsymbol{c}_i$ are chosen. Something worth noting is that while $\vec{\boldsymbol{w}} = \boldsymbol{B}\vec{\boldsymbol{y}}$ can be sent in the clear, the value $\tilde{\boldsymbol{w}} = \vec{\boldsymbol{b}} \cdot \vec{\boldsymbol{y}}$ needs to be sent as part of a commitment because revealing it in the clear would end up revealing some function of the $\boldsymbol{m}_i$.

Private information: For $1 \leq j \leq k$, $V^{(j)} = (\vec{v}_1^{(j)}, \ldots, \vec{v}_m^{(j)}) \in \{0,1\}^{l \times m}$ s.t. $\|\vec{v}_i^{(j)}\|_1 = 1$, $\vec{\boldsymbol{s}}^{(j)}$ with a small norm, and message polynomials $\boldsymbol{m}^{(j)}$

Public information: $\boldsymbol{B}$, $\vec{\boldsymbol{b}}$, $\boldsymbol{T}^{(j)} = \left[\begin{bmatrix} \vec{\boldsymbol{t}}_1^{(j)} \\ \boldsymbol{u}_1^{(j)} \end{bmatrix} \mid \ldots \mid \begin{bmatrix} \vec{\boldsymbol{t}}_n^{(j)} \\ \boldsymbol{u}_n^{(j)} \end{bmatrix} \right]$, where $n = l^m$, s.t

$\boldsymbol{T}^{(j)} \cdot (\vec{v}_1^{(j)} \otimes \ldots \otimes \vec{v}_m^{(j)}) = \begin{bmatrix} \boldsymbol{B} \\ \vec{\boldsymbol{b}} \end{bmatrix} \vec{\boldsymbol{s}}^{(j)} + \begin{bmatrix} \vec{\boldsymbol{0}} \\ \boldsymbol{m}^{(j)} \end{bmatrix}$

Prover		Verifier
$\vec{\boldsymbol{y}} \leftarrow D$ $\vec{\boldsymbol{w}} := \boldsymbol{B}\vec{\boldsymbol{y}}$ $\tilde{\boldsymbol{w}} := \vec{\boldsymbol{b}} \cdot \vec{\boldsymbol{y}}$ $\omega = \mathsf{Com}(\boldsymbol{m}^{(1)}, \ldots, \boldsymbol{m}^{(k)}, V^{(1)}, \ldots, V^{(k)}, \tilde{\boldsymbol{w}}, -\vec{\boldsymbol{w}})$	$\xrightarrow{\omega}$	
		$\boldsymbol{c}^{(1)}, \ldots, \boldsymbol{c}^{(k)} \leftarrow \mathcal{C}$
	$\xleftarrow{\boldsymbol{c}^{(1)}, \ldots, \boldsymbol{c}^{(k)}}$	
$\vec{\boldsymbol{z}} := \vec{\boldsymbol{y}} + \sum \boldsymbol{c}^{(j)} \vec{\boldsymbol{s}}^{(j)}$, and rejection sample Define $\boldsymbol{P}^{(i)} = \boldsymbol{c}^{(i)} \boldsymbol{T}^{(i)}$ Define $\boldsymbol{g}^* := \sum_{j=1}^{k} \boldsymbol{c}^{(j)} \boldsymbol{g}^{(j)} - \tilde{\boldsymbol{g}}^{(w)} + \vec{\boldsymbol{b}} \cdot \vec{\boldsymbol{z}}$ $\pi =$ ZKPoK that $\begin{bmatrix} \boldsymbol{B}\vec{\boldsymbol{z}} - \vec{\boldsymbol{w}} \\ \boldsymbol{g}^* \end{bmatrix}$ is a commitment to $\sum_{j=1}^{k} \boldsymbol{P}^{(j)} \cdot (\vec{v}_1^{(j)} \otimes \cdots \otimes \vec{v}_m^{(j)})$	$\xrightarrow{\vec{\boldsymbol{z}}, \pi}$	
		1. check that $\|\vec{\boldsymbol{z}}\|$ is small 2. verify π

Fig. 6. Given $\boldsymbol{T}^{(j)} \cdot (\vec{v}_1^{(j)} \otimes \ldots \otimes \vec{v}_m^{(j)}) = \begin{bmatrix} \boldsymbol{B} \\ \vec{\boldsymbol{b}} \end{bmatrix} \vec{\boldsymbol{s}}^{(j)} + \begin{bmatrix} \vec{\boldsymbol{0}} \\ \boldsymbol{m}^{(j)} \end{bmatrix}$, the prover creates a BDLOP commitment to all the k $\boldsymbol{m}^{(j)}$ and proves its correctness. The new commitment Com uses public matrices (e.g. $\boldsymbol{A}$, etc. as in Figure 5) which we do not explicitly state in this sketch. The terms comprising $\boldsymbol{g}^*$ are parts of ω, and are described in detail in the protocol in Figure 5 (except with subscripts instead of superscripts).

Aggregation and Set Membership. Converting the protocol from Figure 5 into the one in Figure 6 uses very similar intuition as when converting a signature scheme into a ring signature scheme in Figure 1.

We will now proceed to briefly explain the transition from the protocol in Figure 5 to the one in Figure 6. First, the second verifier check in Figure 5 cannot be done in the clear – that is the verifier cannot know $\vec{\boldsymbol{w}}$. If he knows $\vec{\boldsymbol{w}}$, then he can compute the weighted sum of the committed values $\sum \boldsymbol{c}_i \vec{\boldsymbol{t}}_i$, which would leak information about which commitments were chosen. The prover therefore must commit to $\vec{\boldsymbol{w}}$. So the commitment ω in Figure 6 creates commitments to

$\boldsymbol{m}^{(j)}, \tilde{\boldsymbol{w}}$ exactly like to $\boldsymbol{m}_i, \tilde{\boldsymbol{w}}$ in Figure 5, and also commits to $\vec{\boldsymbol{w}}$ and to $V^{(j)}$, which are needed for the set membership proof.

The prover then sets up the $\vec{\boldsymbol{a}}^*$ and $\boldsymbol{g}^*$ exactly as in Figure 5. Therefore $\boldsymbol{g}^*$ is a commitment to the bottom part of $\sum_{j=1}^{k} \boldsymbol{c}^{(j)} \boldsymbol{T}^{(j)} \cdot (\vec{v}_1^{(j)} \otimes \cdots \otimes \vec{v}_m^{(j)})$. From the second verification equation in Figure 5, we know that the top part of the preceding is $\boldsymbol{B}\vec{\boldsymbol{z}} - \vec{\boldsymbol{w}}$, and we can create a commitment to this value by adding $\boldsymbol{B}\vec{\boldsymbol{z}}$ to the commitment of $-\vec{\boldsymbol{w}}$ that we already have. We therefore have a commitment to $\sum_{j=1}^{k} \boldsymbol{c}^{(j)} \boldsymbol{T}^{(j)} \cdot (\vec{v}_1^{(j)} \otimes \cdots \otimes \vec{v}_m^{(j)})$ and creating the proof π is therefore equivalent to creating a proof for (5) and (6). Showing that this protocol is sound is done the same way as the one in Figure 5 because $\vec{\boldsymbol{z}}$ and π in Figure 6 satisfy the three verification parts in Figure 5.

2 Preliminaries

2.1 Notation

Let $\mathsf{N} \in \mathbb{N}$ be a security parameter and q be an odd prime. We write $x \leftarrow S$ when $x \in S$ is sampled uniformly at random from the finite set S and similarly $x \leftarrow D$ when x is sampled according to the distribution D. For $a < b$ and $n \in \mathbb{N}$, we define $[a, b] := \{a, a+1 \ldots, b\}$ and $[n] := [1, n]$. Given two functions $f, g : \mathbb{N} \to [0, 1]$, we write $f(\mu) \approx g(\mu)$ if $|f(\mu) - g(\mu)| < \mu^{-\omega(1)}$. A function f is negligible if $f \approx 0$. We write $\mathsf{negl}(n)$ to denote an unspecified negligible function in n.

For a power of two d, denote $\mathcal{R}$ and $\mathcal{R}_q$ respectively to be the rings $\mathbb{Z}[X]/(X^d+1)$ and $\mathbb{Z}_q[X]/(X^d+1)$. Bold lower-case letters denote elements in $\mathcal{R}$ or $\mathcal{R}_q$ and bold lower-case letters with arrows represent column vectors with coefficients in $\mathcal{R}$ or $\mathcal{R}_q$. We also write bold upper-case letters for matrices in $\mathcal{R}$ or $\mathcal{R}_q$. By default, for a polynomial denoted as a bold letter, we write its i-th coefficient as its corresponding regular font letter subscript i, e.g. $f_0 \in \mathbb{Z}_q$ is a constant coefficient of $\boldsymbol{f} \in \mathcal{R}_q$.

