

BLOOM: Bimodal Lattice One-Out-of-Many Proofs and Applications

Vadim Lyubashevsky¹ and Ngoc Khanh Nguyen² 

¹ IBM Research Europe, Ruschlikon, Switzerland

² EPFL, Lausanne, Switzerland

Abstract. We give a construction of an efficient one-out-of-many proof system, in which a prover shows that he knows the pre-image for one element in a set, based on the hardness of lattice problems. The construction employs the recent zero-knowledge framework of Lyubashevsky et al. (Crypto 2022) together with an improved, over prior lattice-based one-out-of-many proofs, recursive procedure, and a novel rejection sampling proof that allows to use the efficient bimodal rejection sampling throughout the protocol.

Using these new primitives and techniques, we give instantiations of the most compact lattice-based ring and group signatures schemes. The improvement in signature sizes over prior works ranges between 25% and 2X. Perhaps of even more significance, the size of the user public keys, which need to be stored somewhere publicly accessible in order for ring signatures to be meaningful, is reduced by factors ranging from 7X to 15X. In what could be of independent interest, we also provide noticeably improved proofs for integer relations which, together with one-out-of-many proofs are key components of confidential payment systems.

Keywords: lattices, zero-knowledge, one-out-of-many proofs, ring signatures

1 Introduction

Zero-knowledge proofs are the cornerstone of privacy-enabling cryptography and the ones based on lattice assumptions appear to currently be the most practical potentially quantum-resistant variants. The fundamental hard problem upon which lattice cryptography is based on is finding a vector \vec{s} with a small norm satisfying $A\vec{s} = \vec{t} \bmod p$. Rapid recent progress in the area resulted in the proof size for proving pre-image knowledge in this basic equation to be reduced from being on the order of megabytes [LNSW13] to as short as a dozen kilobytes [YAZ⁺19, ESSL19, BLS19, ALS20, ENS20, LNS21a, LNP22].

A very useful extension of proving knowledge of a pre-image is proving knowledge of a pre-image for one element in a set. That is, given a set $\{\vec{t}_1, \dots, \vec{t}_m\}$, one would like to prove knowledge of a short vector \vec{s} such that $A\vec{s} = \vec{t}_i \bmod p$ without leaking any information about the \vec{s} or the i . This type of a proof is related

to concepts such as set membership proofs and one-out-of-many proofs [GK15].³ These proofs have applications to ring signatures, group signatures, confidential transactions, anonymous credentials, and various other privacy-enhancing cryptographic primitives.

In this work, we improve upon existing lattice-based one-out-of-many proofs and, based on this new building block, construct the most efficient quantum-safe ring and group signatures. Our improvement uses the high-level idea from the recursive algorithm of [LNS21b], but using a different and simpler recursive step which is made possible in part by being able to prove the base case using the recent framework for zero-knowledge proofs [LNP22]. We also give a general improvement to the rejection sampling step present in most lattice-based zero-knowledge proofs. Specifically, we show that when using the zero-knowledge framework from [LNP22], which gives the most efficient linear-size proofs for quadratic relations, one can use the more efficient bimodal rejection sampling procedure [DDLL13] everywhere. In some cases, this requires an extra short commitment, but in the case of the one-out-of-many proof given in the current work, using the more efficient rejection sampling step comes completely for free and noticeably reduces the proof size. We additionally show how to apply this bimodal rejection technique, together with the framework from [LNP22], to create more efficient proofs for integer relations such as addition and multiplication. Like one-out-of-many proofs, proving integer relations is a component of confidential transaction systems, and we believe that our improved tools can be used to make such systems (e.g. [LNS21b, ESZ21]) more efficient.

1.1 Results and Techniques

One-out-of-many proofs. The general equation for a lattice-based one-out-of- m proof can be written as

$$T\vec{v} = A\vec{s} \bmod p \tag{1}$$

where \vec{v} is an m -dimensional unit vector (i.e. a vector consisting of one 1 and the rest zeroes) and \vec{s} is a pre-image to the column in T chosen by the unit vector. For our end application, we would like to prove that \vec{s} has a small norm, but we will give a proof system for slightly more general statements, as the generality is useful in the recursive nature of the protocol. Given a commitment to a vector \vec{v} and \vec{s} , we would like to be able to prove that \vec{v} and \vec{s} satisfy (1), \vec{v} is a unit vector, and \vec{s} additionally satisfies some arbitrary quadratic relations $f_i(\vec{s}) = 0$.⁴ In applications to group and ring signatures, the dimension of \vec{s} is quite small, and so it is enough to only strive for proofs that are linear in the input size.

³ The original formal definition of a one-out-of-many proof from [GK15] is more restricted in that $A\vec{s}$ is a commitment to 0 rather than just an evaluation of a one-way function on \vec{s} . But we do not need to restrict to this definition in this work.

⁴ Being able to prove quadratic relations, of course also allows one to prove that the ℓ_2 -norm of \vec{s} is smaller than some bound.

Note that to keep the size of the proof linear, we still need it to be logarithmic in m since a unit vector of dimension m only has entropy $\log m$.

The proof of knowledge of (1) follows the commit-and-prove paradigm. That is, we commit to the secrets \vec{v} and \vec{s} and then prove that the committed messages satisfy the requisite relations. Keeping the proof linear in $\log m$ is the main technical challenge in building a one-out-of-many proof since one can't simply commit and open the naive m -dimensional representation of \vec{v} . Instead, one can write the unit vector \vec{v} as a tensor product of a logarithmic number of smaller dimensional unit vectors, commit to these unit vectors, and then proceed to recursively prove the relation.

We begin with the base case – proving the knowledge of a unit vector $\vec{v} \in \{0, 1\}^d$ (we can think of d being a constant with respect to m) and a vector \vec{s} satisfying $T\vec{v} = A\vec{s}$ where $f_i(\vec{s}) = 0$ for arbitrary quadratic functions f_i . The most efficient known proof for this statement directly follows from the recent framework of [LNP22] where one commits to the vector \vec{v} and \vec{s} using the “ABDLOP” commitment scheme defined in that work (it is a combination of the Ajtai [Ajt96] and BDLOP [BDL⁺18] commitments – see (11) and Section 2.6) and then proves that the committed values satisfy

$$T\vec{v} = A\vec{s} \pmod{p}, \quad (2)$$

and $f_i(\vec{s}) = 0$, and $\|\vec{v}\| = 1$.⁵

Now suppose, by the inductive hypothesis, that having a commitment to \vec{s}' and (some representation of) $\vec{v}' \in \{0, 1\}^m$, we are able to prove that they satisfy the linear relation

$$T'\vec{v}' = A'\vec{s}' \pmod{p} \quad (3)$$

as well as the quadratic relations $f_i(\vec{s}') = 0$ and that \vec{v}' is a unit vector. We will now show how to prove the relation in (3) for an arbitrary unit vector in $\{0, 1\}^{\mathfrak{d} \cdot m}$. First, observe that any unit vector in $\{0, 1\}^{\mathfrak{d} \cdot m}$ can be written as $\vec{v} \otimes \vec{v}' \in \{0, 1\}^{\mathfrak{d} \cdot m}$, where \vec{v} and \vec{v}' are unit vectors in $\{0, 1\}^{\mathfrak{d}}$ and $\{0, 1\}^m$, respectively. So by writing an arbitrary $\mathfrak{d} \cdot m$ -dimensional unit vector as $\vec{v} \otimes \vec{v}'$, we would like to prove (when having a commitment to $\vec{v}, \vec{v}', \vec{s}$) that

$$T(\vec{v} \otimes \vec{v}') = A\vec{s} \pmod{p}, \quad (4)$$

that $f_i(\vec{s}) = 0$, and that \vec{v}, \vec{v}' are unit vectors. To prove (4), we will prove that

$$\langle \vec{\varphi}_i, T(\vec{v} \otimes \vec{v}') - A\vec{s} \rangle = \vec{0} \pmod{p} \quad (5)$$

where $\vec{\varphi}_i$ for $i = 1, \dots, \mathfrak{l}$ are randomly-chosen challenge vectors in \mathbb{Z}_p^n . Proving one such equation for a randomly-chosen $\vec{\varphi}_i$ would result in the proof having soundness error p^{-1} , and so we would like to prove \mathfrak{l} such equations in order to achieve the desired soundness error of $p^{-\mathfrak{l}}$.

⁵ This only proves that \vec{v} is either a unit vector or a negative unit vector. But this is fine because proving knowledge of \vec{v}, \vec{s} satisfying $\pm T\vec{v} = A\vec{s}$ is equivalent to (1).

To prove (5) for one of l different $\vec{\varphi}_i$, we decompose the matrix T as $T = [T_1 \dots T_{\mathfrak{d}}]$ where $T_i \in \mathbb{Z}_p^{n \times m}$, and observe that by algebraic manipulation, we can rewrite

$$\langle \vec{\varphi}_i, T(\vec{v} \otimes \vec{v}') - A\vec{s} \rangle = \langle \vec{v}, \vec{w}_i \rangle - \vec{\varphi}_i^T A\vec{s} \pmod{p}, \quad (6)$$

where

$$\vec{w}_i = \begin{bmatrix} \vec{\varphi}_i^T T_1 \\ \dots \\ \vec{\varphi}_i^T T_{\mathfrak{d}} \end{bmatrix} \cdot \vec{v}' \in \mathbb{Z}_p^{\mathfrak{d}}. \quad (7)$$

Each of the l different $\vec{\varphi}_i$ leads to an equation of the above form and the prover thus commits to $\vec{w}_1, \dots, \vec{w}_l$ and then using the inductive hypothesis from (3), he can show that

$$\begin{bmatrix} \vec{w}_1 \\ \dots \\ \vec{w}_l \end{bmatrix} = T' \cdot \vec{v}' \pmod{p} \quad (8)$$

where the matrix T' is defined as

$$T' = \begin{bmatrix} \vec{\varphi}_1^T T_1 \\ \dots \\ \vec{\varphi}_1^T T_{\mathfrak{d}} \\ \dots \\ \vec{\varphi}_l^T T_1 \\ \dots \\ \vec{\varphi}_l^T T_{\mathfrak{d}} \end{bmatrix} \in \mathbb{Z}_p^{l \cdot \mathfrak{d} \times m}. \quad (9)$$

The inductive hypothesis also allows him to prove that $f_i(\vec{s}) = 0$ and additionally that (the quadratic functions) $\langle \vec{v}, \vec{w}_i \rangle - \vec{\varphi}_i^T A\vec{s} = 0$. To see that the inductive hypothesis is applicable to proving these statements, we define the vector \vec{s}' in (3) as

$$\vec{s}' = \begin{bmatrix} \vec{w} \\ \vec{s} \\ \vec{v} \end{bmatrix}, \text{ where } \vec{w} = \begin{bmatrix} \vec{w}_1 \\ \dots \\ \vec{w}_l \end{bmatrix} \quad (10)$$

and the matrix A' as $[I \mid 0 \mid 0]$, and thus $T'\vec{v}' = A'\vec{s}' = \vec{w}$. Since we assumed to have commitments to \vec{v}' and we created commitments to all parts of \vec{s}' , we can prove (8) and the aforementioned quadratic relations involving \vec{s} , \vec{w} , and \vec{v} . The main point is that by additionally committing to \vec{w} and \vec{v} , we are able to use the inductive hypothesis to prove relations where the unit vector is \mathfrak{d} times longer.

The commitment scheme used in [LNP22] allows to naturally commit to polynomials in the ring $\mathcal{R}_p = \mathbb{Z}_p[X]/(X^d + 1)$ where, for optimal efficiency, $d = 64$ or 128 . For efficiency of the protocol, we would then like to pack the l commitments to $\vec{w}_i \in \mathbb{Z}_p^{\mathfrak{d}}$ into just one vector \mathbb{Z}_p^d , which can then be represented by one polynomial in \mathcal{R}_p . We can do this in the trivial way as long as $\mathfrak{d} \cdot l \leq d$. Then to compute inner products $\langle \vec{v}, \vec{w}_i \rangle$, we simply put the vector $\vec{v} \in \{0, 1\}^{\mathfrak{d}}$ into a vector \mathbb{Z}_p^d that contains \vec{v} at the top and has the rest of its coefficients set to 0. Then

$\langle \vec{v}, \vec{w}_i \rangle$ is the inner product of an appropriate shift of the vector containing the \vec{v} and the vector containing the \vec{w}_i . A downward shift of a committed vector in \mathbb{Z}_p^d whose bottom coefficients are all 0 is simply a multiplication of the commitment by the polynomial X in \mathcal{R}_p , which can be performed by the verifier.

Thus proving $T \cdot (\vec{v} \otimes \vec{v}') = A\vec{s}$, where $\vec{v} \otimes \vec{v}'$ is a $\mathfrak{d} \cdot m$ dimensional unit vector, requires the commitments needed for proving (3), an additional commitment to \vec{v} , and one more commitment to the \vec{w}_i . Since the base case requires one commitment to \vec{v} (and a commitment to \vec{s}), the total number of commitments to elements in \mathcal{R}_p for proving (1) when \vec{v} is an $m = \mathfrak{d}^k \cdot d$ -dimensional unit vector is $2k + 1$ (and a commitments to \vec{s}), which is logarithmic in m and linear in the dimension of \vec{s} . In particular, we write the unit vector \vec{v} in (1) as $\vec{v}_1 \otimes \dots \otimes \vec{v}_k \otimes \vec{v}_{k+1}$ where $\vec{v}_1, \dots, \vec{v}_k \in \{0, 1\}^{\mathfrak{d}}$ and $\vec{v}_{k+1} \in \{0, 1\}^d$. Then proving (1), that \vec{v}_i are unit vectors, and $f_i(\vec{s}) = 0$ requires creating an ABDLOP commitment to the vectors \vec{v}_i, \vec{s} , and then also creating a commitment to the above-described vector \vec{w}_i at each step of the proof.

Bimodal Gaussian Rejection Sampling Everywhere. The framework of [LNP22] uses the newly-defined ABDLOP commitment scheme to commit to a low-norm polynomial vector \mathbf{s}_1 and an arbitrary-norm polynomial vector \mathbf{m} . To do this, one generates a low-norm randomness \mathbf{s}_2 and outputs the commitment

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_B \end{bmatrix} := \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \mathbf{s}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \mathbf{s}_2 + \begin{bmatrix} \mathbf{0} \\ \mathbf{m} \end{bmatrix} \text{ mod } q.^6 \quad (11)$$

Proving knowledge of \mathbf{s}_1 and \mathbf{m} is done using the ‘‘Fiat-Shamir with Aborts’’ technique [Lyu09, Lyu12] where, upon generating low-norm masking vectors \mathbf{y}_1 and \mathbf{y}_2 , computing $\mathbf{w} = \mathbf{A}_1 \mathbf{y}_1 + \mathbf{A}_2 \mathbf{y}_2$, and receiving a challenge polynomial c , the prover creates the responses $\mathbf{z}_1 = \mathbf{y}_1 + c\mathbf{s}_1$ and $\mathbf{z}_2 = \mathbf{y}_2 + c\mathbf{s}_2$ which satisfy

$$\mathbf{A}_1 \mathbf{z}_1 + \mathbf{A}_2 \mathbf{z}_2 = c\mathbf{t}_A + \mathbf{w} \text{ mod } q. \quad (12)$$

He now needs to perform rejection sampling on the \mathbf{z}_i in order to not leak information about the \mathbf{s}_i . The generic setup from [Lyu12] that results in the smallest-norm \mathbf{z}_i being sent involves \mathbf{y}_i being sampled from a discrete Gaussian distribution while standard deviation is approximately a factor of 12 larger than $\|\mathbf{c}\mathbf{s}_i\|$. With the appropriate rejection sampling step, this results in the polynomial vectors \mathbf{z}_i also being distributed as discrete Gaussians with standard deviation $12 \cdot \|\mathbf{c}\mathbf{s}_i\|$.

It was shown in [DDLL13] that if one first chooses a secret bit $b \in \{-1, 1\}$ and creates $\mathbf{z}_i = \mathbf{y}_i + b\mathbf{c}\mathbf{s}_i$ (which is a bimodal distribution with two peaks at $\pm\mathbf{c}\mathbf{s}_i$), then one can choose the \mathbf{y}_i from a discrete Gaussian distribution with standard deviation $\|\mathbf{c}\mathbf{s}_i\|/\sqrt{2}$ and via appropriate rejection sampling, the standard deviation of \mathbf{z}_i ends up being $\|\mathbf{c}\mathbf{s}_i\|/\sqrt{2}$ as well, which is around 17X smaller than that

⁶ We use modulus q here instead of p in the previous section to signify that the commitment scheme modulus need not be (and is usually not, though they could be related) the same as the modulus that one wants to prove relations over.

in the previous paragraph. The outputs \mathbf{z}_i having a smaller standard deviation means that the proof size will be noticeably smaller, the modulus q can be set smaller as well, which in turn results in a smaller commitment size. The main technical difficulty with using the bimodal distribution is that the bit b needs to remain secret and the \mathbf{z}_i need to satisfy the verification equation (12) irrespective of the b , which implies that we need to have $\mathbf{A}_1 b \mathbf{s}_1 + \mathbf{A}_2 b \mathbf{s}_2 = \mathbf{c} t_A \bmod q$. This set-up exists in the special case of the BLISS signature scheme [DDLL13] where the modulus is set to be even, and one can also force it by modifying the equation being proved as in [TWZ20], but these techniques would extol a high extra cost on the output size of the zero-knowledge proofs.

In our work, we show how one can use the bimodal rejection sampling technique for masking the secret vectors \mathbf{s}_1 and \mathbf{s}_2 in (11) either completely for free, or at a small increase in the commitment size. We note that using a rejection sampling procedure that had similar properties as bimodal rejection sampling was already employed in [LNS21a] on the *randomness vector* \mathbf{s}_2 . The rejection step there allowed for a smaller standard deviation at the cost of leaking one bit of \mathbf{s}_2 . This leakage is not a problem because the commitment scheme from (11) (and the related commitment in [LNS21a]) is used inside a commit-and-prove approach to constructing zero-knowledge proofs. In particular, when trying to prove some relation (e.g. (1)), the prover commits to the secret values (\vec{s} and \vec{v} in the case of (1)) using the commitment in (11) and proves relations about \mathbf{s}_1 and \mathbf{m} , which in turn implies the initial relation he set out to prove. The important part is that the commitment scheme is only used once – if the prover is to perform another proof, he will create another commitment with different randomness \mathbf{s}_2 . Thus leaking a small part of the randomness \mathbf{s}_2 is not a problem as long as the Module-LWE problem upon which the hiding of the commitment scheme is based on remains hard.

Our work improves on [LNS21a] in two ways. First, we show that directly using the bimodal rejection sampling on \mathbf{z}_2 (despite \mathbf{z}_2 not being distributed according to a bimodal distribution) only leaks a few more (i.e. $\log q$) bits of \mathbf{s}_2 but ends up saving a factor of 2 in the rejection sampling probability. Leaking $\log q$ bits still keeps the entropy of \mathbf{s}_2 very high. This problem has been investigated both theoretically and in practice [AP12, ASA16, DDGR20, BJRW21], and it does not seem that the LWE problem is weakened if a few bits of the randomness are leaked. The more interesting improvement is in also being able to use the bimodal rejection sampling on \mathbf{z}_1 . Here one cannot leak anything because \mathbf{s}_1 is where the actual secret message is stored.⁷ The new idea is to create a commitment to $b\mathbf{s}_1$, for a random $b \in \{-1, 1\}$ instead of to \mathbf{s}_1 . We can show that creating such a commitment and then outputting $\mathbf{z}_1 = \mathbf{y}_1 + cb\mathbf{s}_1$ after applying

⁷ One might be tempted to put the secret message into the \mathbf{m} part of the commitment, which does not leak even if there is leakage in \mathbf{s}_2 , but this results in a much less efficient commitment scheme because the dimension of the commitment grows linearly with the dimension of \mathbf{m} , whereas \mathbf{s}_1 has no effect on the size of the commitment. See Section 2.6.

the bimodal rejection sampling does not leak anything about \mathbf{s}_1 or b as long as the commitment is only used once.

There is, however, an obvious problem with committing to $b\mathbf{s}_1$ – the relation that one may want to prove about the message committed to in \mathbf{s}_1 may not hold true when the messages are negated. There are two cases here – the simple case is that what we would like to prove still holds true for the negated messages. In our one-out-of-many proof, we commit to unit vectors $\vec{v}_1, \dots, \vec{v}_{k+1}$ and a low-norm vector \vec{s} such that $\vec{v}_1 \otimes \dots \otimes \vec{v}_{k+1} = \vec{v}$ such that (1) is satisfied. Note that if we're tensoring an even number of \vec{v}_i , then $\bigotimes_i \vec{v}_i = \bigotimes_i -\vec{v}_i$, and also $\|\vec{v}_i\| = 1 = \|- \vec{v}_i\|$ and $\|\vec{s}\| = \beta = \|- \vec{s}\|$.⁸ Therefore proving (1) as well as $\|\vec{v}_i\| = 1$ and $\|\vec{s}\| = \beta$ can be done regardless of whether we committed to the positive or negative of these values.

In the case that we would like to prove some relations on some secret vector \vec{s} which are not sign-independent, we commit to a bit $b \in \{-1, 1\}$ in the \mathbf{m} part of the ABDLOP commitment and to $\vec{s}' = b\vec{s}$, and then prove knowledge of a vector \vec{s}' and a bit $b \in \{-1, 1\}$ satisfying $f_i(b\vec{s}') = 0$. Very importantly, note that $f_i(b\vec{s}')$ is still a quadratic equation because all the quadratic terms in f_i remain the same (since multiplication by $b \in \{-1, 1\}$ does not change them), and it is only the linear terms that get multiplied by b , thus becoming quadratic. We then prove that if the bit b is chosen randomly in $\{-1, 1\}$, then one can use the bimodal rejection sampling on $\mathbf{z}_1 = \mathbf{y}_1 + cb\mathbf{s}_1$ without leaking anything about the secret \mathbf{s}_1 .

Applications to Ring and Group Signatures. Being able to prove (1) immediately gives us a construction of a ring signature scheme. In particular, every user has a secret/public key pair \vec{s}_i, \vec{t}_i , where $\|\vec{s}_i\|$ is small and $A\vec{s}_i = \vec{t}_i \bmod p$. Given a matrix T whose columns are the public keys of a group of users, the signature of user i of a message μ is a zero-knowledge proof of knowledge (with μ being used as an input into the random oracle of the Fiat-Shamir transform) of a unit vector \vec{v} and a short vector \vec{s} satisfying (1). The full details are given in the full version of the paper. In Figure 1, we compare an instantiation of our ring signature with other known potentially quantum-safe ones. Once the group size within which one wishes to hide is larger than a few hundred members, the size of our signature is the smallest, even including the isogeny-based construction of [BKP20]. Additionally, the size of the public key of our ring signatures can be as small as 128 bytes per user, which is a significant reduction over all prior lattice-based ring signatures.⁹ Having small public keys is important because the

⁸ If we want to prove that $\|\vec{s}\| \leq \beta$, then we could create another commitment to a vector $\vec{s}' \in \mathcal{R}_q$ such that $\|\vec{s}\|^2 + \|\vec{s}'\|^2 = \beta^2$ – the existence of such an \vec{s}' is guaranteed by the four squares theorem.

⁹ It is of course possible to reduce the public key size of any scheme by hashing it as $\text{pk}' = H(\text{pk})$ for some cryptographic hash function H with the resulting pk' being as small as 32 bytes. This technique is fine for regular signatures, where one can reveal pk as part of the signature; but ring signatures will require a zero-knowledge proof

	sig. sizes for N :			asymptotic sig. size	hardness assumption	(user) public key size
	2^6	2^{12}	2^{21}			
Raptor [LAZ19]	81	5161	–	$O(N)$	NTRU	0.9
DualRing-LB [YEL ⁺ 21]	6	106	–	$O(N)$	MSIS, MLWE	[2.8, 3.4]
Falaf [BKP20]	32	35	39	$O(\log N)$	MSIS, MLWE	1.9
MatRiCT [EZS ⁺ 19]	31	59	148	$O(\log^{1.7} N)$	MSIS, MLWE	[3.4, 22.7]
MatRiCT+ [ESZ21]	11	18	40(?)	$O(\log^{1.7} N)$	MSIS, MLWE	3
SMILE [LNS21b]	18	19	22	$O(\log N)$	MSIS, E-MLWE	2
Calamari [BKP20]	8	14	23	$O(\log N)$	CSIDH-512	0.06
This Work	13	14	16	$O(\log N)$	MSIS, E-MLWE	0.13

Fig. 1: Comparison of the different post-quantum ring signature schemes with approximately 128 bits of security. Sizes from previous constructions are either taken from the corresponding prior work or from the recent survey by Buser et al. [BDE⁺22]. All the values are given in KB. Here, N is the size of the ring. The signature sizes for [ESZ21, LNS21b] only approximately correspond to the ring sizes (e.g. 18KB signature size is for the ring of 2^{10} users and not 2^{12}). For DualRing-LB and MatRiCT (and MatRiCT+) the public key size grows in the number of users. For MatRiCT+, the public key size for the ring of size 1024 is provided. Further, we extrapolate the signature size for MatRiCT+ with 2^{21} users from the smaller examples and from MatRiCT. In our construction, we rely on the Extended-MLWE problem introduced in Definition 5. Note that this is a different version of the E-MLWE problem compared to the one in [LNS21a] which is used in SMILE [LNS21b] (see Section 3.1 for more details).

public keys of *all* users need to be stored somewhere accessible by everyone who wishes to use the ring signature.

The reason for the significant reduction in the public key size over the previous lattice schemes is that we were able to adapt the new framework from [LNP22] as the base case of our recursive one-out-of-many proof. In prior ring signatures (e.g. [EZS⁺19, ESZ21, LNS21b]) the signer had knowledge of \vec{v} and \vec{s} satisfying (1), but for efficiency would only prove knowledge of an \vec{s}' and an additional low-norm polynomial c with $\|\vec{s}'\| \gg \|\vec{s}\|$ satisfying $T\vec{v} = cA\vec{s}' \bmod p$ (where the right-hand side operations are over a polynomial ring \mathcal{R}_p .) Being able to prove knowledge of a vector that has the exact norm of \vec{s} and not have an additional multiplication by c allows us to use a much smaller modulus p , which in turn also allows to reduce the number of rows in A . The public key size can, in fact, be essentially as small as the outputs in the hash function SWIFFT [LMR08].¹⁰

that $\mathbf{pk}' = H(\mathbf{pk})$, which will make the signatures orders of magnitude larger and slower.

¹⁰ If we make p too small, then the signature size will increase because a smaller p requires more “garbage terms” in the zero-knowledge proof to increase soundness. In our parameter settings, we chose a particular compromise between the public key

One can construct group signatures in a somewhat similar manner as ring signatures. A technique employed in [EZS⁺19, ESZ21] has the public key of each member stored in the matrix T , as in the ring signature, and the secret key of user i are the \vec{v} and \vec{s} from (1). The keys are generated by choosing a small-norm random \vec{s}_i and then putting $A\vec{s}_i = \vec{t}_i$ into the public matrix T . To sign, the group member does the same thing as in the ring signature (with an additional encryption and proof required by the group signature). Our group signature works in the same fashion except that the secret key / public key pairs are generated by first generating $\vec{t}_i = H(i)$, for some cryptographic hash function H , and then using a trapdoor for the matrix A to generate a short \vec{s}_i such that $A\vec{s}_i = \vec{t}_i$. The main advantage of generating the keys in this manner is that the public key size no longer needs to be linear in the number of group members, and can just consist of the matrix A (since everyone can now generate T themselves). The disadvantage is that using a trap-door sampler to generate \vec{s}_i results in $\|\vec{s}_i\|$ being larger. But because our one-out-of-many proof system can prove exact norms, the proof size does not increase by too much. Using GPV-type trapdoor sampling [GPV08, DP16] along with an optimized NTRU trapdoor [HHGP⁺03, DLP14] and the parameters used in the Falcon signature scheme [FHK⁺20], we give an instantiation (in the full version of the paper) of a lattice-based group signature scheme with the smallest public key and signature sizes (see Figure 2). The only exception when one would want to use a different scheme is in the case that the group sizes are very large – in that case [LNPS21, LNP22] has an advantage over all others in the table due to the fact that the running time for signature generation and verification are independent of the group size, rather than linear.