2.2 Cyclotomic Rings

Suppose q splits into l prime ideals of degree d/l in $\mathcal{R}$. This means $X^d + 1 \equiv \boldsymbol{\varphi}_1 \ldots \boldsymbol{\varphi}_l \pmod{q}$ with irreducible polynomials $\boldsymbol{\varphi}_j$ of degree d/l modulo q. We assume that $\mathbb{Z}_q$ contains a primitive $2l$-th root of unity $\zeta \in \mathbb{Z}_q$ but no elements whose order is a higher power of two, i.e. $q - 1 \equiv 2l \pmod{4l}$. Therefore, we have

$$X^d + 1 \equiv \prod_{j \in \mathbb{Z}_l} \left(X^{\frac{d}{l}} - \zeta^{2j+1} \right) \pmod{q}. \tag{13}$$

Let $\mathcal{M}_q := \{\boldsymbol{p} \in \mathbb{Z}_q[X] : \deg(\boldsymbol{p}) < d/l\}$ be the $\mathbb{Z}_q$-module of polynomials of degree less than d/l. We define the Number Theoretic Transform (NTT) of a

polynomial $\boldsymbol{p} \in \mathcal{R}_q$ as follows:

$$\mathsf{NTT}(\boldsymbol{p}) := \begin{bmatrix} \hat{\boldsymbol{p}}_0 \\ \vdots \\ \hat{\boldsymbol{p}}_{l-1} \end{bmatrix} \in \mathcal{M}_q^l \text{ where } \mathsf{NTT}(\boldsymbol{p})_j = \hat{\boldsymbol{p}}_j = \boldsymbol{p} \bmod (X^{\frac{d}{l}} - \zeta^{2j+1}).$$

Furthermore, we expand the definition of NTT to vectors of polynomials $\vec{\boldsymbol{p}} \in \mathcal{R}_q^k$, where the NTT operation is applied to each coefficient of $\vec{\boldsymbol{p}}$, resulting in a vector in $\mathcal{M}_q^{kl}$.

We also define the inverse NTT operation. Namely, for a vector $\vec{v} \in \mathcal{M}_q^l$, $\mathsf{NTT}^{-1}(\vec{v})$ is the polynomial $\boldsymbol{p} \in \mathcal{R}_q$ such that $\mathsf{NTT}(\boldsymbol{p}) = \vec{v}$.

Let $\vec{v} = (v_0, \ldots, v_{l-1}), \vec{w} = (w_0, \ldots, w_{l-1}) \in \mathcal{M}_q^l$. Then, we define the component-wise product $\vec{v} \circ \vec{w}$ to be the vector $\vec{u} = (u_0, \ldots, u_{l-1}) \in \mathcal{M}_q^l$ such that

$$u_j = v_j w_j \bmod (X^{\frac{d}{l}} - \zeta^{2j+1})$$

for $j \in \mathbb{Z}_l$. By definition, we have the following property of the inverse NTT operation:

$$\mathsf{NTT}^{-1}(\vec{v}) \cdot \mathsf{NTT}^{-1}(\vec{w}) = \mathsf{NTT}^{-1}(\vec{v} \circ \vec{w}).$$

Similarly, we define the *inner product*:

$$\langle \vec{v}, \vec{w} \rangle = \sum_{j=0}^{l-1} \left(v_j w_j \bmod (X^{\frac{d}{l}} - \zeta^{2j+1}) \right).$$

We remark that this operation is not an inner product in the strictly mathematical sense (e.g. it is not linear). However, it has a few properties which are characteristic for an inner product. For instance, given arbitrary vectors $\vec{x}, \vec{y}, \vec{z} \in \mathcal{M}_q^l$ and scalar $c \in \mathbb{Z}_q$ we have: $\langle \vec{x}, \vec{y} \rangle = \langle \vec{y}, \vec{x} \rangle$ (symmetry), $\langle \vec{x} + \vec{y}, \vec{z} \rangle = \langle \vec{x}, \vec{z} \rangle + \langle \vec{y}, \vec{z} \rangle$ (distributive law) and $\langle c\vec{x}, \vec{y} \rangle = c\langle \vec{x}, \vec{z} \rangle$. We also highlight that the definition of $\langle \cdot, \cdot \rangle$ depends on the factors of $X^d + 1$ modulo q.

We generalise the newly introduced operations to work for vectors $\vec{v} = (\vec{v}_1, \ldots, \vec{v}_k)$ and $\vec{w} = (\vec{w}_1, \ldots, \vec{w}_k) \in \mathcal{M}_q^{kl}$ of length being a multiple of l in the usual way. In particular $\langle \vec{v}, \vec{w} \rangle = \sum_{i=1}^{k} \langle \vec{v}_i, \vec{w}_i \rangle$.

Eventually, for a matrix $A \in \mathcal{M}_q^{n \times kl}$ with rows $\vec{a}_1, \ldots, \vec{a}_n \in \mathcal{M}_q^{kl}$ and a vector $\vec{v} \in \mathcal{M}_q^{kl}$, we define the matrix-vector operation:

$$A\vec{v} = \begin{pmatrix} \langle \vec{a}_1, \vec{v} \rangle \\ \vdots \\ \langle \vec{a}_n, \vec{v} \rangle \end{pmatrix} \in \mathcal{M}_q^n.$$

In proving linear relations, we will need the following simple lemma.

Lemma 2.1. *Let $n, k \in \mathbb{N}$. Then, for any $A \in \mathcal{M}_q^{nl \times kl}, \vec{v} \in \mathcal{M}_q^{nl}$ and $\vec{s} \in \mathbb{Z}_q^{kl}$ we have*

$$\langle A\vec{s}, \vec{v} \rangle = \langle \vec{s}, A^T \vec{v} \rangle.$$

Proof. We prove the statement for $k = n = 1$. The proof can then be easily using the definition of an inner product. Let $\vec{a}_i$ be the $(i+1)$-th row of A and $a_{i,j} \in \mathcal{M}_q$ be its $(j+1)$-th coefficient. Similarly, we define s_i and v_i to be the $(i+1)$-th coefficient of $\vec{s}$ and $\vec{v}$ respectively. Then, by definition we have:

$$\begin{aligned}
\langle A\vec{s}, \vec{v}\rangle &= \sum_{i=0}^{l-1} \langle \vec{a}_i, \vec{s}\rangle v_i \bmod (X^{\frac{d}{l}} - \zeta^{2i+1}) \\
&= \sum_{i=0}^{l-1} \left(\sum_{j=0}^{l-1} a_{i,j} s_j \bmod (X^{\frac{d}{l}} - \zeta^{2j+1}) \right) v_i \bmod (X^{\frac{d}{l}} - \zeta^{2i+1}) \\
&= \sum_{i=0}^{l-1} \sum_{j=0}^{l-1} a_{i,j} s_j v_i \bmod (X^{\frac{d}{l}} - \zeta^{2i+1}) \\
&= \sum_{j=0}^{l-1} s_j \left(\sum_{i=0}^{l-1} a_{i,j} v_i \bmod (X^{\frac{d}{l}} - \zeta^{2i+1}) \right) \\
&= \langle \vec{s}, A^T \vec{v}\rangle.
\end{aligned} \tag{14}$$

Here, the crucial step was the observation that for $\vec{s} \in \mathbb{Z}_q^l$ and any $i, j \in \mathbb{Z}_l$ we have:

$$a_{i,j} s_j \bmod (X^{\frac{d}{l}} - \zeta^{2j+1}) = a_{i,j} s_j,$$

i.e. there is no reduction modulo the polynomial when multiplying by a scalar. □

Last but not least, we recall the following lemma from [10].

Lemma 2.2. *Let* $\boldsymbol{p} = p_0 + p_1 X + \ldots + p_{d-1} X^{d-1} \in \mathcal{R}_q$. *Then,*

$$\frac{1}{l} \sum_{i=0}^{l} \mathsf{NTT}(\boldsymbol{p})_i = \sum_{i=0}^{d/l-1} p_i X^i.$$

For our constructions in this work, the practical hardness of either of the problems against known attacks is not affected by the parameter m. Therefore, we sometimes simply write M-SIS$_{\kappa,B}$ or M-LWE$_{\lambda,\chi}$. The parameters κ and λ denote the *module ranks* for M-SIS and M-LWE, respectively. Also, when χ is a uniform distribution for the set $[-\mu, \mu]$, we simply denote M-LWE$_{\lambda,\mu}$.

2.3 Probability Distributions

In this paper we sample the coefficients of the random polynomials in the commitment scheme using the distribution χ on $\{-1, 0, 1\}$ where ± 1 both have probability 5/16 and 0 has probability 6/16 identically as in [6, 1, 10].

Discrete Gaussian distribution. We now define the discrete Gaussian distribution used for the rejection sampling.