Other Applications. In addition to ring and group signatures, lattice-based one-out-of-many proofs have recently found applications in the constructions of confidential transaction protocols [EZS⁺19, LNS21b, ESZ21]. These constructions also used other primitives, notably proofs of addition that were used to make sure that the amounts in the transactions match up. As a side contribution, in the full version of the paper, we also show how to use the new framework of [LNP22] in conjunction with the bimodal rejection sampling technique to construct more efficient proofs of integer addition and multiplication, which improve upon the constructions from [LNS20, ESZ21] that are used in the aforementioned instantiations. We believe that the improved one-out-of-many proof and proof of addition from this paper should noticeably shorten the confidential transaction proof sizes. We leave the integration of these tools as well as the full implementation of the confidential transaction system to future work.

Acknowledgements. We would like to thank the anonymous reviewers for useful feedback. This work was conducted when the second author was at IBM Research Europe, and it was supported by the EU H2020 ERC Project 101002845 PLAZA.

size and signature size, but one could make the public key size even smaller at the expense of a few extra kilobytes in the signature size.

	signature sizes for N :			asymptotic sig. size	anonymity	master public key size
	2^6	2^{10}	2^{21}			
[LNPS21, LNP22]	90	90	90	$O(1)$	CPA	48
[BDK ⁺ 21][Lattice]	89	91	96	$O(\log N)$	CCA	$2.9 \cdot (N + 1)$
MatRiCT [Ezs ⁺ 19]	34	60	148	$O(\log^{1.7} N)$	CPA	$[3.8, 24] \cdot N$
MatRiCT+ [ESZ21]	12	19	45(?)	$O(\log^{1.7} N)$	CPA	$3 \cdot N$
[BDK ⁺ 21][Isogeny]	6.6	9.0	15.5	$O(\log N)$	CCA	$0.06 \cdot (N + 1)$
This Work	17	18	20	$O(\log N)$	CPA	3.2

Fig. 2: Comparison of different post-quantum group signature with approximately 128 bits of security. All the values are given in KB. Here, N denotes the size of the group. We note that the schemes from [LNPS21, LNP22] not only do offer constant size signatures but also enjoy signing and verification complexity independent of the group size N which is not the case for all the other works (including ours). Further, the constructions in [BDK⁺21, EZS⁺19, ESZ21] are dynamic which results in the linear public key size in the number of users N . For MatRiCT+, the public key size for the group of size 1024 is provided. Since the signature size for MatRiCT+ was not explicitly provided for group size 2^{21} , we set the value to be three times smaller than for MatRiCT which seems to be the case for smaller examples. Finally, we remark that our scheme can achieve CCA anonymity by following the Naor-Yung paradigm [NY90], i.e. encrypting the same message under two different public keys and adding a NIZK proof that both ciphertext encrypt the same message. We estimate that with this modification our group signature sizes will be around 30KB.

2 Preliminaries

2.1 Notation

Denote \mathbb{Z}_p to be the ring of integers modulo p . Let $q = q_1 \cdot \dots \cdot q_n$ be a product of n odd primes where $q_1 < q_2 < \dots < q_n$. Usually, we pick $n = 1$ or $n = 2$. In this paper we pick each $q_i = 5 \pmod{8}$. We write $\vec{v} \in \mathbb{Z}_q^m$ to denote vectors over a ring \mathbb{Z}_q . Matrices over \mathbb{Z}_q will be written as regular capital letters. By default, all vectors are column vectors. For simplicity, we denote $\vec{u}^2 = \vec{u} \circ \vec{u}$. 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 . Further, denote $[n] := \{1, \dots, n\}$.

2.2 Cyclotomic Rings

For a power of two d and a positive integer p , denote \mathcal{R} and \mathcal{R}_p respectively to be the rings $\mathbb{Z}[X]/(X^d + 1)$ and $\mathbb{Z}_p[X]/(X^d + 1)$. Lower-case letters denote elements in \mathcal{R} or \mathcal{R}_p and bold lower-case (resp. upper-case) letters represent column vectors (resp. matrices) with coefficients in \mathcal{R} or \mathcal{R}_p . For a polynomial $f \in \mathcal{R}_p$, denote $\vec{f} \in \mathbb{Z}_p^d$ to be the coefficient vector of f . By default, we write its i -th coefficient as its corresponding regular font letter subscript i , e.g. $f_{d/2} \in \mathbb{Z}_p$

is the coefficient corresponding to $X^{d/2}$ of $f \in \mathcal{R}_p$. For the constant coefficient, however, we will denote $\tilde{f} := f_0 \in \mathbb{Z}_p$.

The ring \mathcal{R} has a group of automorphisms $\text{Aut}(\mathcal{R})$ that is isomorphic to \mathbb{Z}_{2d}^\times . Let $\sigma_i \in \text{Aut}(\mathcal{R}_q)$ be defined by $\sigma_i(X) = X^i$. For readability, we denote for an arbitrary vector $\mathbf{m} \in \mathcal{R}^k$:

$$\sigma_i(\mathbf{m}) := (\sigma_i(m_1), \dots, \sigma_i(m_k))$$

and similarly $\sigma_i(\mathbf{R})$ for any matrix \mathbf{R} . When we write $\langle \mathbf{u}, \mathbf{v} \rangle \in \mathbb{Z}$ for $\mathbf{u}, \mathbf{v} \in \mathcal{R}^k$, we mean the inner product of their corresponding coefficient vectors.

We recall the result by Lyubashevsky et al. [LNP22] which says that for specific primes p , if $c \in \mathcal{R}_p$ satisfies $\sigma_{-1}(c) = c$ and c is non-zero then c is invertible over \mathcal{R}_p .

Lemma 1 ([LNP22]). *Let $p \equiv 5 \pmod{8}$ be a prime. Then all non-zero $c \in \mathcal{R}_p$ satisfying $\sigma_{-1}(c) = c$ are invertible.*

In this paper, we will only be interested in the $\sigma := \sigma_{-1}$ automorphism. The main reason is the following observation.

Lemma 2 ([LNP22]). *Let $\mathbf{u}, \mathbf{v} \in \mathcal{R}_q^k$. Then, the constant coefficient of $\sigma(\mathbf{u})^T \mathbf{v}$ is equal to $\langle \mathbf{u}, \mathbf{v} \rangle$.*

Thus, one reduces inner product arguments $\langle \mathbf{u}, \mathbf{v} \rangle = a$ to proving that $\sigma(\mathbf{u})^T \mathbf{v} - a \in \mathcal{R}_q$ has a vanishing constant coefficient.

We introduce the following notation:

$$\langle x \rangle_\sigma := (x, \sigma(x)) \in \mathcal{R}_q^2 \text{ for } x \in \mathcal{R}_q.$$

Similarly, for a vector $\mathbf{x} = (x_1, \dots, x_n)$, define $\langle \mathbf{x} \rangle_\sigma = (\langle x_1 \rangle_\sigma, \dots, \langle x_n \rangle_\sigma) \in \mathcal{R}_q^{2n}$. We will use the following simple properties.

Lemma 3. *For any $\mathbf{x}, \mathbf{y} \in \mathcal{R}_q^n$ and any $c \in \mathcal{R}_q$ such that $\sigma(c) = c$:*

$$\langle \mathbf{x} \parallel \mathbf{y} \rangle_\sigma = \langle \mathbf{x} \rangle_\sigma \parallel \langle \mathbf{y} \rangle_\sigma \quad \text{and} \quad \langle \mathbf{x} + c\mathbf{y} \rangle_\sigma = \langle \mathbf{x} \rangle_\sigma + c\langle \mathbf{y} \rangle_\sigma.$$

Next, we recall the definition of the discrete Gaussian distribution over \mathcal{R} .

Definition 1. *The discrete Gaussian distribution on \mathcal{R}^ℓ centered around $\mathbf{v} \in \mathcal{R}^\ell$ with standard deviation $\mathfrak{s} > 0$ is given by*

$$D_{\mathbf{v}, \mathfrak{s}}^\ell(\mathbf{z}) = \frac{e^{-\|\mathbf{z}-\mathbf{v}\|^2/2\mathfrak{s}^2}}{\sum_{\mathbf{z}' \in \mathcal{R}^\ell} e^{-\|\mathbf{z}'\|^2/2\mathfrak{s}^2}}.$$

When it is centered around $\mathbf{0} \in \mathcal{R}^\ell$ we write $D_{\mathfrak{s}}^\ell = D_{\mathbf{0}, \mathfrak{s}}^\ell$.

We will use the standard tail bound result from [Ban93, Lemma 1.5(i)].

Lemma 4. *Let $\mathbf{z} \leftarrow D_{\mathfrak{s}}^m$. Then $\Pr \left[\|\mathbf{z}\| > t \cdot \mathfrak{s}\sqrt{md} \right] < \left(te^{\frac{1-t^2}{2}} \right)^{md}$.*

2.3 Module-SIS and Module-LWE Problems

Security of the [BDL⁺18] commitment scheme used in our protocols relies on the well-known computational lattice problems, namely Module-LWE (MLWE) and Module-SIS (MSIS) [LS15]. Both problems are defined over \mathcal{R}_q .

Definition 2 (MSIS $_{\kappa,m,B}$). *Given $\mathbf{A} \leftarrow \mathcal{R}_q^{\kappa \times m}$, the Module-SIS problem with parameters $\kappa, m > 0$ and $0 < B < q$ asks to find $\mathbf{z} \in \mathcal{R}_q^m$ such that $\mathbf{A}\mathbf{z} = \mathbf{0}$ over \mathcal{R}_q and $0 < \|\mathbf{z}\| \leq B$. An algorithm \mathcal{A} is said to have advantage ϵ in solving MSIS $_{\kappa,m,B}$ if*

$$\Pr[0 < \|\mathbf{z}\|_\infty \leq B \wedge \mathbf{A}\mathbf{z} = \mathbf{0} \mid \mathbf{A} \leftarrow \mathcal{R}_q^{\kappa \times m}; \mathbf{z} \leftarrow \mathcal{A}(\mathbf{A})] \geq \epsilon.$$

Definition 3 (MLWE $_{m,\lambda,\chi}$). *The Module-LWE problem with parameters $m, \lambda > 0$ and an error distribution χ over \mathcal{R} asks the adversary \mathcal{A} to distinguish between the following two cases: 1) $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ for $\mathbf{A} \leftarrow \mathcal{R}_q^{m \times \lambda}$, a secret vector $\mathbf{s} \leftarrow \chi^\lambda$ and error vector $\mathbf{e} \leftarrow \chi^m$, and 2) $(\mathbf{A}, \mathbf{b}) \leftarrow \mathcal{R}_q^{m \times \lambda} \times \mathcal{R}_q^m$. Then, \mathcal{A} is said to have advantage ϵ in solving MLWE $_{m,\lambda,\chi}$ if*

$$\begin{aligned} & \left| \Pr[b = 1 \mid \mathbf{A} \leftarrow \mathcal{R}_q^{m \times \lambda}; \mathbf{s} \leftarrow \chi^\lambda; \mathbf{e} \leftarrow \chi^m; b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})] \right. \\ & \left. - \Pr[b = 1 \mid \mathbf{A} \leftarrow \mathcal{R}_q^{m \times \lambda}; \mathbf{b} \leftarrow \mathcal{R}_q^m; b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b})] \right| \geq \epsilon. \end{aligned} \quad (13)$$

2.4 Rejection Sampling

In lattice-based zero-knowledge proofs, the prover will want to output a vector \mathbf{z} whose distribution should be independent of a secret message/randomness vector \mathbf{r} , so that \mathbf{z} cannot be used to gain any information on the prover's secret. During the protocol, the prover computes $\mathbf{z} = \mathbf{y} + \mathbf{c}\mathbf{r}$ where \mathbf{r} is either a secret vector or randomness used to commit to the prover's secret, $c \leftarrow \mathcal{C}$ is a challenge polynomial, and \mathbf{y} is a “masking” vector. In order to remove the dependency of \mathbf{z} on \mathbf{r} , one applies *rejection sampling*. We summarise the two most common techniques for rejection sampling described in [Lyu12, DDLL13].

Lemma 5 (Rejection Sampling [Lyu12, DDLL13]). *Let $V \subseteq \mathcal{R}^\ell$ be a set of polynomials with norm at most T and $\rho: V \rightarrow [0, 1]$ be a probability distribution. Fix the standard deviation $\mathfrak{s} = \gamma T$. Then, the following statements hold.*

1. *Let $M = \exp(14/\gamma + 1/(2\gamma^2))$. Now, sample $\mathbf{v} \leftarrow \rho$ and $\mathbf{y} \leftarrow D_{\mathfrak{s}}^\ell$, set $\mathbf{z} = \mathbf{y} + \mathbf{v}$, and run $b \leftarrow \text{Rej}_1(\mathbf{z}, \mathbf{v}, \mathfrak{s})$ as defined in Fig. 3. Then, the probability that $b = 0$ is at least $(1 - 2^{-128})/M$ and the distribution of (\mathbf{v}, \mathbf{z}) , conditioned on $b = 0$, is within statistical distance of 2^{-128} of the product distribution $\rho \times D_{\mathfrak{s}}^\ell$.*
2. *Let $M = \exp(1/(2\gamma^2))$. Now, sample $\mathbf{v} \leftarrow \rho, \beta \leftarrow \{0, 1\}$ and $\mathbf{y} \leftarrow D_{\mathfrak{s}}^\ell$, set $\mathbf{z} = \mathbf{y} + (-1)^\beta \mathbf{v}$, and run $b \leftarrow \text{Rej}_2(\mathbf{z}, \mathbf{v}, \mathfrak{s})$ as defined in Fig. 3. Then, the probability that $b = 0$ is equal to $1/M$ and the distribution of (\mathbf{v}, \mathbf{z}) , conditioned on $b = 0$, is identical to the product distribution $\rho \times D_{\mathfrak{s}}^\ell$.*

$\text{Rej}_1(\vec{z}, \vec{v}, \mathfrak{s})$	$\text{Rej}_2(\vec{z}, \vec{v}, \mathfrak{s})$
01 $u \leftarrow [0, 1)$	01 $u \leftarrow [0, 1)$
02 If $u > \frac{1}{M} \cdot \exp\left(\frac{-2\langle \vec{z}, \vec{v} \rangle + \ \vec{v}\ ^2}{2\sigma^2}\right)$	02 If $u > \frac{1}{M \exp\left(-\frac{\ \vec{v}\ ^2}{2\sigma^2}\right) \cosh\left(\frac{\langle \vec{z}, \vec{v} \rangle}{\sigma^2}\right)}$
03 return 1 (i.e. <i>reject</i>)	03 return 1 (i.e. <i>reject</i>)
04 Else	04 Else
05 return 0 (i.e. <i>accept</i>)	05 return 0 (i.e. <i>accept</i>)

Fig. 3: Two rejection sampling algorithms: the one used generally in previous works [Lyu12] (left) and the bimodal Gaussian one [DDLL13] (right).

2.5 Challenge Space

We recall the specific challenge space used in [LNP22]. Namely, we fix $\eta > 0$ and a power-of-two k and set the challenge space \mathcal{C} as:

$$\mathcal{C} := \left\{ c \in S_\kappa : \sigma_{-1}(c) = c \wedge \sqrt[2k]{\|c^{2k}\|_1} \leq \eta \right\}. \quad (14)$$

Roughly speaking, the first condition, i.e. $\sigma_{-1}(c) = c$, is needed to prove quadratic equations in the committed messages which might additionally involve automorphisms, e.g. $m_1 m_2 = \sigma_{-1}(m_3)$ where m_1, m_2, m_3 are the secret messages. On the other hand, the second condition allows us to use [LNP22, Lemma 2.15] and deduce that if $\|\mathbf{r}\| \leq \alpha$ and $c \in \mathcal{C}$ then $\|\mathbf{c}\mathbf{r}\| \leq \eta\alpha$.

Further, we denote $\bar{\mathcal{C}} := \{c - c' : c, c' \in \mathcal{C} \text{ and } c \neq c'\}$ to be the set of differences of any two distinct elements in \mathcal{C} . We will choose the constant η such that (experimentally) the probability for $c \leftarrow S_\kappa$ to satisfy $\sqrt[2k]{\|c^{2k}\|_1} \leq \eta$ is at least 99%.

For security of our protocols, we need the invertibility property of the challenge space \mathcal{C} , i.e. the difference of any two distinct elements of \mathcal{C} is invertible over \mathcal{R}_q . To this end, we apply Lemma 1 and thus we only need the condition $\kappa < q_1/2$. Secondly, to achieve negligible soundness error, we will need $|\mathcal{C}|$ to be exponentially large. In Table 4 we propose example parameters to instantiate the challenge space \mathcal{C} .

d	κ	η	$ \mathcal{C} $
64	8	140	2^{129}
128	2	59	2^{147}

Fig. 4: Example parameters to instantiate the challenge space $\mathcal{C} := \{c \in S_\kappa : \sigma_{-1}(c) = c \wedge \sqrt[2k]{\|c^{2k}\|_1} \leq \eta\}$ for a modulus q such that its smallest prime divisor q_1 is greater than 16. In our examples we picked $k = 32$.

2.6 ABDLOP Commitment

We recall the ABDLOP commitment scheme defined in [LNP22], which is a generalisation of the Ajtai [Ajt96] and BDLOP [BDL⁺18] constructions. Concretely,

to commit to a message vector $\mathbf{s}_1 \in \mathcal{R}_q^{m_1}$ with small coefficients as well as a “full-fledged” polynomial vector $\mathbf{m} \in \mathcal{R}_q^\ell$, we sample a randomness vector $\mathbf{s}_2 \leftarrow \chi^{m_2}$, where χ is a probability distribution over \mathcal{R}_q , and compute:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_B \end{bmatrix} := \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \mathbf{s}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \mathbf{s}_2 + \begin{bmatrix} \mathbf{0} \\ \mathbf{m} \end{bmatrix}$$

where $\mathbf{A}_1 \leftarrow \mathcal{R}_q^{n \times m_1}$, $\mathbf{A}_2 \leftarrow \mathcal{R}_q^{n \times m_2}$, $\mathbf{B} \leftarrow \mathcal{R}_q^{\ell \times m_2}$. We observe that when $\ell = 0$ (resp. $m_1 = 0$) then this construction ends up being the Ajtai (resp. BDLOP) commitment scheme. In particular, the commitment size does not depend on the length m_1 of \mathbf{s}_1 (but it does on ℓ). Hence, our strategy is to commit to long vectors with small coefficients in the “Ajtai” part \mathbf{s}_1 and commit to a few *garbage* polynomials used for the proofs in the “BDLOP” part \mathbf{m} .

An opening of the commitment is a triple $(\mathbf{s}_1, \mathbf{m}, \mathbf{s}_2)$ ¹¹. As usual in lattice-based cryptography, we also consider relaxed openings of a commitment which are defined as follows.

Definition 4. *A relaxed opening of the ABDLOP commitment $(\mathbf{t}_A, \mathbf{t}_B)$ is a tuple $(\mathbf{s}_1, \mathbf{m}, \mathbf{s}_2, c)$ which satisfies:*

$$\begin{aligned} \mathbf{A}_1 \mathbf{s}_1 + \mathbf{A}_2 \mathbf{s}_2 &= \mathbf{t}_A \\ \mathbf{A}_2 \mathbf{s}_2 + \mathbf{m} &= \mathbf{t}_B \\ c &\in \bar{\mathcal{C}} \text{ as defined in Section 2.5} \\ \|\mathbf{c}\mathbf{s}_1\| &\leq B_1 \quad \text{and} \quad \|\mathbf{c}\mathbf{s}_2\| \leq B_2. \end{aligned}$$

As shown in [LNP22, Lemma 3.1], the ABDLOP commitment is binding with respect to relaxed openings under the Module-SIS assumption.

Lemma 6 ([LNP22]). *The ABDLOP commitment is computationally binding with respect to relaxed openings under the $\text{MSIS}_{n, m_1+m_2, B}$ assumption where $B := 4\eta\sqrt{B_1^2 + B_2^2}$.*

The hiding property of the ABDLOP commitment scheme follows from the fact that under the Module-LWE assumption that $\begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \mathbf{s}_2$ looks pseudorandom.

2.7 Framework for Proving Lattice Statements

The recently proposed framework by Lyubashevsky et al. [LNP22] can be used to prove various relations in the committed messages. Concretely, one can prove knowledge of the secret messages $(\mathbf{s}_1, \mathbf{m}) \in \mathcal{R}_q^{m_1+\ell}$ which satisfy all the following conditions:

1. *Quadratic relations over \mathcal{R}_q with automorphisms.* For $i \in [N]$ and public triples $(\mathbf{R}_{i,2}, \mathbf{r}_{i,1}, r_{i,0}) \in \mathcal{R}_q^{2(m_1+\ell) \times 2(m_1+\ell)} \times \mathcal{R}_q^{2(m_1+\ell)} \times \mathcal{R}_q$, we have:

$$\langle \mathbf{s}_1 \parallel \mathbf{m} \rangle_\sigma^T \mathbf{R}_{i,2} \langle \mathbf{s}_1 \parallel \mathbf{m} \rangle_\sigma + \mathbf{r}_{i,1}^T \langle \mathbf{s}_1 \parallel \mathbf{m} \rangle_\sigma + r_{i,0} = 0. \quad (15)$$

¹¹ Message \mathbf{m} does not need to be included in the opening since it can be deterministically computed from \mathbf{t}_B and \mathbf{s}_2 .

2. *Quadratic relations over \mathbb{Z}_q with automorphisms.* For $i \in [M]$ and public triples $(\mathbf{R}'_{i,2}, \mathbf{r}'_{i,1}, r'_{i,0}) \in \mathcal{R}_q^{2(m_1+\ell) \times 2(m_1+\ell)} \times \mathcal{R}_q^{2(m_1+\ell)} \times \mathcal{R}_q$:

$$\text{const. coeff. of } \langle \mathbf{s}_1 \parallel \mathbf{m} \rangle_\sigma^T \mathbf{R}'_{i,2} \langle \mathbf{s}_1 \parallel \mathbf{m} \rangle_\sigma + \mathbf{r}'_{i,1}^T \langle \mathbf{s}_1 \parallel \mathbf{m} \rangle_\sigma + r'_{i,0} \text{ equals } 0. \quad (16)$$

3. *Shortness in the infinity norm.* For public $\mathbf{P}_s \in \mathcal{R}_q^{n_{\text{bin}} \times m_1}$, $\mathbf{P}_m \in \mathcal{R}_q^{n_{\text{bin}} \times \ell}$ and $\mathbf{f} \in \mathcal{R}_q^{n_{\text{bin}}}$, the following polynomial vector has binary coefficients

$$\mathbf{P}_s \mathbf{s}_1 + \mathbf{P}_m \mathbf{m} + \mathbf{f} \in \{0, 1\}^{n_{\text{bin}} \cdot d}. \quad (17)$$

4. *Shortness in the Euclidean norm.* For $i \in [Z]$, public bound $\mathcal{B}_i < \sqrt{q}$ and $\mathbf{E}_s^{(i)} \in \mathcal{R}_q^{n_i \times m_1}$, $\mathbf{E}_m^{(i)} \in \mathcal{R}_q^{n_i \times \ell}$ and $\mathbf{v}^{(i)} \in \mathcal{R}_q^{n_i}$, we have:

$$\|\mathbf{E}_s^{(i)} \mathbf{s}_1 + \mathbf{E}_m^{(i)} \mathbf{m} + \mathbf{v}^{(i)}\| \leq \mathcal{B}_i.$$

This is equivalent to additionally proving knowledge of the binary polynomial $\vartheta_i \in \mathcal{R}$ such that

$$\langle \text{pow}(\mathcal{B}_i^2), \vartheta_i \rangle = \mathcal{B}_i^2 - \left\| \mathbf{E}_s^{(i)} \mathbf{s}_1 + \mathbf{E}_m^{(i)} \mathbf{m} + \mathbf{v}^{(i)} \right\|^2 \text{ over } \mathbb{Z} \quad (18)$$

where $\text{pow}(n) := \sum_{i=0}^{\lfloor \log n \rfloor} (2X)^i \in \mathcal{R}$ for $n \leq 2^{d-1}$.

3 Shorter Proofs via Bimodal Gaussians

In order to provide zero-knowledge (or more precisely, simulatability) for proving relations in the ABDLOP committed messages $(\mathbf{s}_1, \mathbf{m})$ under the randomness \mathbf{s}_2 , one applies the rejection sampling technique. In the original protocols presented in [LNP22], the standard rejection sampling [Lyu12] is used for \mathbf{s}_1 and the more recent one [LNS21a] for \mathbf{s}_2 . In this section we describe how one can apply bimodal Gaussian rejection sampling [DDLL13] on both the message and randomness which significantly reduces the standard deviations, and consequently the proof size, compared to [LNP22].

3.1 Bimodal Gaussian Rejection Sampling on the Randomness

In our constructions, we apply a rejection sampling procedure to mask a secret vector \vec{v} by first sampling \vec{y} from a discrete Gaussian with standard deviation \mathfrak{s} , and then computing $\vec{z} := \vec{v} + \vec{y}$. By Lemma 5, if we additionally run $\text{Rej}_1(\vec{z}, \vec{v}, \mathfrak{s})$, then the distribution of \vec{z} is indistinguishable to the one where we simply sample \vec{z} from a discrete Gaussian and output \vec{z} with certain (known) probability. Here, it is important that one could generate \vec{z} without having any information on \vec{v} .

Now, suppose that instead of Rej_1 , we run Rej_2 which is used for bimodal Gaussian rejection sampling [DDLL13]. It is now a natural question to ask whether there is a way to simulate the \vec{z} by having as little information on \vec{v} as possible. We answer this question positively and show that this distribution is simulatable given only the inner product $\langle \vec{z}, \vec{v} \rangle$ of \vec{z} and \vec{v} . We summarise our observation with the following lemma.

Lemma 7. Let $\vec{v} \in \mathbb{Z}^m$ be a vector of norm T . Fix $\mathfrak{s} \geq \gamma T$ and $M = \exp\left(\frac{1}{2\gamma^2}\right)$. Then the distributions of the outputs of $\mathcal{A}(\vec{v})$ and $\mathcal{F}(\vec{v})$ defined in Figure 5 are identical. Moreover, the probability that \mathcal{A} outputs something is exactly $1/M$.

$\mathcal{A}(\vec{v})$ 01 $\vec{y} \leftarrow D_{\mathfrak{s}}^m$ 02 $\vec{z} := \vec{y} + \vec{v}$ 03 output (\vec{z}, \vec{v}) with prob. $\frac{\exp\left(\frac{\ \vec{v}\ ^2}{2\mathfrak{s}^2}\right)}{M \cosh(\langle \vec{z}, \vec{v} \rangle / \mathfrak{s}^2)}$	$\mathcal{F}(\vec{v})$ 01 $\vec{y} \leftarrow D_{\mathfrak{s}}^m$ 02 $(\vec{z}_+, \vec{z}_-) := (\text{sign}(\langle \vec{y}, \vec{v} \rangle) \cdot \vec{y}, -\text{sign}(\langle \vec{y}, \vec{v} \rangle) \cdot \vec{y})$ 03 $p := \frac{\exp\left(\frac{2\langle \vec{y}, \vec{v} \rangle}{\mathfrak{s}^2}\right)}{\exp\left(\frac{2\langle \vec{y}, \vec{v} \rangle}{\mathfrak{s}^2}\right) + 1}$ 04 $\vec{z} := \begin{cases} \vec{z}_+ & \text{with prob. } p \\ \vec{z}_- & \text{with prob. } 1 - p \end{cases}$ 05 output (\vec{z}, \vec{v}) with prob. $\frac{1}{M}$
--	--

Fig. 5: Algorithms \mathcal{A} and \mathcal{F} for Lemma 7. We define $\text{sign}(x) = 1$ if $x \geq 0$ and -1 otherwise.