Definition 2.3. *The* discrete Gaussian distribution *on $\mathcal{R}^\ell$ centered around $\vec{\boldsymbol{v}} \in \mathcal{R}^\ell$ with standard deviation $\mathfrak{s} > 0$ is given by*

$$D^{\ell d}_{\boldsymbol{v},\mathfrak{s}}(\vec{\boldsymbol{z}}) = \frac{e^{-\|\vec{\boldsymbol{z}}-\vec{\boldsymbol{v}}\|^2/2\mathfrak{s}^2}}{\sum_{\vec{\boldsymbol{z}}'\in\mathcal{R}^\ell} e^{-\|\vec{\boldsymbol{z}}'\|^2/2\mathfrak{s}^2}}.$$

When it is centered around $\vec{\mathbf{0}} \in \mathcal{R}^\ell$ we write $D^{\ell d}_{\mathfrak{s}} = D^{\ell d}_{\vec{\mathbf{0}},\mathfrak{s}}$

2.4 BDLOP Commitment Scheme

We recall the BDLOP commitment scheme from [3]. Suppose that we want to commit to a message vector $\vec{\boldsymbol{m}} = (\boldsymbol{m}_1, \ldots, \boldsymbol{m}_n) \in \mathcal{R}_q^n$ for $n \geq 1$ and that module ranks of κ and λ are required for M-SIS and M-LWE security, respectively. Then, in the key generation, a matrix $\boldsymbol{B}_0 \leftarrow \mathcal{R}_q^{\kappa\times(\kappa+\lambda+n)}$ and vectors $\vec{\boldsymbol{b}}_1, \ldots, \vec{\boldsymbol{b}}_n \leftarrow \mathcal{R}_q^{\kappa+\lambda+n}$ are generated and output as public parameters. Note that one could choose to generate $\boldsymbol{B}_0, \vec{\boldsymbol{b}}_1, \ldots, \vec{\boldsymbol{b}}_n$ in a more structured way as in [3] since it saves some computation. However, for readability, we write the commitment matrices in the "Knapsack" form as above. In our case, the hiding property of the commitment scheme is established via the duality between the Knapsack and M-LWE problems. We refer to [13, Appendix C] for a more detailed discussion.

To commit to the message $\vec{\boldsymbol{m}}$, we first sample $\vec{\boldsymbol{r}} \leftarrow \chi^{d\cdot(\kappa+\lambda+n)}$. Now, there are two parts of the commitment scheme: the binding part and the message encoding part. In particular, we compute

$$\begin{aligned}\vec{\boldsymbol{t}}_0 &= \boldsymbol{B}_0\vec{\boldsymbol{r}} \bmod q,\\ \boldsymbol{t}_i &= \langle\vec{\boldsymbol{b}}_i, \vec{\boldsymbol{r}}\rangle + \boldsymbol{m}_i \bmod q,\end{aligned}$$

for $i \in [n]$, where $\vec{\boldsymbol{t}}_0$ forms the binding part and each $\boldsymbol{t}_i$ encodes a message polynomial $\boldsymbol{m}_i$. In this paper, when we write that we compute a BDLOP commitment to a vector $\vec{m} = (\vec{m}_1, \ldots, \vec{m}_n) \in \mathcal{M}_q^{nl}$, we mean that we commit to the vector of polynomials $\vec{\boldsymbol{m}} = (\mathsf{NTT}^{-1}(\vec{m}_1), \ldots, \mathsf{NTT}^{-1}(\vec{m}_n)) \in \mathcal{R}_q^n$ as above.

Next, we define the notion of a weak opening of the commitment [1].

Definition 2.4. *A* weak opening *for the commitment $\vec{\boldsymbol{t}} = \vec{\boldsymbol{t}}_0 \parallel \boldsymbol{t}_1 \parallel \cdots \parallel \boldsymbol{t}_n$ consists of a polynomial $\bar{\boldsymbol{c}} \in \mathcal{R}_q$, a randomness vector $\vec{\boldsymbol{r}}^* \in \mathcal{R}_q$ and messages $\boldsymbol{m}_1^*, \ldots, \boldsymbol{m}_n^* \in \mathcal{R}_q$ such that*

$$\begin{aligned}&\|\bar{\boldsymbol{c}}\|_1 \leq 2d \text{ and } \bar{\boldsymbol{c}} \text{ is invertible over } \mathcal{R}_q\\ &\|\bar{\boldsymbol{c}}\vec{\boldsymbol{r}}^*\|_2 \leq 2\beta,\\ &\boldsymbol{B}_0\vec{\boldsymbol{r}}^* = \vec{\boldsymbol{t}}_0,\\ &\langle\vec{\boldsymbol{b}}_i, \vec{\boldsymbol{r}}^*\rangle + \boldsymbol{m}_i^* = \boldsymbol{t}_i \text{ for } i \in [n].\end{aligned}$$

Attema et al. [1] show that the commitment scheme is still binding with respect to weak openings if M-SIS$_{\kappa,8d\beta}$ is hard.

3 Efficient Lattice-Based Set Membership Proof

In this section we construct an efficient logarithmic-size ring signature protocol using recent results [1, 10, 18, 19] as the building blocks. Security analysis of the interactive protocol as well as ring signature instantiation are described in the full version of the paper [20].

3.1 Overview

In order to showcase our main techniques, let us consider the following set membership problem. Namely, suppose we would like to prove knowledge of a secret element $\vec{w}_i \in \mathcal{M}_q^{kl}$, for some $k \in \mathbb{N}$, such that $\vec{w} \in S$, where S is a public set $S = \{\vec{p}_1, \ldots, \vec{p}_n\} \subseteq \mathcal{M}_q^{kl}$ of size $n = l^m$ which is a power of l.

We now use the observation from [12, 14, 5] that $\vec{w} \in S$ if and only if there exists a binary vector $\vec{v} \in \{0,1\}^n$ with exactly one 1 such that $P\vec{v} = \vec{w}$ where $P \in \mathcal{M}_q^{kl \times n}$ is the matrix with i-th column being $\vec{p}_i$. One could then directly prove knowledge of $\vec{w}$ and $\vec{v}$ which satisfy conditions above using e.g. the protocol from [10, 18]. However, the proof size grows significantly when n gets bigger. In order to overcome this limitation, [14, 5] observe that vector $\vec{v}$ can be uniquely decomposed into smaller vectors $\vec{v}_1, \ldots, \vec{v}_m \in \{0,1\}^l$ which have exactly one 1 each and

$$\vec{v} = \vec{v}_1 \otimes \vec{v}_2 \otimes \cdots \otimes \vec{v}_m. \tag{15}$$

In the end, we want to commit to $\vec{w}$ and smaller vectors $\vec{v}_1, \ldots, \vec{v}_m$ and prove

$$P(\vec{v}_1 \otimes \cdots \otimes \vec{v}_m) = \vec{w} \tag{16}$$

along with

$$\vec{v}_i \circ (\vec{v}_i - \vec{1}) = \vec{0} \text{ and } \langle \vec{1}, \vec{v}_i \rangle = 1 \text{ for } i \in [m] \tag{17}$$

where for an integer $a \in \mathbb{Z}_q$, $\vec{a} := (a, \ldots, a) \in \mathbb{Z}_q^l$. We highlight that Equation 16 is over the $\mathbb{Z}_q$-module $\mathcal{M}_q$ (see Section 2.2).

We now present a new recursive approach to prove (16) and (17) efficiently. For readability, we first introduce the following notation:

$$\begin{aligned} &\vec{u}_j := \vec{v}_j \otimes \cdots \otimes \vec{v}_m \text{ for } j \in [m], \\ &P_1 := P \text{ and } \vec{x}_1 = (\vec{x}_{1,1}, \ldots, \vec{x}_{1,k}) := \vec{w}. \end{aligned}$$

We start by sending the BDLOP commitments (as described in Sections 1.1 and 2.4) to $\vec{v}_1, \ldots, \vec{v}_m, \vec{w}_1, \ldots, \vec{w}_k$ to the verifier:

$$\begin{aligned} \vec{\boldsymbol{t}}_0 &= \boldsymbol{B}_0 \vec{\boldsymbol{r}} \bmod q, \\ \boldsymbol{t}_i &= \langle \vec{\boldsymbol{b}}_i, \vec{\boldsymbol{r}} \rangle + \mathsf{NTT}^{-1}(\vec{v}_i) \bmod q \text{ for } i \in [m] \\ \boldsymbol{t}_{m+i} &= \langle \vec{\boldsymbol{b}}_{m+i}, \vec{\boldsymbol{r}} \rangle + \mathsf{NTT}^{-1}(\vec{x}_i) \bmod q \text{ for } i \in [k]. \end{aligned}$$

Then, a verifier $\mathcal{V}$ sends a challenge $\vec{\gamma}_1 = (\vec{\gamma}_{1,1}, \ldots, \vec{\gamma}_{1,k}) \leftarrow \mathcal{M}_q^{kl}$. Clearly, if (16) holds then we have

$$\langle P_1(\vec{v}_1 \otimes \vec{u}_2) - \vec{x}_1, \vec{\gamma}_1 \rangle = 0.$$

Otherwise, the probability that the inner product above is equal to zero is exactly $q^{-d/l}$ which is negligible.