Proof. Fix $\vec{v} \in V$ and $\vec{z} \in \mathbb{Z}^m$ and let

$$p := \frac{\exp\left(\frac{2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right)}{\exp\left(\frac{2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right) + 1}.$$

By definition of \mathcal{A} , $\mathcal{A}(\vec{v}, \vec{z})$ is equal to

$$D_{\mathfrak{s}}^m(\vec{z} - \vec{v}) \cdot \frac{\exp\left(\frac{\|\vec{v}\|^2}{2\mathfrak{s}^2}\right)}{M \cosh\left(\frac{\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right)} = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{2 \exp\left(\frac{2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right)}{M \left(\exp\left(\frac{2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right) + 1\right)} = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{2p}{M}$$

Now, we focus on $\mathcal{F}(\vec{v})$. We see that by construction, $\langle \vec{z}_+, \vec{v} \rangle \geq 0$ and $\langle \vec{z}_-, \vec{v} \rangle \leq 0$. Let us consider three separate cases. First, suppose \vec{z} satisfies $\langle \vec{z}, \vec{v} \rangle > 0$. Informally, we want to compute the probability that $\vec{y} = \pm \vec{z}$ and \mathcal{F} picks \vec{z}_+ . Then,

$$\mathcal{F}(\vec{v}, \vec{z}) = 2D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{\exp\left(\frac{2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right)}{\exp\left(\frac{2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right) + 1} \cdot \frac{1}{M} = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{2p}{M}.$$

Further, suppose $\langle \vec{z}, \vec{v} \rangle < 0$. Informally, we compute the probability that $\vec{y} = \pm \vec{z}$ and \mathcal{F} picks \vec{z}_- . Then,

$$\mathcal{F}(\vec{v}, \vec{z}) = 2D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{1}{\exp\left(\frac{-2\langle \vec{z}, \vec{v} \rangle}{\mathfrak{s}^2}\right) + 1} \cdot \frac{1}{M} = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{2p}{M}.$$

Finally, assume $\langle \vec{z}, \vec{v} \rangle = 0$ and thus $p = 1/2$. Then, $\mathcal{F}(\vec{v}, \vec{z})$ is simply the probability that $(\vec{y} = \vec{z} \wedge \mathcal{F} \text{ outputs } \vec{z}_+)$ or $(\vec{y} = -\vec{z} \wedge \mathcal{F} \text{ outputs } \vec{z}_-)$. Hence,

$$\mathcal{F}(\vec{v}, \vec{z}) = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{1}{2M} + D_{\mathfrak{s}}^m(-\vec{z}) \cdot \frac{1}{2M} = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{1}{M} = D_{\mathfrak{s}}^m(\vec{z}) \cdot \frac{2p}{M}.$$

Therefore, we proved that for every \vec{z} , $\mathcal{A}(\vec{v}, \vec{z}) = \mathcal{F}(\vec{v}, \vec{z})$.

Finally, the second part of the statement follows from a simple observation that \mathcal{F} outputs something with probability exactly $1/M$. \square

Extended-MLWE Revisited. We observe that the only information about \vec{v} needed in order to run the simulator \mathcal{F} in the security proof is the value of $\langle \vec{y}, \vec{v} \rangle$. Hence, we reduce the simulatability property of our protocols to the hardness of the so-called Extended-MLWE. Here, as usual, an adversary needs to distinguish between the tuples $(\mathbf{B}, \mathbf{Bs})$ and (\mathbf{B}, \mathbf{u}) , where \mathbf{u} is a uniformly random vector but this time it is also given a “hint” of the form $(c, \mathbf{y}, \langle \mathbf{cs}, \mathbf{y} \rangle)$ where c and \mathbf{y} are sampled from some known distributions. For simplicity, we will describe the problem in a “knapsack” form.

Definition 5 (Extended-MLWE). *The Extended-MLWE problem with parameters n, m and distribution χ, ξ_c, ξ_y over \mathcal{R} asks the adversary \mathcal{A} to distinguish between the two cases: 1) $(\mathbf{B}, \mathbf{Bs}, c, \mathbf{y}, \langle \mathbf{cs}, \mathbf{y} \rangle)$ and 2) $(\mathbf{B}, \mathbf{u}, c, \mathbf{y}, \langle \mathbf{cs}, \mathbf{y} \rangle)$ for $\mathbf{B} \leftarrow \mathcal{R}_q^{m \times (n+m)}$, a secret vector $\mathbf{s} \leftarrow \chi^{n+m}$, uniformly random vector $\mathbf{u} \in \mathcal{R}_q^m$ and $(c, \mathbf{y}) \leftarrow \xi_c \times \xi_y^{n+m}$. Then, \mathcal{A} is said to have advantage ϵ in solving Extended-MLWE $_{n,m,\chi,\xi_c,\xi_y}$ if*

$$\left| \Pr \left[b = 1 \mid \mathbf{B} \leftarrow \mathcal{R}_q^{m \times (n+m)}; \mathbf{s} \leftarrow \chi^{n+m}; (c, \mathbf{y}) \leftarrow \xi_c \times \xi_y^{n+m}; b \leftarrow \mathcal{A}(\mathbf{B}, \mathbf{Bs}, c, \mathbf{y}, \langle \mathbf{cs}, \mathbf{y} \rangle) \right] \right. \\ \left. - \Pr \left[b = 1 \mid \mathbf{B} \leftarrow \mathcal{R}_q^{m \times (n+m)}; \mathbf{s} \leftarrow \chi^{n+m}; (c, \mathbf{y}) \leftarrow \xi_c \times \xi_y^{n+m}; \mathbf{u} \leftarrow \mathcal{R}_q^m; b \leftarrow \mathcal{A}(\mathbf{B}, \mathbf{u}, c, \mathbf{y}, \langle \mathbf{cs}, \mathbf{y} \rangle) \right] \right| \geq \epsilon.$$

We say that Extended-MLWE $_{n,m,\chi,\xi_c,\xi_y}$ is hard if for all PPT adversaries \mathcal{A} , the advantage in solving Extended-MLWE $_{n,m,\chi,\xi_c,\xi_y}$ is negligible.

We note that the (Module-)LWE problem with various side information has already been discussed in prior work e.g. [DGK⁺10, AP12, DDGR20]. As far as we are aware, this new variant of MLWE is the closest to the Extended Module-LWE problems defined by Lyubashevsky et al. [LNS21a], Alperin-Sheriff and Apon [ASA16], Alperin-Sheriff and Peikert [AP12] and Boudgoust et al. [BJRW21].

We observe that [ASA16] describes a similar problem with the two differences: (i) there is no c involved (assume that $c = 1$) and (ii) the hint is an arbitrary \mathbb{Q} -linear function on the “error” part \mathbf{e} of the secret \mathbf{s} (in particular it could be $\langle \mathbf{e}, \mathbf{y} \rangle \in \mathbb{Z}$ where $\mathbf{y} \leftarrow \xi_y^m$). Alperin-Sheriff and Apon show that their Extended-MLWE problem can be reduced to plain MLWE if the errors come from a discrete Gaussian with a large enough standard deviation. The proof strategy was later extended by Boudgoust et al. [BJRW21] who define another Extended-MLWE problem. This time, however, the hint becomes a whole polynomial $\langle \mathbf{e}, \mathbf{y} \rangle \in \mathcal{R}$. Finally, the only difference between our problem and the one in [LNS21a] is that the adversary is given the whole inner product $\langle \mathbf{cs}, \mathbf{y} \rangle$ instead of its sign.

If we consider our Extended-MLWE without any polynomial ring structure, then the problem becomes almost identical to the one introduced by Alperin-Sheriff and Peikert [AP12] (if we again assume $c = 1$). The authors additionally show that it is possible to reduce such a problem to plain LWE with the reduction loss $O(|\langle \vec{s}, \vec{y} \rangle|)$.

Applications. As an example, we show how to use the new rejection sampling strategy in the protocol for proving linear equations in the committed messages [LNP22][Section 3], however this approach can also be applied in all the protocols from [LNP22]. Let $(\mathbf{t}_A, \mathbf{t}_B)$ be the ABDLOP commitment to the message pair $(\mathbf{s}_1, \mathbf{m}) \in \mathcal{R}_q^{m_1} \times \mathcal{R}_q^\ell$ under randomness \mathbf{s}_2 , i.e.

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_B \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \cdot \mathbf{s}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \cdot \mathbf{s}_2 + \begin{bmatrix} \mathbf{0} \\ \mathbf{m} \end{bmatrix}. \quad (19)$$

Suppose the prover wants to prove knowledge of the message $(\mathbf{s}_1, \mathbf{m})$ such that

$$\mathbf{R}_1 \mathbf{s}_1 + \mathbf{R}_m \mathbf{m} = \mathbf{u}$$

where $\mathbf{R}_1 \in \mathcal{R}_q^{N \times m_1}$, $\mathbf{R}_m \in \mathcal{R}_q^{N \times \ell}$ and $\mathbf{u} \in \mathcal{R}_q^N$.

We present the commit-and-prove protocol in Figure 6 for proving linear relations. The only difference between this protocol and [LNP22, Fig. 4] is that for \mathbf{z}_2 we apply the new rejection sampling algorithm described above.

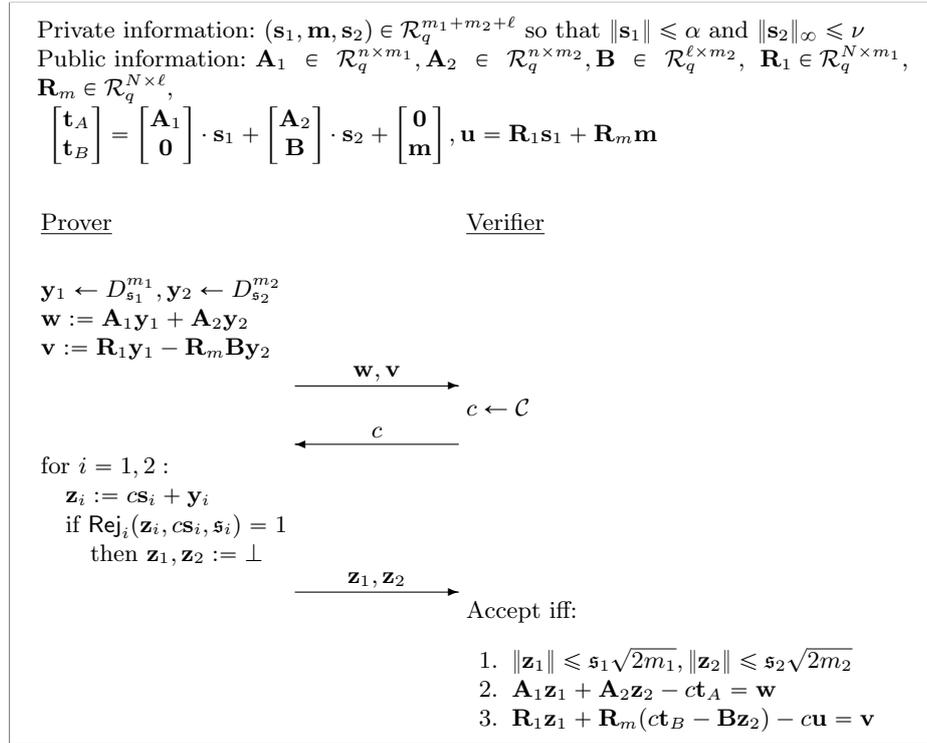


Fig. 6: Proof of knowledge $\Pi^{(1)}((\mathbf{s}_2, \mathbf{s}_1, \mathbf{m}), (f_1, f_2, \dots, f_N))$ of $(\mathbf{s}_1, \mathbf{s}_2, \bar{c}) \in \mathcal{R}_q^{m_1} \times \mathcal{R}_q^{m_2} \times \bar{\mathcal{C}}$ satisfying (i) $\mathbf{A}_1 \mathbf{s}_1 + \mathbf{A}_2 \mathbf{s}_2 = \mathbf{t}_A$, $\mathbf{B} \mathbf{s}_2 + \mathbf{m} = \mathbf{t}_B$ (ii) $\|\mathbf{s}_i \bar{c}\| \leq 2s_i \sqrt{2m_i d}$ for $i = 1, 2$ and (iii) $\mathbf{R}_1 \mathbf{s}_1 + \mathbf{R}_m \mathbf{m} = \mathbf{u}$. Functions Rej_i are defined in Fig. 3.

3.2 Bimodal Gaussian Rejection Sampling on the Message

This subsection focuses on applying bimodal Gaussian rejection sampling on the message vector \mathbf{s}_1 . First of all, we cannot apply Lemma 7 since it would potentially leak certain information about the message \mathbf{s}_1 which, unlike \mathbf{s}_2 , is not freshly sampled every time a new proof is generated. Instead, we follow the original methodology from [DDL13].

Concretely, let us focus on the protocol in Fig. 6. If one were to naively apply bimodal rejection sampling on $c\mathbf{s}_1$ then the masked opening of $c\mathbf{s}_1$ would become:

$$\mathbf{z}_1 := \mathbf{y}_1 + bc\mathbf{s}_1 \text{ where } b \leftarrow \{-1, 1\}.$$

As before, we set $\mathbf{z}_2 := \mathbf{y}_2 + c\mathbf{s}_2$. Hence, if we keep $\mathbf{w} := \mathbf{A}_1\mathbf{y}_1 + \mathbf{A}_2\mathbf{y}_2$ then by construction:

$$\mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 = \mathbf{w} + c(\mathbf{A}_1b\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2).$$

Note that $\mathbf{A}_1b\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2$ is a top part of the ABDLOP commitment to $(b\mathbf{s}_1, \mathbf{m})$ under randomness \mathbf{s}_2 . Thus, it is a natural approach to simply commit to $(b\mathbf{s}_1, \mathbf{m})$ and prove the quadratic equation. However, this comes with a big obstacle, i.e. we still need to prove the underlying relation in \mathbf{s}_1, \mathbf{m} even though we committed to $b\mathbf{s}_1$ and \mathbf{m} . It might cause a problem even in the simple case of linear relations. Indeed, initially we want to prove that $\mathbf{R}_1\mathbf{s}_1 + \mathbf{R}_m\mathbf{m} = \mathbf{u}$. Since we committed to $b\mathbf{s}_1$ and not \mathbf{s}_1 , it makes sense to try and prove the equivalent statement:

$$\mathbf{R}_1(b\mathbf{s}_1) + \mathbf{R}_m(b\mathbf{m}) = b\mathbf{u}. \quad (20)$$

This suggests that we should also commit to $b\mathbf{m}$ and not \mathbf{m} . However, it does not solve the issue completely since vector \mathbf{u} is still multiplied by a (secret) sign b . Hence, the intuitive solution would be to also commit to b in the ABDLOP commitment, prove $b \in \{-1, 1\}$ and the linear relation (20) in $b\mathbf{s}_1, b\mathbf{m}$ and b . Therefore, the cost of such an approach is at least committing to an extra polynomial.

We show that for certain types of statements we can circumvent committing to b and still apply bimodal Gaussian rejection sampling. Namely, we focus on *sign-invariant* relations.

Definition 6. Let $R \subseteq \{0, 1\}^* \times \mathcal{R}^{m_1+\ell}$ be a binary relation. We say that R is *sign-invariant* if for every pair (u, \mathbf{w}) we have: $R(u, \mathbf{w}) = 1 \iff R(u, -\mathbf{w}) = 1$.

Suppose we want to prove knowledge of $(\mathbf{s}_1, \mathbf{m}) \in \mathcal{R}_q^{m_1+\ell}$ such that $(u, (\mathbf{s}_1, \mathbf{m})) \in R$ where R is a sign-invariant relation. Then, we can sample a fresh sign $b \leftarrow \{-1, 1\}$ and commit to $(b\mathbf{s}_1, b\mathbf{m})$ using the ABDLOP commitment. Further, we simply prove that $R(u, (b\mathbf{s}_1, b\mathbf{m})) = 1$ which implies that $R(u, (\mathbf{s}_1, \mathbf{m})) = 1$.

Concrete instantiation. We demonstrate our intuition with the following example. Namely, we want to prove knowledge of $(\mathbf{s}_1, \mathbf{m})$ which satisfies:

$$\sigma(\mathbf{s}_1)^T \mathbf{s}_1 + \sigma(\mathbf{m})^T \mathbf{m} = 0.$$

Clearly, $(b\mathbf{s}_1, b\mathbf{m})$ satisfies the relation above for $b \in \{-1, 1\}$. As described before, we first sample a sign $b \leftarrow \{-1, 1\}$, randomness vector $\mathbf{s}_2 \leftarrow \chi$ and compute

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_B \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \cdot b\mathbf{s}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \cdot \mathbf{s}_2 + \begin{bmatrix} \mathbf{0} \\ b\mathbf{m} \end{bmatrix}. \quad (21)$$

Then, we simply follow the multiplicative proof from [LNP22, Section 4] to prove that

$$\sigma(b\mathbf{s}_1)^T(b\mathbf{s}_1) + \sigma(b\mathbf{m})^T(b\mathbf{m}) = 0.$$

Concretely, consider the *masked opening* $\mathbf{z}_1 := \mathbf{y}_1 + bcs_1$ of \mathbf{s}_1 . Note that

$$\sigma(\mathbf{z}_1)^T \mathbf{z}_1 = c^2 \sigma(b\mathbf{s}_1)^T(b\mathbf{s}_1) + c(\sigma(\mathbf{y}_1)^T(bs) + \mathbf{y}_1^T \sigma(bs)) + \sigma(\mathbf{y}_1)^T \mathbf{y}_1$$

and hence the coefficient corresponding to the quadratic term c^2 is what we are interested in. Here, we used the property of the challenge space \mathcal{C} that $c = \sigma(c)$ for $c \in \mathcal{C}$. We cannot do the same argument with $b\mathbf{m}$ since no masked opening of $b\mathbf{m}$ was sent. However, we observe that the verifier can compute $\mathbf{t}_B - \mathbf{B}\mathbf{z}_2 = -\mathbf{B}\mathbf{y}_2 + c(b\mathbf{m})$ which is of the similar form as the masked opening of $b\mathbf{s}_1$. Then

$$\begin{aligned} & \sigma(\mathbf{t}_B - \mathbf{B}\mathbf{z}_2)^T(\mathbf{t}_B - \mathbf{B}\mathbf{z}_2) \\ &= c^2(b\mathbf{m})^T(b\mathbf{m}) - c(\sigma(\mathbf{B}\mathbf{y}_2)^T(b\mathbf{m}) + (\mathbf{B}\mathbf{y}_2)^T \sigma(b\mathbf{m})) + \sigma(\mathbf{B}\mathbf{y}_2)^T \mathbf{B}\mathbf{y}_2. \end{aligned}$$

Therefore, we want to prove that the term in front of c^2 in the following expression disappears, i.e

$$\sigma(\mathbf{z}_1)^T \mathbf{z}_1 + \sigma(\mathbf{t}_B - \mathbf{B}\mathbf{z}_2)^T(\mathbf{t}_B - \mathbf{B}\mathbf{z}_2) = cg_1 + g_0$$

where

$$\begin{aligned} g_1 &:= \sigma(\mathbf{y}_1)^T(bs) + \mathbf{y}_1^T \sigma(bs) - \sigma(\mathbf{B}\mathbf{y}_2)^T(b\mathbf{m}) - (\mathbf{B}\mathbf{y}_2)^T \sigma(b\mathbf{m}) \\ g_0 &:= \sigma(\mathbf{y}_1)^T \mathbf{y}_1 + \sigma(\mathbf{B}\mathbf{y}_2)^T \mathbf{B}\mathbf{y}_2. \end{aligned} \quad (22)$$

The idea is then to additionally send a commitment $t = \mathbf{b}^T \mathbf{s}_2 + g_1$ to g_1 and send

$$v := g_0 + \mathbf{b}^T \mathbf{y}_2 = \sigma(\mathbf{y}_1)^T \mathbf{y}_1 + \sigma(\mathbf{B}\mathbf{y}_2)^T \mathbf{B}\mathbf{y}_2 + \mathbf{b}^T \mathbf{y}_2 \quad (23)$$

in the clear. Then, the verifier can check that:

$$v \stackrel{?}{=} \sigma(\mathbf{z}_1)^T \mathbf{z}_1 + \sigma(\mathbf{t}_B - \mathbf{B}\mathbf{z}_2)^T(\mathbf{t}_B - \mathbf{B}\mathbf{z}_2) + (\mathbf{b}^T \mathbf{z}_2 - ct). \quad (24)$$

We present the protocol for proving this relation in Fig. 7 and summarise its security properties in the full version of the paper.

Dealing with relations which are not sign-invariant. Typically, relations do not have the property that they are sign-invariant. In this case, to apply bimodal Gaussian rejection sampling on the message \mathbf{s}_1 one needs to be more careful. As hinted in the discussion above, one solution would be to commit to

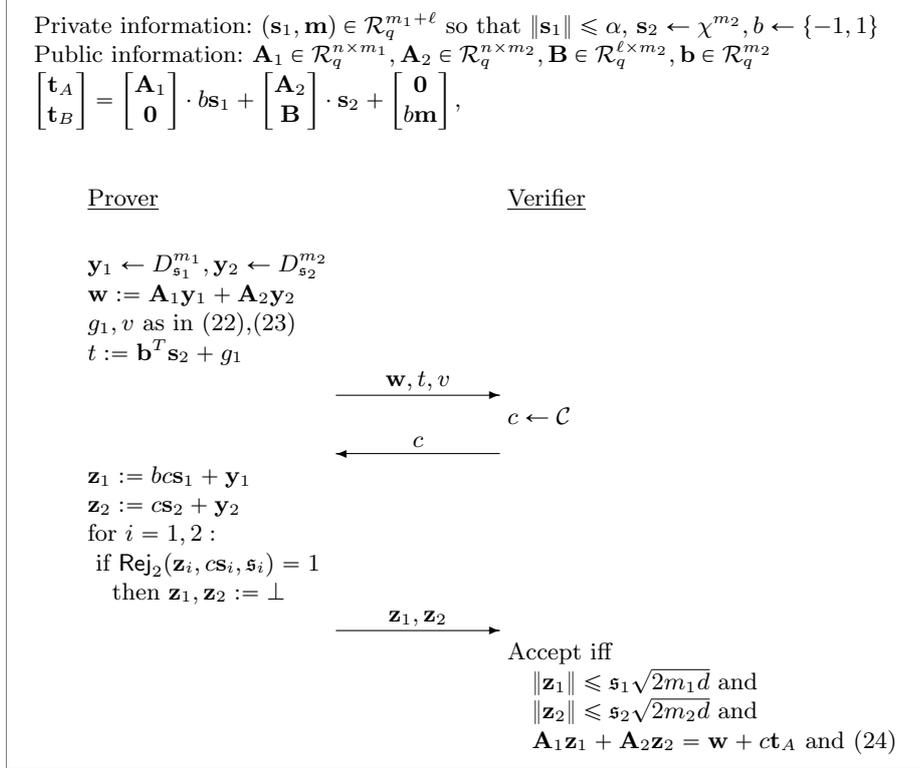


Fig. 7: Commit-and-prove protocol $\Pi_{\text{quad}}(\mathbf{s}_2, \mathbf{s}_1, \mathbf{m})$ for messages $(\mathbf{s}_1, \mathbf{m}) \in \mathcal{R}_q^{m_1+\ell}$, randomness $\mathbf{s}_2 \in \mathcal{R}_q^{m_2}$ and $\bar{c} \in \bar{\mathcal{C}}$ which satisfy: $\mathbf{A}_1\mathbf{s}_1 + \mathbf{A}_2\mathbf{s}_2 = \mathbf{t}_A$, $\mathbf{B}\mathbf{s}_2 + \mathbf{m} = \mathbf{t}_B$ (ii) $\|\mathfrak{s}_i\bar{c}\| \leq 2\mathfrak{s}_i\sqrt{2m_id}$ for $i = 1, 2$ and (iii) $\sigma(\mathbf{s}_1)^T\mathbf{s}_1 + \sigma(\mathbf{m})^T\mathbf{m} = 0$.

the sign b (in the BDLOP part of the ABDLOP commitment) and prove that $b \in \{-1, 1\}$ ¹². Then, for example, to prove an arbitrary quadratic equation with automorphisms (15), we commit to $(b\mathbf{s}_1, b\mathbf{m} \parallel b)$ and equivalently prove:

$$\langle b\mathbf{s}_1 \parallel b\mathbf{m} \rangle_{\sigma}^T \mathbf{R}_{i,2} \langle b\mathbf{s}_1 \parallel b\mathbf{m} \rangle_{\sigma} + b\mathbf{r}_{i,1}^T \langle b\mathbf{s}_1 \parallel b\mathbf{m} \rangle_{\sigma} + r_{i,0} = 0$$

which is a quadratic equation in $b\mathbf{s}_1, b\mathbf{m}$ and b . Then, in the soundness argument of [LNP22] we would extract $\bar{\mathbf{s}}_1, \bar{\mathbf{m}}$ and $\bar{b} \in \{-1, 1\}$ which satisfy

$$\langle \bar{\mathbf{s}}_1 \parallel \bar{\mathbf{m}} \rangle_{\sigma}^T \mathbf{R}_{i,2} \langle \bar{\mathbf{s}}_1 \parallel \bar{\mathbf{m}} \rangle_{\sigma} + \bar{b}\mathbf{r}_{i,1}^T \langle \bar{\mathbf{s}}_1 \parallel \bar{\mathbf{m}} \rangle_{\sigma} + r_{i,0} = 0.$$

Finally, since we proved that \bar{b} is a sign, we define $(\mathbf{s}_1^*, \mathbf{m}^*) := (\bar{b}\bar{\mathbf{s}}_1, \bar{b}\bar{\mathbf{m}})$ and deduce that

$$\langle \mathbf{s}_1^* \parallel \mathbf{m}^* \rangle_{\sigma}^T \mathbf{R}_{i,2} \langle \mathbf{s}_1^* \parallel \mathbf{m}^* \rangle_{\sigma} + \mathbf{r}_{i,1}^T \langle \mathbf{s}_1^* \parallel \mathbf{m}^* \rangle_{\sigma} + r_{i,0} = 0$$

which is what we wanted to extract at the very beginning.

¹² Proving that b is a sign has already been covered in [LNP22, Section 5.1].

Similarly, to prove (17), we want to prove instead

$$\mathbf{P}_s(b\mathbf{s}_1) + \mathbf{P}_m(b\mathbf{m}) + b\mathbf{f} \in \{0, b\}^{n_{\text{bin}}d}.$$

We observe that $x \in \{0, b\} = \{\frac{b-1}{2}, \frac{b+1}{2}\}$ if and only if $x - \frac{b-1}{2} \in \{0, 1\}$. Hence, the statement above is equivalent to:

$$\mathbf{P}_s(b\mathbf{s}_1) + \mathbf{P}_m(b\mathbf{m}) + b(\mathbf{f} - 2^{-1} \cdot \mathbf{1}) + 2^{-1} \cdot \mathbf{1} \in \{0, 1\}^{n_{\text{bin}}d}$$

where $\mathbf{1} \in \mathcal{R}_q^{n_{\text{bin}}}$ is the polynomial vector with the coefficient vector $\vec{\mathbf{1}}$. Hence, we reduced our problem to proving that a linear combination of $b\mathbf{s}_1$, $b\mathbf{m}$ and b has binary coefficients. We conclude that using similar techniques, one can transform all the relations in \mathbf{s}_1, \mathbf{m} described in Section 2.7 to equivalent ones in $b\mathbf{s}_1, b\mathbf{m}$ and $b \in \{-1, 1\}$.