Now, by Lemma 2.1 and using the fact that each $\vec{v}_i \in \mathbb{Z}_q^l$, we have:

$$\begin{aligned}\langle P_1(\vec{v}_1 \otimes \vec{u}_2) - \vec{x}_1, \vec{\gamma}_1\rangle &= \langle \vec{v}_1 \otimes \vec{u}_2, P_1^T\vec{\gamma}_1\rangle - \langle \vec{x}_1, \vec{\gamma}_1\rangle \\ &= \sum_{i=1}^{l} v_{1,i}\langle \vec{u}_2, P_{1,i}^T\vec{\gamma}_1\rangle - \langle \vec{x}_1, \vec{\gamma}_1\rangle \\ &= \sum_{i=1}^{l} v_{1,i}\gamma_1^T P_{1,i}\vec{u}_2 - \langle \vec{x}_1, \vec{\gamma}_1\rangle \\ &= \vec{v}_1^T P_2\vec{u}_2 - \langle \vec{x}_1, \vec{\gamma}_1\rangle = \langle \vec{v}_1, P_2\vec{u}_2\rangle - \langle \vec{x}_1, \vec{\gamma}_1\rangle\end{aligned} \tag{18}$$

where we denote

$$P_1 = \left(P_{1,1}\ P_{1,2}\ \cdots\ P_{1,l}\right) \in \mathcal{M}_q^{l\times l^m}$$

and the matrix P_2 is defined as

$$P_2 := \begin{pmatrix}\gamma_1^T P_{1,1}\\ \vdots \\ \gamma_1^T P_{1,l}\end{pmatrix} \in \mathcal{M}_q^{l\times l^{m-1}}. \tag{19}$$

Let us define the following vectors:

$$\vec{x}_2 := P_2\vec{u}_2 \in \mathcal{M}_q^l \text{ and } \vec{y}_1 := \vec{v}_1 \circ \vec{x}_2 - \sum_{i=1}^{k} \vec{x}_{1,i} \circ \vec{\gamma}_{1,i}. \tag{20}$$

First, we prove that $\vec{x}_2$ is constructed correctly. Note that by definition of $\vec{u}_2$ we have

$$\vec{x}_2 = P_2(\vec{v}_2 \otimes \cdots \otimes \vec{v}_m)$$

which is of the form (16) but with one less tensor. Hence, in order to prove this equation, we recursively follow the argument above. Then, assuming one can prove (20) for $\vec{x}_2$, by Lemma 2.2 we know that $\langle P_1(\vec{v}_1 \otimes \cdots \otimes \vec{v}_m) - \vec{x}_1, \vec{\gamma}_1\rangle = 0$ if and only if $\boldsymbol{y}_1 := \mathsf{NTT}^{-1}(\vec{y}_1)$ has the first d/l coefficients equal to zero. We present how to prove this property for $\boldsymbol{y}_1$ below.

Let us fix $j = 2$. Suppose that $j < m$. Then, in order to show that $\vec{x}_2$ from (20) is well-formed, we apply the exact strategy as before. Namely, we send a commitment to $\vec{x}_j$:

$$\boldsymbol{t}_{m+k+j-1} = \langle \vec{\boldsymbol{b}}_{m+k+j-1}, \vec{\boldsymbol{r}}\rangle + \mathsf{NTT}^{-1}(\vec{x}_j).$$

Then, given a challenge $\vec{\gamma}_j \leftarrow \mathcal{M}_q^l$, we deduce as in Equation 18 that

$$\langle P_j(\vec{v}_j \otimes \vec{u}_{j+1}) - \vec{x}_j, \vec{\gamma}\rangle = \langle \vec{v}_j, P_{j+1}\vec{u}_{j+1}\rangle - \langle \vec{x}_j, \vec{\gamma}_j\rangle$$

where

$$P_j = \left(P_{j,1}\ P_{j,2}\ \cdots\ P_{j,l}\right) \in \mathcal{M}_q^{l\times l^{m-j+1}}$$

and the matrix P_{j+1} is defined as

$$P_{j+1} := \begin{pmatrix} \gamma_j^T P_{j,1} \\ \vdots \\ \gamma_j^T P_{j,l} \end{pmatrix} \in \mathcal{M}_q^{l \times l^{m-j}}. \tag{21}$$

Next, we define vectors $\vec{x}_{j+1}, \vec{y}_j \in \mathcal{M}_q^l$:

$$\vec{x}_{j+1} := P_{j+1}\vec{u}_{j+1} \text{ and } \vec{y}_j := \vec{v}_j \circ \vec{x}_{j+1} - \vec{x}_j \circ \vec{\gamma}_j. \tag{22}$$

Now, in order to prove well-formedness of $\vec{x}_{j+1}$ we simply run the argument from this paragraph for $j := j+1$. Assuming that $\vec{x}_{j+1}$ is constructed correctly, we also need to prove that the coefficients of $\vec{y}_j$ sum up to 0, i.e. the first d/l coefficients of $\boldsymbol{y}_j = \mathsf{NTT}^{-1}(\vec{y}_j)$ are all zeroes. Below we describe how it can be done for all the $\boldsymbol{y}_j$'s simultaneously.

Eventually, for $j = m$ we want to prove that $\vec{x}_m = P_m\vec{u}_m = P_m\vec{v}_m$ which is a simple linear proof from [10]. We also want to show $\langle \vec{1}, \vec{v}_i \rangle = 1$ for $i \in [m]$. All these relations can be combined into one linear equation:

$$\begin{pmatrix} 0 & 0 & \cdots & 0 & P_m \\ B & 0 & \cdots & 0 & 0 \\ 0 & B & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & B & 0 \end{pmatrix} \begin{pmatrix} \vec{v}_1 \\ \vdots \\ \vec{v}_m \end{pmatrix} = \begin{pmatrix} \vec{x}_m \\ \vec{e}_1 \\ \vdots \\ \vec{e}_1 \end{pmatrix} \tag{23}$$

where

$$B = \begin{pmatrix} 1 & \cdots & 1 \\ 0 & \cdots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \cdots & 0 \end{pmatrix} \in \mathbb{Z}_q^{l \times l} \text{ and } \vec{e}_1 = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \in \mathbb{Z}_q^l.$$

Let us denote $P_m \in \mathcal{M}_q^{(m+1)l \times ml}$ to be the matrix on the left-hand side of Equation 23.

We proceed to proving (23). First, we get a challenge vector

$$\vec{\gamma}_m = (\vec{\gamma}_{m,1}, \ldots, \vec{\gamma}_{m,m+1}) \leftarrow \mathcal{M}_q^{(m+1)l}$$

from $\mathcal{V}$ and deduce that:

$$\left\langle \tilde{P}_m \begin{pmatrix} \vec{v}_1 \\ \vdots \\ \vec{v}_m \end{pmatrix} - \begin{pmatrix} \vec{x}_m \\ \vec{e}_1 \\ \vdots \\ \vec{e}_1 \end{pmatrix}, \vec{\gamma}_m \right\rangle = \left\langle \begin{pmatrix} \vec{v}_1 \\ \vdots \\ \vec{v}_m \end{pmatrix}, \tilde{P}_m^T \vec{\gamma}_m \right\rangle - \langle \vec{x}_m, \vec{\gamma}_{m,1} \rangle - \sum_{i=1}^{m} \langle \vec{e}_1, \vec{\gamma}_{m,i+1} \rangle.$$

Let $\vec{x}_{m+1} = (\vec{x}_{m+1,1}, \ldots, \vec{x}_{m+1,m}) := \tilde{P}_m^T \vec{\gamma}_m \in \mathcal{M}_q^{ml}$ and

$$\vec{y}_m := \left(\sum_{i=1}^{m} \vec{v}_i \circ \vec{x}_{m+1,i} \right) - \vec{x}_m \circ \vec{\gamma}_{m,1} - \vec{e}_1 \circ \sum_{i=1}^{m} \vec{\gamma}_{m,i}. \tag{24}$$

Note that in this case $\vec{x}_{m+1}$ is public (as opposed to $\vec{x}_1, \ldots, \vec{x}_m$). Then, as before we get that $\boldsymbol{y}_m = y_{m,0} + y_{m,1}X + \ldots + y_{m,d-1}X^{d-1} = \mathsf{NTT}^{-1}(\vec{y}_m)$ satisfies:

$$y_{m,0} + \ldots + y_{m,d/l-1}X^{d/l-1} = \frac{1}{l}\left\langle \tilde{P}_m \begin{pmatrix} \vec{v}_1 \\ \vdots \\ \vec{v}_m \end{pmatrix} - \begin{pmatrix} \vec{x}_m \\ \vec{e}_1 \\ \vdots \\ \vec{e}_1 \end{pmatrix}, \vec{\gamma}_m \right\rangle.$$

Therefore, we need to argue that $\boldsymbol{y}_m$ has the first d/l polynomial coefficients equal to 0.