4 Efficient One-out-of-Many Proofs

In this section we construct an efficient logarithmic-size one-out-of-many proof [GK15] with applications to lattice-based ring and group signatures using techniques from [LNP22] as the building block. In the full version of the paper we show how to further reduce the proof size using the techniques developed in Section 3, and eventually describe our ring signature construction.

The one-out-of-many proof considers the following problem. Informally, we want to prove knowledge of an opening to some commitment contained in a public set S without revealing any information about the commitment itself. In the lattice setting, we would like to prove knowledge of a short vector such that $\mathbf{A}\mathbf{s} \in S$, where S is a public set $S = \{\mathbf{t}_1, \dots, \mathbf{t}_N\} \subseteq \mathcal{R}_q^n$ of size $N = d \cdot \mathfrak{d}^k$. In this section we assume that $\mathbf{s} \in \{0, 1\}^{md}$ has binary coefficients and $d = \mathfrak{l} \cdot \mathfrak{d}$ for $\mathfrak{l} \in \mathbb{N}$. For simplicity, we can already instantiate some of these parameters as $(d, \mathfrak{d}, \mathfrak{l}) = (64, 8, 16)$.

We now use the observation from [ESS⁺19, GK15, BCC⁺15] that $\mathbf{A}\mathbf{s} \in S$ if and only if there exists a binary vector $\vec{v} \in \{0, 1\}^N$ with exactly one 1, i.e. a unit vector, such that

$$[\vec{t}_1 \ \vec{t}_2 \ \dots \ \vec{t}_N] \vec{v} = A\vec{s} \quad (25)$$

where $A = \text{rot}(\mathbf{A}) \in \mathbb{Z}_q^{nd \times md}$ is the the rotation matrix of \mathbf{A} . One could then directly prove knowledge of \vec{s} and \vec{v} which satisfy conditions above using the protocol from Section 2.7. However, the proof size grows linearly in N since we would commit to the whole vector \vec{v} .

In order to circumvent this limitation, [GK15, BCC⁺15] observe that vector \vec{v} can be uniquely decomposed into unit vectors $\vec{v}_1, \dots, \vec{v}_k \in \{0, 1\}^{\mathfrak{d}}$ and $\vec{v}_{k+1} \in \{0, 1\}^d$ such that

$$\vec{v} = \vec{v}_1 \otimes \vec{v}_2 \otimes \dots \otimes \vec{v}_{k+1} := \vec{v}_1 \otimes (\vec{v}_2 \otimes (\dots \otimes (\vec{v}_k \otimes \vec{v}_{k+1}))). \quad (26)$$

For notational convenience, let us define the set of polynomials \mathcal{X} in \mathcal{R}_q with their coefficient vectors being a unit vector. Concretely, \mathcal{X} is defined as follows:

$$\mathcal{X} := \{1, X, X^2, \dots, X^{d-1}\}.$$

In the end, we want to commit to \mathbf{s} and polynomials $u_1, \dots, u_k, v_{k+1} \in \mathcal{X}$ such that $\vec{u}_i = \vec{v}_i \parallel 0^{d-\mathfrak{d}} \in \mathbb{Z}_q^{d \cdot 13}$ for $i \in [k]$ and prove

$$T(\vec{v}_1 \otimes \dots \otimes \vec{v}_{k+1}) = A\vec{s} \quad (27)$$

where $T \in \mathbb{Z}_q^{nd \times N}$ is the matrix on the left-hand side of (25). We formally define the corresponding relation:

$$R_{\text{oom}} := \left\{ ((T, A), (\mathbf{s}, u_1, \dots, u_k, v_{k+1})) : \mathbf{s} \in \{0, 1\}^{md} \wedge T(\vec{v}_1 \otimes \dots \otimes \vec{v}_{k+1}) = A\vec{s} \right. \\ \left. \wedge u_1, \dots, u_k, v_{k+1} \in \mathcal{X} \text{ where } \vec{u}_i := \vec{v}_i \parallel 0^{d-\mathfrak{d}} \right\}.$$

We now describe a commit-and-prove system for relation R_{oom} using the ABDLOP commitment. Suppose that $k \geq 1$, otherwise one can prove this relation directly using the framework from [LNP22].

First, note that proving $u_1, \dots, u_k, v_{k+1} \in \mathcal{X}$ and $\mathbf{s} \in \{0, 1\}^{md}$ can be done directly using the techniques from Section 2.7 hence we focus first on (27). Our strategy to prove this equation with $k - 1$ tensor products would be somehow to reduce it to proving an equation of the same form with only $k - 2$ tensor products. Then, by recursion, we will end up with a system of linear equations with no tensor products involved and thus we can apply the methods presented in Section 2.7.

The key idea to reduce the number of tensor products is to ask the verifier for \mathfrak{l} challenges $\vec{\varphi}_1, \dots, \vec{\varphi}_{\mathfrak{l}} \in \mathbb{Z}_q^{nd}$ and then prove that:

$$\langle T(\vec{v}_1 \otimes \dots \otimes \vec{v}_{k+1}) - A\vec{s}, \vec{\varphi}_i \rangle = 0 \quad \text{for } i = 1, 2, \dots, \mathfrak{l}.$$

Note that if (27) was not true, then these \mathfrak{l} equations above would hold with probability at most $q_1^{-\mathfrak{l}}$. Now, if we write

$$T := [T_{0,1} \ T_{0,2} \ \dots \ T_{0,\mathfrak{d}}] \quad \text{where each } T_{0,i} \in \mathbb{Z}_q^{nd \times d\mathfrak{d}^{k-1}}$$

then by simple algebraic manipulation we obtain

$$\begin{aligned} \langle T(\vec{v}_1 \otimes \dots \otimes \vec{v}_{k+1}) - A\vec{s}, \vec{\varphi}_i \rangle &= \langle \vec{v}_1 \otimes \dots \otimes \vec{v}_{k+1}, \vec{\varphi}_i^T T \rangle - \langle \vec{s}, A^T \vec{\varphi}_i \rangle \\ &= \langle \vec{v}_1, T_{1,i}(\vec{v}_2 \otimes \dots \otimes \vec{v}_{k+1}) \rangle - \langle \vec{s}, A^T \vec{\varphi}_i \rangle \end{aligned}$$

where

$$T_{1,i} := \begin{bmatrix} \vec{\varphi}_i^T T_{0,1} \\ \vdots \\ \vec{\varphi}_i^T T_{0,\mathfrak{d}} \end{bmatrix} \in \mathbb{Z}_q^{\mathfrak{d} \times d\mathfrak{d}^{k-1}}.$$

Now, let us define $\vec{w}_i := T_{1,i}(\vec{v}_2 \otimes \dots \otimes \vec{v}_{k+1})$ and $w \in \mathcal{R}_q$ such that

$$\vec{w} = \vec{w}_1 \parallel \dots \parallel \vec{w}_{\mathfrak{l}} \in \mathbb{Z}_q^d.$$

¹³ Alternatively, $u_i \in \{1, X, X^2, \dots, X^{\mathfrak{d}-1}\}$.

Next, we commit to w and show that for all i ,

$$\langle \vec{v}_1, \vec{w}_i \rangle - \langle \vec{s}, A^T \vec{\varphi}_i \rangle = 0 \quad \text{and} \quad \vec{w}_i := T_{1,i}(\vec{v}_2 \otimes \cdots \otimes \vec{v}_{k+1}).$$

We observe that the first statement is equivalent to proving that the constant coefficient of

$$X^{(i-1)\mathfrak{d}} u_1 \sigma(w) - \sigma(\mathbf{a}_i)^T \mathbf{s}$$

is equal to zero where the coefficient vector of $\mathbf{a}_i \in \mathcal{R}_q^m$ is exactly $\vec{a}_i := A^T \vec{\varphi}_i$.

Lemma 8. *Let $i \in [l]$. Then, the constant coefficient of $X^{(i-1)\mathfrak{d}} u_1 \sigma(w) \in \mathcal{R}_q$ is equal to $\langle \vec{v}_1, \vec{w}_i \rangle$.*

Proof. First, we note that $\langle \vec{v}_1, \vec{w}_i \rangle = \langle X^{(i-1)\mathfrak{d}} u_1, w \rangle$. Here, we used the fact that the coefficient vector of u_1 is of the form $\vec{v}_1 \parallel 0^{d-\mathfrak{d}}$. Then, by Lemma 2, $\langle X^{(i-1)\mathfrak{d}} u_1, w \rangle$ is the constant coefficient of $X^{(i-1)\mathfrak{d}} u_1 \sigma(w)$.

On the other hand, the second statement can be combined for all i and written as:

$$\vec{w} = \begin{bmatrix} T_{1,1} \\ \vdots \\ T_{1,l} \end{bmatrix} (\vec{v}_2 \otimes \cdots \otimes \vec{v}_{k+1}). \quad (28)$$

Thus, we reduce the one-out-of-many problem to proving knowledge of a tuple $(\mathbf{s}, u_1, \dots, u_k, v_{k+1}, w)$ which satisfies the following conditions: (i) $\mathbf{s} \in \{0, 1\}^{md}$, (ii) $T_1(\vec{v}_2 \otimes \cdots \otimes \vec{v}_{k+1}) = \vec{w}$, (iii) for all $i \in [l]$, the constant coefficient of $X^{(i-1)\mathfrak{d}} u_1 \sigma(w) - \sigma(\mathbf{a}_i)^T \mathbf{s}$ is zero, and (iv) $u_1, \dots, u_k, v_{k+1} \in \mathcal{X}$ where

$$\vec{u}_i := \vec{v}_i \parallel 0^{d-\mathfrak{d}} \text{ for } i \in [k] \quad \text{and} \quad T_1 := \begin{bmatrix} T_{1,1} \\ \vdots \\ T_{1,l} \end{bmatrix} \in \mathbb{Z}_q^{d \times d \mathfrak{d}^{k-1}}.$$

Note that the second statement only involves $k-2$ tensor products.

We can define the correspond relation as:

$$R := \left\{ \begin{array}{l} ((T_1, (\mathbf{a}_i)_{i \in [l]}), (\mathbf{s}, u_1, \dots, u_k, v_{k+1}, w)) : \mathbf{s} \in \{0, 1\}^{md} \wedge T_1(\vec{v}_2 \otimes \cdots \otimes \vec{v}_{k+1}) = \vec{w} \\ \wedge \forall i \in [l], \text{ const coeff. of } X^{(i-1)\mathfrak{d}} u_1 \sigma(w) - \sigma(\mathbf{a}_i)^T \mathbf{s} \text{ is zero} \\ \wedge u_1, \dots, u_k, v_{k+1} \in \mathcal{X} \text{ where } \vec{u}_i := \vec{v}_i \parallel 0^{d-\mathfrak{d}} \end{array} \right\}.$$

Intermediate relations. We construct a commit-and-prove system for relation R using recursion. Namely, take $1 \leq j \leq k$ and consider the following generalised relation

$$R_j := \left\{ \begin{array}{l} \left((T_j \in \mathbb{Z}_q^{d \times d \mathfrak{d}^{k-j}}, (\mathbf{a}_i)_{i \in [l]}, (\varphi_{\ell,i})_{\ell \in [j-1], i \in [l]}), (\mathbf{s}, u_1, \dots, u_k, v_{k+1}, w_1, \dots, w_j) \right) : \\ \mathbf{s} \in \{0, 1\}^{md} \wedge T_j(\vec{v}_{j+1} \otimes \cdots \otimes \vec{v}_{k+1}) = \vec{w}_j \\ \wedge \forall i \in [l], \text{ const coeff. of } X^{(i-1)\mathfrak{d}} u_1 \sigma(w_1) - \sigma(\mathbf{a}_i)^T \mathbf{s} \text{ is zero} \\ \wedge \forall \ell \in [j-1], i \in [l], \text{ const coeff. of } X^{(i-1)\mathfrak{d}} u_{\ell+1} \sigma(w_{\ell+1}) - \sigma(\varphi_{\ell,i}) w_\ell \text{ is zero} \\ \wedge u_1, \dots, u_k, v_{k+1} \in \mathcal{X} \text{ where } \vec{u}_i := \vec{v}_i \parallel 0^{d-\mathfrak{d}} \end{array} \right\}. \quad (29)$$

We highlight that in R_j elements $\varphi_{\ell,i}$ are polynomials in \mathcal{R}_q . Also, it is easy to see that $R_1 = R$.

Base case. We first show how to prove R_k only using the methods described in Section 2.7. In the following, we say that a statement is of Type- n if it corresponds to the Statement n in Section 2.7.

To begin with, using the ABDLOP commitment we commit to

$$\mathbf{s}_1 := \mathbf{s} \parallel u_1 \parallel \cdots \parallel u_k \parallel v_{k+1} \in \mathcal{R}_q^{m+k+1}, \quad \mathbf{m} := (w_1, \dots, w_k) \in \mathcal{R}_q^k.$$

Then, proving $\mathbf{s} \in \{0, 1\}^{md}$ and $u_1, \dots, u_k, v_{k+1} \in \{0, 1\}^d$ is of Type-3. Next, by Lemma 2, proving $T_k \vec{v}_{k+1} = \vec{w}_k$ and $\langle \vec{1}, \vec{v}_i \rangle = 1$ for $i \in [k+1]$ is of Type-2. Further, it is easy to see that proving the constant coefficients of $X^{(i-1)\mathfrak{d}} u_1 \sigma(w_1) - \sigma(\mathbf{a}_i)^T \mathbf{s}$ and $X^{(i-1)\mathfrak{d}} u_{i+1} \sigma(w_{i+1}) - \sigma(\varphi_{i,i}) w_i$ vanish is of Type-2. Finally, proving that $\vec{u}_i = \vec{v}_i \parallel 0^{d-\mathfrak{d}}$ for $i \in [k]$ is equivalent to proving that the constant coefficient of $X^{-j} u_i$ is zero for $\mathfrak{d} \leq j \leq d$, which is of Type-2.