Finally, what have left to prove is that (i) polynomials $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m$ have the first d/l coefficients equal to zero and (ii) vectors $\vec{v}_i$ are binary. We first focus on (i) and adapt the strategy shown in [10]. At the beginning, we will commit to a uniformly random polynomial $\boldsymbol{g}$ which has the first d/l coefficients equal to zero:

$$\boldsymbol{t}_{k+2m} = \langle \vec{\boldsymbol{b}}_{k+2m}, \vec{\boldsymbol{r}} \rangle + \boldsymbol{g}.$$

Then, we will reveal the polynomial

$$\boldsymbol{h} = \boldsymbol{g} + \boldsymbol{y}_1 + \ldots + \boldsymbol{y}_m. \tag{25}$$

Hence, the verifier manually checks the the first d/l coefficients of $\boldsymbol{h}$ are indeed zeroes. On the other hand, to prove (25) we follow the approach for proving multiplicative relations from [1].

Let $\vec{\boldsymbol{y}} \leftarrow D^{(\kappa+\lambda+k+2m)}$ be the masking vector. That is, given a challenge polynomial $\boldsymbol{c} \leftarrow C$ from a challenge distribution C (defined in Section 3.2), the prover will output a masked opening $\vec{\boldsymbol{z}}$ of the randomness $\vec{\boldsymbol{r}}$ defined as: $\vec{\boldsymbol{z}} = \vec{\boldsymbol{y}} + \boldsymbol{c}\vec{\boldsymbol{r}}$. Then, define polynomials $\boldsymbol{f}_\eta$ as:

$$\boldsymbol{f}_\eta = \begin{cases} \langle \vec{\boldsymbol{b}}_\eta, \vec{\boldsymbol{y}} \rangle - \boldsymbol{c}\boldsymbol{v}_\eta \text{ if } \eta \in [m] \\ \langle \vec{\boldsymbol{b}}_{m+i}, \vec{\boldsymbol{y}} \rangle - \boldsymbol{c}\boldsymbol{x}_{1,i} \text{ for } \eta = m+i; i \in [k] \\ \langle \vec{\boldsymbol{b}}_{m+k+j}, \vec{\boldsymbol{y}} \rangle - \boldsymbol{c}\boldsymbol{x}_{j+1} \text{ for } \eta = m+k+j; j \in [m-1] \\ \langle \vec{\boldsymbol{b}}_{k+2m}, \vec{\boldsymbol{y}} \rangle - \boldsymbol{c}\boldsymbol{g} \text{ if } \eta = k+2m \end{cases}$$

where $\boldsymbol{x}_j = \mathsf{NTT}^{-1}(\vec{x}_j)$ and similarly for $\boldsymbol{v}_i$ and $\boldsymbol{\gamma}_j$. Note that $\boldsymbol{f}_\eta = \langle \vec{\boldsymbol{b}}_\eta, \vec{\boldsymbol{z}} \rangle - \boldsymbol{c}\vec{\boldsymbol{t}}_\eta$ for all η and thus can be calculated by the verifier.

First, let us focus on $\boldsymbol{y}_1$. By definition we have (see (20)):

$$\boldsymbol{F}_1 := \boldsymbol{f}_1\boldsymbol{f}_{m+k+1} + \boldsymbol{c}\sum_{i=1}^{k}\boldsymbol{\gamma}_{1,i}\boldsymbol{f}_{m+i} = \boldsymbol{\omega}_1 + \boldsymbol{\psi}_1\boldsymbol{c} + \boldsymbol{y}_1\boldsymbol{c}^2$$

where polynomials $\boldsymbol{\omega}_1, \boldsymbol{\psi}_1$ are defined as follows

$$\boldsymbol{\omega}_1 := \langle \vec{\boldsymbol{b}}_1, \vec{\boldsymbol{y}} \rangle \langle \vec{\boldsymbol{b}}_{m+k+1}, \vec{\boldsymbol{y}} \rangle$$
$$\boldsymbol{\psi}_1 := \sum_{i=1}^{k}\boldsymbol{\gamma}_{1,i}\langle \vec{\boldsymbol{b}}_{m+i}, \vec{\boldsymbol{y}} \rangle - \langle \vec{\boldsymbol{b}}_1, \vec{\boldsymbol{y}} \rangle \boldsymbol{x}_2 - \langle \vec{\boldsymbol{b}}_{m+k+1}, \vec{\boldsymbol{y}} \rangle \boldsymbol{v}_1$$

Now, by Definition of $\boldsymbol{y}_j$ (see (22)), for fixed $j \in [2, m-1]$ we have:

$$\boldsymbol{F}_j := \boldsymbol{f}_j \boldsymbol{f}_{m+k+j} + \boldsymbol{c}\boldsymbol{\gamma}_j \boldsymbol{f}_{m+k+j-1} = \boldsymbol{\omega}_j + \boldsymbol{\psi}_j \boldsymbol{c} + \boldsymbol{y}_j \boldsymbol{c}^2$$

where

$$\begin{aligned} \boldsymbol{\omega}_j &:= \langle \vec{\boldsymbol{b}}_j, \vec{\boldsymbol{y}} \rangle \langle \vec{\boldsymbol{b}}_{m+k+j}, \vec{\boldsymbol{y}} \rangle \\ \boldsymbol{\psi}_j &:= \boldsymbol{\gamma}_j \langle \vec{\boldsymbol{b}}_{m+k+j-1}, \vec{\boldsymbol{y}} \rangle - \langle \vec{\boldsymbol{b}}_j, \vec{\boldsymbol{y}} \rangle \boldsymbol{x}_{j+1} - \langle \vec{\boldsymbol{b}}_{m+k+j}, \vec{\boldsymbol{y}} \rangle \boldsymbol{v}_j. \end{aligned} \tag{26}$$

In case of $j = m$, we transform Equation 24 into:

$$\boldsymbol{F}_m := \boldsymbol{c} \left(- \sum_{i=1}^{m} \boldsymbol{x}_{m+1,i} \boldsymbol{f}_i + \boldsymbol{\gamma}_{m,1} \boldsymbol{f}_{k+2m-1} - \boldsymbol{e}_1 \sum_{i=1}^{m} \boldsymbol{\gamma}_{m,i} \right) = \boldsymbol{\psi}_m \boldsymbol{c} + \boldsymbol{y}_m \boldsymbol{c}^2$$

where

$$\boldsymbol{\psi}_m := - \sum_{i=1}^{m} \boldsymbol{x}_{m+1,i} \langle \vec{\boldsymbol{b}}_i ., \vec{\boldsymbol{y}} \rangle + \boldsymbol{\gamma}_{m,1} \langle \vec{\boldsymbol{b}}_{k+2m-1}, \vec{\boldsymbol{y}} \rangle - \boldsymbol{e}_1 \sum_{i=1}^{m} \boldsymbol{\gamma}_{m,i}. \tag{27}$$

Clearly, all $\boldsymbol{F}_j$ can be computed by the verifier. Therefore, if we denote

$$\boldsymbol{\omega}_{\mathsf{sm}} := \sum_{i=1}^{m-1} \boldsymbol{\omega}_i \text{ and } \boldsymbol{\psi}_{\mathsf{sm}} := \sum_{i=1}^{m} \boldsymbol{\psi}_i - \langle \vec{\boldsymbol{b}}_{k+2m}, \vec{\boldsymbol{y}} \rangle \tag{28}$$

then we obtain:

$$\sum_{j=1}^{m} \boldsymbol{F}_j - \boldsymbol{c}\boldsymbol{f}_{k+2m} - \boldsymbol{c}^2 \boldsymbol{h} = \boldsymbol{\omega}_{\mathsf{sm}} + \boldsymbol{\psi}_{\mathsf{sm}} \boldsymbol{c} + (\boldsymbol{y}_1 + \ldots + \boldsymbol{y}_m + \boldsymbol{g} - \boldsymbol{h}) \boldsymbol{c}^2.$$

Hence, we want to prove that the coefficient corresponding to the quadratic term of $\sum_{j=1}^{m} \boldsymbol{F}_j - \boldsymbol{c}\boldsymbol{f}_{k+2m} - \boldsymbol{c}^2 \boldsymbol{h}$ vanishes.

Recall that we still need to prove (ii), i.e. all $\vec{v}_i$'s are binary. We first get challenges $\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_m \leftarrow \mathcal{R}_q$ from the verifier. Then, we observe that

$$\sum_{i=1}^{m} \boldsymbol{\alpha}_i (\boldsymbol{f}_i^2 + \boldsymbol{c}\boldsymbol{f}_i) = \boldsymbol{\omega}_{\mathsf{bin}} + \boldsymbol{\psi}_{\mathsf{bin}} \boldsymbol{c} + \left(\sum_{i=1}^{m} \boldsymbol{\alpha}_i \boldsymbol{v}_i (\boldsymbol{v}_i - \boldsymbol{1}) \right) \boldsymbol{c}^2$$

where

$$\boldsymbol{\omega}_{\mathsf{bin}} := \sum_{i=1}^{m} \boldsymbol{\alpha}_i \langle \vec{\boldsymbol{b}}_i, \vec{\boldsymbol{y}} \rangle^2 \text{ and } \boldsymbol{\psi}_{\mathsf{bin}} := \sum_{i=1}^{m} \boldsymbol{\alpha}_i \langle \vec{\boldsymbol{b}}_i, \vec{\boldsymbol{y}} \rangle (\boldsymbol{1} - 2\boldsymbol{v}_i). \tag{29}$$

Therefore, we combine (i) and (ii) by proving that the quadratic term in

$$\boldsymbol{\alpha}_0 \left(\sum_{j=1}^{m} \boldsymbol{F}_j - \boldsymbol{c}\boldsymbol{f}_{k+2m} - \boldsymbol{c}^2 \boldsymbol{h} \right) + \sum_{i=1}^{m} \boldsymbol{\alpha}_i (\boldsymbol{f}_i^2 + \boldsymbol{c}\boldsymbol{f}_i) \tag{30}$$

is equal to zero. In order to do so, we commit to the garbage polynomial

$$\boldsymbol{t}_{k+2m+1} = \langle \vec{\boldsymbol{b}}_{k+2m+1}, \vec{\boldsymbol{r}} \rangle + \boldsymbol{\psi}_{\mathsf{bin}} + \boldsymbol{\alpha}_0 \boldsymbol{\psi}_{\mathsf{sm}}$$

and additionally send $\boldsymbol{\omega} := \langle \vec{\boldsymbol{b}}_{k+2m+1}, \vec{\boldsymbol{y}} \rangle + \boldsymbol{\omega}_{\mathsf{bin}} + \boldsymbol{\alpha}_0 \boldsymbol{\omega}_{\mathsf{sm}}$. Then, the verifier computes $\boldsymbol{f}_{k+2m+1} = \langle \vec{\boldsymbol{b}}_{k+2m+1}, \vec{\boldsymbol{z}} \rangle - \boldsymbol{c}\boldsymbol{t}_{k+2m+1}$ and checks whether:

$$\boldsymbol{\alpha}_0 \left(\sum_{j=1}^{m} \boldsymbol{F}_j - \boldsymbol{c}\boldsymbol{f}_{k+2m} - \boldsymbol{c}^2 \boldsymbol{h} \right) + \sum_{i=1}^{m} \boldsymbol{\alpha}_i (\boldsymbol{f}_i^2 + \boldsymbol{c}\boldsymbol{f}_i) + \boldsymbol{f}_{k+2m+1} \stackrel{?}{=} \boldsymbol{\omega}.$$

3.2 Main Protocol

We present our main lattice-based one-out-of-many proof using the techniques from Section 3.1 and show how it can be turned into an efficient, logarithmic-sized ring signature.

Similarly as in the previous works [12, 13], the secret key of a user is a vector $\vec{\boldsymbol{s}} \leftarrow [-\mu, \mu]^{\ell d}$ of short polynomials over $\mathcal{R}_q$ and the corresponding public key $\vec{\boldsymbol{pk}} \in \mathcal{R}_q^k$ is defined as $\vec{\boldsymbol{pk}} := \boldsymbol{A}\vec{\boldsymbol{s}}$ for a public matrix $\boldsymbol{A} \in \mathcal{R}_q^{k \times \ell}$. Suppose there are $n = l^m$ users in the ring [7] and for $\iota \in [n]$, let $\vec{\boldsymbol{pk}}_\iota$ be the public key corresponding to the ι-th user. Then, during the signing process, user ι wants to prove knowledge of a short vector $\vec{\boldsymbol{s}}$ such that

$$\boldsymbol{A}\vec{\boldsymbol{s}} \in \{\vec{\boldsymbol{pk}}_1, \ldots, \vec{\boldsymbol{pk}}_n\}$$

without revealing any information about its index ι.

We present the main protocol in Fig. 7 with verification equations in Fig. 9. User $\iota \in [n]$, which acts as a prover $\mathcal{P}$, starts by decomposing the index vector $\vec{v} = (0, \ldots, 0, 1, 0, \ldots, 0) \in \{0,1\}^n$, where the ι-th coefficient is equal to 1, into m smaller vectors of length l as in (15). Note that each $\vec{v}_i \in \mathbb{Z}_q^l$ satisfies (17). At the same time, $\mathcal{P}$ samples a masking $\vec{\boldsymbol{y}}' \leftarrow D_{\mathsf{s}'}^{\ell d}$ and computes $\vec{\boldsymbol{w}}' = (\boldsymbol{w}'_1, \ldots, \boldsymbol{w}'_k) = \boldsymbol{A}\vec{\boldsymbol{y}}' \in \mathcal{R}_q^k$. Furthermore, for the linear proof $\mathcal{P}$ generates a random $\boldsymbol{g} \in \mathcal{R}_q$ such that $g_0 = \ldots = g_{d/l-1} = 0$. Now, the prover sends the BDLOP commitments to $\vec{v}_i$ as well as to $\vec{\boldsymbol{w}}'$ and $\boldsymbol{g}$. Namely, it generates a randomness vector $\vec{\boldsymbol{r}} \leftarrow \chi^{(\lambda+\kappa+2m+1)d}$ and sends:

$$\begin{aligned}
\vec{\boldsymbol{t}}_0 &= \boldsymbol{B}_0 \vec{\boldsymbol{r}} \bmod q, \\
\boldsymbol{t}_i &= \langle \vec{\boldsymbol{b}}_i, \vec{\boldsymbol{r}} \rangle + \mathsf{NTT}^{-1}(\vec{v}_i) \ \text{ for } i \in [m] \\
\boldsymbol{t}_{m+i} &= \langle \vec{\boldsymbol{b}}_{m+i}, \vec{\boldsymbol{r}} \rangle + \boldsymbol{w}'_i \text{ for } i \in [k]. \\
\boldsymbol{t}_{k+2m} &= \langle \vec{\boldsymbol{b}}_{k+2m}, \vec{\boldsymbol{r}} \rangle + \boldsymbol{g}
\end{aligned}$$

[7] If there are less than l^m users then we simply add the zero vectors as public keys so that the ring has exactly l^m elements. Then the proof that the prover knows a short preimage to one of the columns implies that they must know a preimage to one of the actual public keys because knowing a preimage for one of the zero columns would constitute a SIS solution.

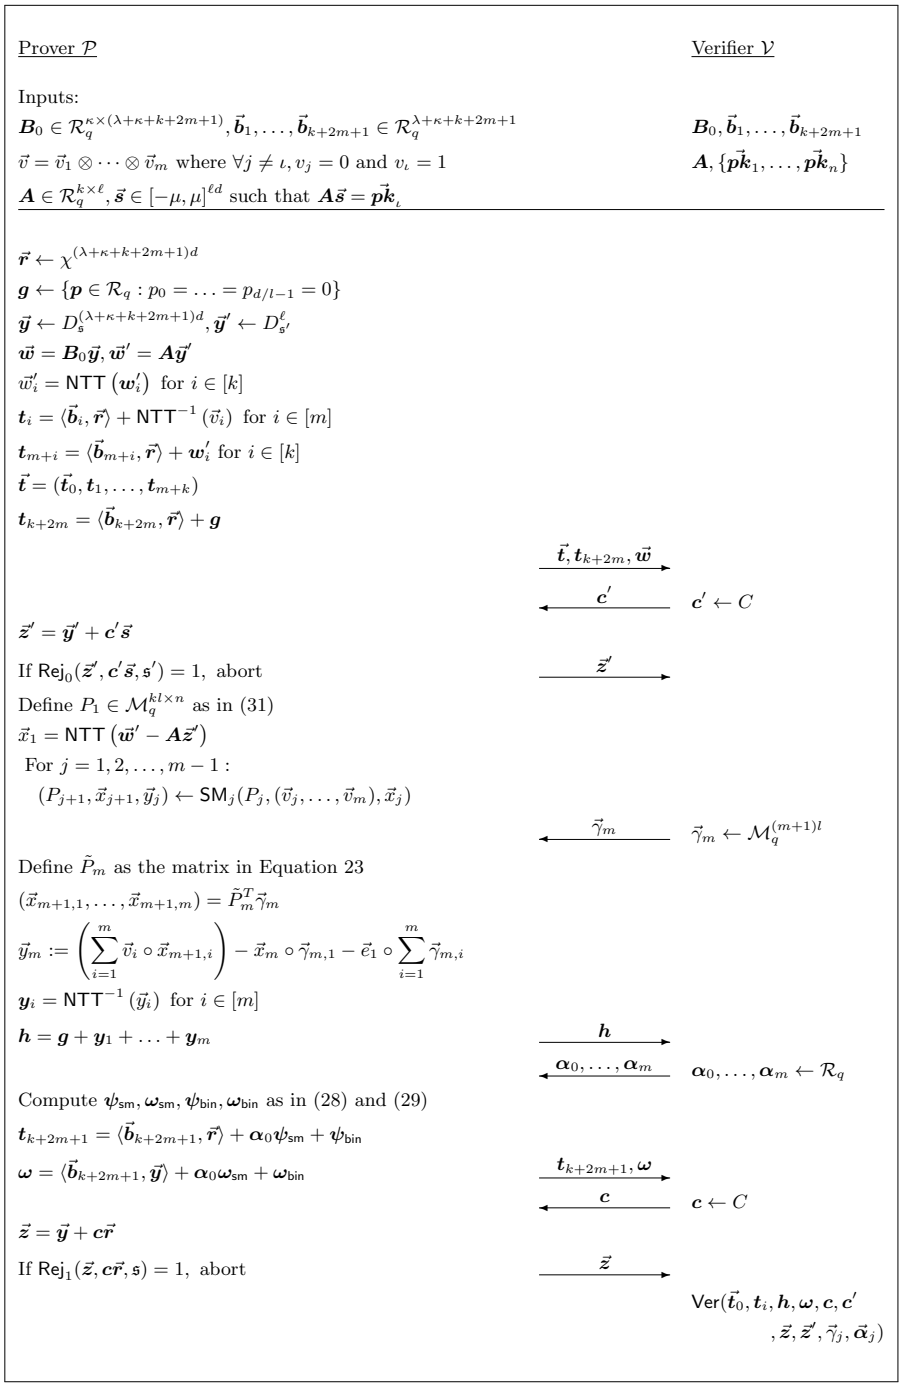

Fig. 7. Interactive protocol for our ring signature construction. Verifications equations Ver and the sub-protocol $\mathsf{SM}_j(P_j, (\vec{v}_j, \ldots, \vec{v}_m), \vec{x}_j)$ are defined in Fig. 9 and 8 respectively. We note that Rej_i, for $i = 0, 1$, are the rejection sampling algorithms from [17] and [19] respectively. See [20, Appendix A.3] for more details.

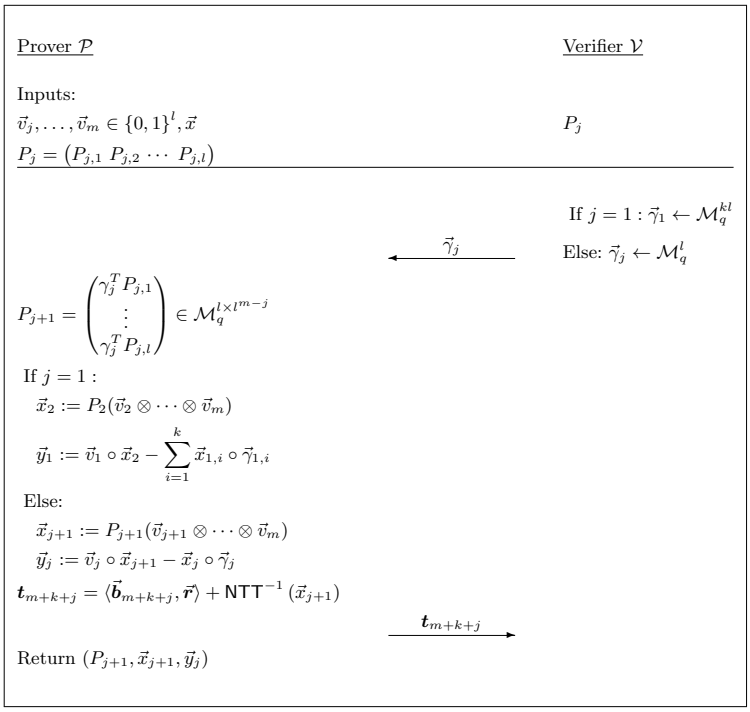

Fig. 8. The sub-protocol $\mathsf{SM}_j(P_j, (\vec{v}_j, \ldots, \vec{v}_m), \vec{x}_j)$ used in Fig. 7.

Additionally, $\mathcal{P}$ computes $\vec{\boldsymbol{w}} = \boldsymbol{B}_0 \vec{\boldsymbol{y}}$ for $\vec{\boldsymbol{y}}$ sampled from $D_{\mathfrak{s}}^{(\kappa+k+2m+1)d}$. Then, $\mathcal{P}$ sends

$$(\vec{\boldsymbol{t}}_0, \boldsymbol{t}_1, \ldots, \boldsymbol{t}_{m+k}, \boldsymbol{t}_{k+2m}, \vec{\boldsymbol{w}})$$

to the verifier.

The verifier $\mathcal{V}$ outputs a challenge polynomial $\boldsymbol{c}' \leftarrow C$. Next, $\mathcal{P}$ computes $\vec{\boldsymbol{z}}' = \vec{\boldsymbol{y}}' + \boldsymbol{c}'\vec{\boldsymbol{s}}$ and applies the rejection sampling algorithm. If it does not abort, $\mathcal{P}$ returns $\vec{\boldsymbol{z}}'$.

Let $P \in \mathcal{M}_q^{kl \times n}$ be the matrix defined as

$$P = \left(\mathsf{NTT}\left(-\boldsymbol{c}' \cdot \vec{\boldsymbol{pk}}_1\right) \middle| \cdots \middle| \mathsf{NTT}\left(-\boldsymbol{c}' \cdot \vec{\boldsymbol{pk}}_n\right)\right), \tag{31}$$

i.e. the i-th column of P is equal to $\mathsf{NTT}\left(-\boldsymbol{c}' \cdot \vec{\boldsymbol{pk}}_i\right) \in \mathcal{M}_q^{kl}$. Clearly, it can be computed by the verifier. Also, define

$$\vec{w} = \mathsf{NTT}\left(\boldsymbol{w}' - \boldsymbol{A}\vec{\boldsymbol{z}}'\right) \in \mathcal{M}_q^{kl}.$$

Then, user ι wants to prove that $P(\vec{v}_1 \otimes \cdots \otimes \vec{v}_m) = \vec{w}$. Obviously, the verifier can manually construct a commitment to $\vec{w}$ by subtracting $(\boldsymbol{t}_{m+1}, \ldots, \boldsymbol{t}_{m+k})$ by $\boldsymbol{A}\vec{\boldsymbol{z}}'$. One observes that this is the equation of type (16) and it is where we apply

$\underline{\mathsf{Ver}(\vec{\boldsymbol{t}}_0, \boldsymbol{t}_1, \ldots, \boldsymbol{t}_{k+2m+1}, \boldsymbol{h}, \boldsymbol{\omega}, \boldsymbol{c}, \boldsymbol{c}', \vec{\boldsymbol{z}}, \vec{\boldsymbol{z}}', \vec{\gamma}_1, \ldots, \vec{\gamma}_m, \boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_m)}$

$\|\vec{\boldsymbol{z}}'\|_2 \overset{?}{<} \beta' = \mathfrak{s}'\sqrt{2\ell d}$

$\|\vec{\boldsymbol{z}}\|_2 \overset{?}{<} \beta = \mathfrak{s}\sqrt{2(\lambda+\kappa+k+2m+1)d}$

$\boldsymbol{B}_0\vec{\boldsymbol{z}} \overset{?}{=} \vec{\boldsymbol{w}} + \boldsymbol{c}\vec{\boldsymbol{t}}_0$

$(\boldsymbol{t}_{m+1}, \ldots, \boldsymbol{t}_{m+k}) = (\boldsymbol{t}_{m+1}, \ldots, \boldsymbol{t}_{m+k}) - \boldsymbol{A}\vec{\boldsymbol{z}}' \in \mathcal{R}_q^k$

$\forall j \in [k+2m+1], \boldsymbol{f}_j = \langle \vec{\boldsymbol{b}}_j, \vec{\boldsymbol{z}} \rangle - \boldsymbol{c}\boldsymbol{t}_j$

$\forall i \in [m+1], \boldsymbol{\gamma}_{m,i} := \mathsf{NTT}^{-1}(\vec{\gamma}_{1,i})\,; \forall j \in [1,k], \boldsymbol{\gamma}_{1,j} := \mathsf{NTT}^{-1}(\vec{\gamma}_{1,j})$

$\forall j \in [2, m-1], \boldsymbol{\gamma}_j = \mathsf{NTT}^{-1}(\vec{\gamma}_j)$

$(\boldsymbol{x}_{m+1,1}, \ldots, \boldsymbol{x}_{m+1,m}) := \mathsf{NTT}^{-1}\left(\tilde{P}_m^T \vec{\gamma}_m\right)$ where $\tilde{P}_m$ is the matrix in (23)

$\boldsymbol{e}_1 := \mathsf{NTT}^{-1}((1, 0, \ldots, 0))$

$\boldsymbol{F}_1 := \boldsymbol{f}_1 \boldsymbol{f}_{m+k+1} + \boldsymbol{c}\sum_{i=1}^{k} \boldsymbol{\gamma}_{1,i} \boldsymbol{f}_{m+i}$

$\forall j \in [2, m-1], \boldsymbol{F}_j := \boldsymbol{f}_j \boldsymbol{f}_{m+k+j} + \boldsymbol{c}\boldsymbol{\gamma}_j \boldsymbol{f}_{m+k+j-1}$

$\boldsymbol{F}_m := \boldsymbol{c}\left(-\sum_{i=1}^{m} \boldsymbol{x}_{m+1,i}\boldsymbol{f}_i + \boldsymbol{\gamma}_{m,1}\boldsymbol{f}_{k+2m-1} - \boldsymbol{e}_1 \sum_{i=1}^{m} \boldsymbol{\gamma}_{m,i}\right)$

$\boldsymbol{\alpha}_0\left(\sum_{j=1}^{m} \boldsymbol{F}_j - \boldsymbol{c}\boldsymbol{f}_{k+2m} - \boldsymbol{c}^2\boldsymbol{h}\right) + \sum_{i=1}^{m} \boldsymbol{\alpha}_i(\boldsymbol{f}_i^2 + \boldsymbol{c}\boldsymbol{f}_i) + \boldsymbol{f}_{k+2m+1} \overset{?}{=} \boldsymbol{\omega}$

For $i = 0, \ldots, d/l - 1$:

  $h_i \overset{?}{=} 0$

Fig. 9. Verification equations for the protocol in Fig. 7.

the strategy described in Section 3.1. Namely, for $j = 1, 2, 3, \ldots, m-1$, we run a two-round sub-protocol $\mathsf{SM}_j(P_j, (\vec{v}_j, \ldots, \vec{v}_m), \vec{x}_j)$ defined in Fig. 8 which does the following. The verifier $\mathcal{V}$ starts by sending a challenge vector $\vec{\gamma}_j$. Then, $\mathcal{P}$ computes the matrix P_{j+1} and vectors $\vec{x}_{j+1}, \vec{y}_j \in \mathcal{M}_q^l$ as defined in the previous section. Eventually, it outputs the commitment to $\vec{x}_{j+1}$:

$$\boldsymbol{t}_{m+k+j} = \langle \vec{\boldsymbol{b}}_{m+k+j}, \vec{\boldsymbol{r}} \rangle + \mathsf{NTT}^{-1}(\vec{x}_{j+1}).$$

In the end, the sub-protocol returns

$$(P_{j+1}, \vec{x}_{j+1}, \vec{y}_j) \leftarrow \mathsf{SM}_j(P_j, (\vec{v}_j, \ldots, \vec{v}_m), \vec{x}_j).$$

After executing the SM sub-protocol $m-1$ times, the verifier sends $\vec{\gamma}_m \leftarrow \mathcal{M}_q^{(m+1)l}$. Then, in order to prove Equation 23, $\mathcal{P}$ first computes $\vec{y}_m$ as in Equation 24 and outputs the polynomial $\boldsymbol{h} = \boldsymbol{g} + \boldsymbol{y}_1 + \ldots + \boldsymbol{y}_m$, where $\boldsymbol{y}_i = \mathsf{NTT}^{-1}(\vec{y}_i)$ for $i \in [m]$.

Next, $\mathcal{V}$ sends uniform polynomials $\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_m \leftarrow \mathcal{R}_q$. Then, $\mathcal{P}$ returns a commitment

$$\boldsymbol{t}_{k+2m+1} = \langle \vec{\boldsymbol{b}}_{k+2m+1}, \vec{\boldsymbol{y}} \rangle + \boldsymbol{\psi}$$

to the garbage polynomial $\boldsymbol{\psi} = \boldsymbol{\psi}_{\mathsf{bin}} + \boldsymbol{\alpha}_0\boldsymbol{\psi}_{\mathsf{sm}}$ along with $\boldsymbol{\omega} := \langle \vec{\boldsymbol{b}}_{k+2m+1}, \vec{\boldsymbol{y}} \rangle + \boldsymbol{\omega}_{\mathsf{bin}} + \boldsymbol{\alpha}_0\boldsymbol{\omega}_{\mathsf{sm}}$ (where their components are defined in (28) and (29)).

Finally, the verifier picks a challenge $\boldsymbol{c} \leftarrow C$ and outputs $\boldsymbol{c}$. Here, the coefficients of a challenge $\boldsymbol{c} \leftarrow C$ are independently identically distributed with $P(0) = 1/2$ and $\Pr(1) = \Pr(-1) = 1/4$ [8]. Then, prover $\mathcal{P}$ computes $\vec{\boldsymbol{z}} = \vec{\boldsymbol{y}} + \boldsymbol{c}\vec{\boldsymbol{r}}$ and applies rejection sampling. If it does not abort, $\mathcal{P}$ returns $\vec{\boldsymbol{z}}$.

[8] We will make use of the properties of C described in [1]. We refer to [20, Appendix A.1] for more details.

Acknowledgements

We would like to thank anonymous reviews for useful feedback. This work was supported by the SNSF ERC Transfer Grant CRETP2-166734 FELICITY.

References

1. Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In *CRYPTO (2)*, volume 12171 of *Lecture Notes in Computer Science*, pages 470–499. Springer, 2020.
2. Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In *CRYPTO*, pages 669–699, 2018.
3. Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In *SCN*, pages 368–385, 2018.
4. Ward Beullens, Shuichi Katsumata, and Federico Pintore. Calamari and falafl: Logarithmic (linkable) ring signatures from isogenies and lattices. In *ASIACRYPT (2)*, volume 12492 of *Lecture Notes in Computer Science*, pages 464–492. Springer, 2020.
5. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, Jens Groth, and Christophe Petit. Short accountable ring signatures based on DDH. In *ESORICS (1)*, volume 9326 of *Lecture Notes in Computer Science*, pages 243–265. Springer, 2015.
6. Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In *CRYPTO (1)*, volume 11692 of *Lecture Notes in Computer Science*, pages 176–202. Springer, 2019.
7. Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In *ACM Conference on Computer and Communications Security*, pages 574–591. ACM, 2018.
8. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO (1)*, pages 40–56, 2013.
9. 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.
10. Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In *ASIACRYPT (2)*, volume 12492 of *Lecture Notes in Computer Science*, pages 259–288. Springer, 2020.
11. Muhammed F. Esgin, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In *CRYPTO (1)*, volume 11692 of *Lecture Notes in Computer Science*, pages 115–146. Springer, 2019.
12. Muhammed F. Esgin, Ron Steinfeld, Amin Sakzad, Joseph K. Liu, and Dongxi Liu. Short lattice-based one-out-of-many proofs and applications to ring signatures. In *ACNS*, volume 11464 of *Lecture Notes in Computer Science*, pages 67–88. Springer, 2019.

13. Muhammed F. Esgin, Raymond K. Zhao, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. Matrict: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In *CCS*, pages 567–584. ACM, 2019.
14. Jens Groth and Markulf Kohlweiss. One-out-of-many proofs: Or how to leak a secret and spend a coin. In *EUROCRYPT*, pages 253–280, 2015.
15. Xingye Lu, Man Ho Au, and Zhenfei Zhang. Raptor: A practical lattice-based (linkable) ring signature. In *ACNS*, volume 11464 of *Lecture Notes in Computer Science*, pages 110–130. Springer, 2019.
16. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
17. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
18. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Practical lattice-based zero-knowledge proofs for integer relations. In *CCS*, pages 1051–1070. ACM, 2020.
19. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Shorter lattice-based zero-knowledge proofs via one-time commitments. In *Public Key Cryptography (1)*, volume 12710 of *Lecture Notes in Computer Science*, pages 215–241. Springer, 2021.
20. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Smile: Set membership from ideal lattices with applications to ring signatures and confidential transactions. Cryptology ePrint Archive, Report 2021/564, 2021. `https://eprint.iacr.org/2021/564`.
21. Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In *EUROCRYPT (1)*, volume 10820 of *Lecture Notes in Computer Science*, pages 204–224. Springer, 2018.
22. Shen Noether. Ring signature confidential transactions for monero. *IACR Cryptol. ePrint Arch.*, 2015:1098, 2015.
23. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2017. https://csrc.nist.gov/projects/post-quantum-cryptography/ round-1-submissions.
24. Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In *ASIACRYPT*, volume 2248 of *Lecture Notes in Computer Science*, pages 552–565. Springer, 2001.
25. Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In *CRYPTO (1)*, volume 11692 of *Lecture Notes in Computer Science*, pages 147–175. Springer, 2019.