From now on, we will call the commit-and-prove protocol for relation R_k described above as Π_k .

Recursive step. Let us assume we have a commit-and-prove system Π_{j+1} for relation R_{j+1} where $2 \leq j+1 \leq k$. Now we want to use it to prove relation R_j . We observe that the only statement which is included in R_j but not in R_{j+1} is

$$T_j(\vec{v}_{j+1} \otimes \cdots \otimes \vec{v}_{k+1}) = \vec{w}_j. \quad (30)$$

We prove this equation as before. Namely, we ask the verifier for \mathfrak{l} challenges $\vec{\varphi}_{j,1}, \dots, \vec{\varphi}_{j,\mathfrak{l}} \in \mathbb{Z}_q^d$ and then prove that:

$$\langle T_j(\vec{v}_{j+1} \otimes \cdots \otimes \vec{v}_{k+1}) - \vec{w}_j, \vec{\varphi}_{j,i} \rangle = 0 \quad \text{for } i = 1, 2, \dots, \mathfrak{l}.$$

Note that if (27) was not true, then these \mathfrak{l} equations above would hold with probability at most $q_1^{-\mathfrak{l}}$. Now, if we write

$$T_j := [T_{j,1} \ T_{j,2} \ \cdots \ T_{j,\mathfrak{d}}] \quad \text{where each } T_{j,i} \in \mathbb{Z}_q^{d \times d\mathfrak{d}^{k-j-1}}$$

then we have

$$\begin{aligned} \langle T_j(\vec{v}_{j+1} \otimes \cdots \otimes \vec{v}_{k+1}) - \vec{w}_j, \vec{\varphi}_{j,i} \rangle &= \langle \vec{v}_{j+1} \otimes \cdots \otimes \vec{v}_{k+1}, \vec{\varphi}_{j,i}^T T_j \rangle - \langle \vec{w}_j, \vec{\varphi}_{j,i} \rangle \\ &= \langle \vec{v}_{j+1} \otimes \cdots \otimes \vec{v}_{k+1}, \vec{\varphi}_{j,i}^T T_j \rangle - \langle \vec{w}_j, \vec{\varphi}_{j,i} \rangle \\ &= \langle \vec{v}_{j+1}, T_{j+1,i}(\vec{v}_{j+2} \otimes \cdots \otimes \vec{v}_{k+1}) \rangle - \langle \vec{w}_j, \vec{\varphi}_{j,i} \rangle \end{aligned}$$

where

$$T_{j+1,i} := \begin{bmatrix} \vec{\varphi}_{j,i}^T T_{j,1} \\ \vdots \\ \vec{\varphi}_{j,i}^T T_{j,\mathfrak{d}} \end{bmatrix} \in \mathbb{Z}_q^{\mathfrak{d} \times \mathfrak{d}^{k-j-1}}.$$

Now, let us define $\vec{w}_{j+1,i} := T_{j+1,i}(\vec{v}_{j+2} \otimes \cdots \otimes \vec{v}_{k+1})$ and $w_{j+1} \in \mathcal{R}_q$ so that

$$\vec{w}_{j+1} = \vec{w}_{j+1,1} \parallel \cdots \parallel \vec{w}_{j+1,\mathfrak{l}} \in \mathbb{Z}_q^d.$$

Then, we need to show that for all i ,

$$\langle \vec{w}_{j+1}, \vec{w}_{j+1,i} \rangle - \langle \vec{w}_j, \vec{\varphi}_{j,i} \rangle = 0 \quad \text{and} \quad \vec{w}_{j+1,i} = T_{j+1,i}(\vec{v}_{j+2} \otimes \cdots \otimes \vec{v}_{k+1}).$$

The first statement is equivalent to proving that the constant coefficient of

$$X^{(i-1)\mathfrak{d}} u_{j+1} \sigma(w_{j+1}) - \sigma(\varphi_{j,i}) w_j$$

is equal to zero. The second statement, however, can be combined for all i and written as:

$$\vec{w}_{j+1} = T_{j+1}(\vec{v}_{j+2} \otimes \cdots \otimes \vec{v}_{k+1}) \quad \text{where} \quad T_{j+1} := \begin{bmatrix} T_{j+1,1} \\ \vdots \\ T_{j+1,\ell} \end{bmatrix} \in \mathbb{Z}_q^{d \times \mathfrak{d}^{k-j-1}}. \quad (31)$$

Therefore, we reduced proving (30) to proving that

$$\begin{aligned} & - X^{(i-1)\mathfrak{d}} u_{j+1} \sigma(w_{j+1}) - \sigma(\varphi_{j,i}) w_j \text{ is equal to zero} \\ & - \vec{w}_{j+1} = T_{j+1}(\vec{v}_{j+2} \otimes \cdots \otimes \vec{v}_{k+1}) \end{aligned}$$

which in combination with other relations in R_j , it directly reduces to proving relations in R_{j+1} .

In Fig. 8 we give a commit-and-prove protocol for relation R_j which uses Π_j as a black-box. One observes by discussion above that the correctness error for Π_j is the same as for Π_k (which can be calculated directly from [LNP22]). Simulatability follows from the fact that running Π_j , and thus Π_{j+1} up to Π_k as subroutines, involves only sending intermediate commitments t_i to w_i which can be simulated by the Extended-MLWE assumption. Finally, one can prove by induction that the knowledge soundness error for the protocol Π_j is $(k-j) \cdot q_1^{-l} + \varepsilon_k$ where ε_k is the knowledge soundness error for Π_k and is computed as in [LNP22]. The expected runtime of the extractor, that has black-box access to a (potentially malicious) prover which runs in time T , is $2^{k-j} \cdot \text{poly}(T)$. Consequently, we can only consider values k which are logarithmic in the security parameter. Due to space constraints, we refer to the full version of the paper for more details.

Back to R_{oom} . At the very beginning of this section we showed how to reduce proving relation R_{oom} to proving R_1 . Later on, we described a commit-and-prove protocol for R_1 . Hence, we combine these two results to obtain a commit-and-prove protocol Π_{oom} in Fig. 9 for the one-out-of-many relation R_{oom} . Arguing similarly as above, the correctness error for Π_{oom} is the same as for Π_k and the soundness error is at most $kq_1^{-l} + \varepsilon_k$.

References

- AFK21. Thomas Attema, Serge Fehr, and Michael Kloof. Fiat-shamir transformation of multi-round interactive proofs. Cryptology ePrint Archive, Paper 2021/1377, 2021. <https://eprint.iacr.org/2021/1377>.

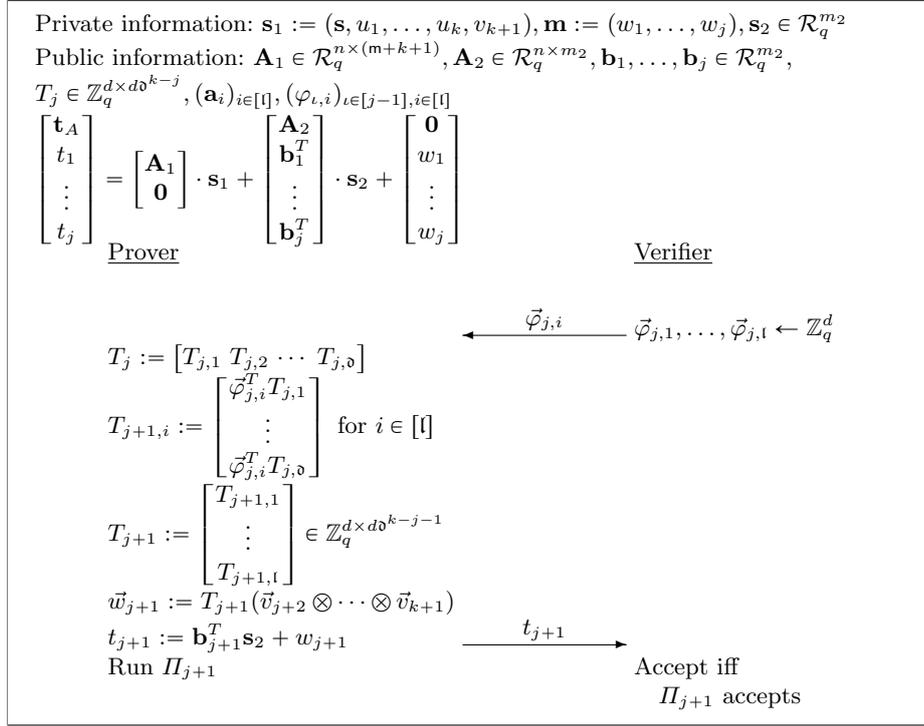


Fig. 8: Commit-and-prove protocol Π_j for the relation R_j where $j < k$.

- Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
- ALS20. 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.
- AP12. Jacob Alperin-Sheriff and Chris Peikert. Circular and KDM security for identity-based encryption. In *Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 334–352. Springer, 2012.
- ASA16. Jacob Alperin-Sheriff and Daniel Apon. Dimension-preserving reductions from lwe to lwr. *Cryptology ePrint Archive*, Report 2016/589, 2016. <https://ia.cr/2016/589>.
- Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, Dec 1993.
- BCC⁺15. 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.
- BDE⁺22. Maxime Buser, Rafael Dowsley, Muhammed F. Esgin, Clémentine Gritti, Shabnam Kasra Kermanshahi, Veronika Kuchta, Jason T. LeGrow, Joseph K. Liu, Raphael C.-W. Phan, Amin Sakzad, Ron Steinfeld, and Jiangshan Yu. A survey on exotic signatures for post-quantum blockchain: Challenges & research directions. *Cryptology ePrint Archive*, Paper 2022/1151, 2022. <https://eprint.iacr.org/2022/1151>.

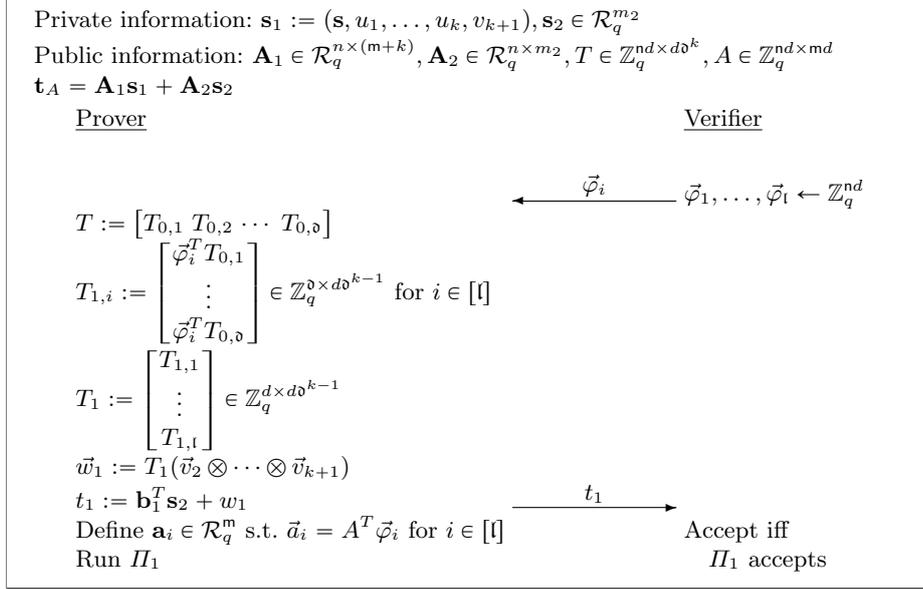


Fig. 9: Commit-and-prove protocol Π_{oom} for the relation R_{oom} .

- BDK⁺21. Ward Beullens, Samuel Dobson, Shuichi Katsumata, Yi-Fu Lai, and Federico Pintore. Group signatures and more from isogenies and lattices: Generic, simple, and efficient. *IACR Cryptol. ePrint Arch.*, page 1366, 2021.
- BDL⁺18. Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In *SCN*, pages 368–385, 2018.
- BJRW21. Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. On the hardness of module-lwe with binary secret. In *CT-RSA*, volume 12704 of *Lecture Notes in Computer Science*, pages 503–526. Springer, 2021.
- BKP20. Ward Beullens, Shuichi Katsumata, and Federico Pintore. Calamari and falaf: Logarithmic (linkable) ring signatures from isogenies and lattices. In *ASIACRYPT (2)*, volume 12492 of *Lecture Notes in Computer Science*, pages 464–492. Springer, 2020.
- BLS19. 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.
- DDGR20. Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi. LWE with side information: Attacks and concrete security estimation. In *CRYPTO (2)*, volume 12171 of *Lecture Notes in Computer Science*, pages 329–358. Springer, 2020.
- DDLL13. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO (1)*, pages 40–56, 2013.
- DGK⁺10. Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Public-key encryption schemes with auxiliary in-

- puts. In Daniele Micciancio, editor, *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, 2010.
- DLP14. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In *ASIACRYPT*, pages 22–41, 2014.
- DP16. Léo Ducas and Thomas Prest. Fast fourier orthogonalization. In *ISSAC*, pages 191–198, 2016.
- ENS20. Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In *ASIACRYPT (2)*, pages 259–288, 2020.
- ESLL19. 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)*, pages 115–146. Springer, 2019.
- ESS⁺19. 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.
- ESZ21. Muhammed F. Esgin, Ron Steinfeld, and Raymond K. Zhao. Matric+: More efficient post-quantum private blockchain payments. *IACR Cryptol. ePrint Arch.*, page 545, 2021.
- EZS⁺19. Muhammed F. Esgin, Raymond K. Zhao, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. Matric: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In *CCS*, pages 567–584. ACM, 2019.
- FHK⁺20. Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Prest, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, , and Zhenfei Zhang. Falcon: Fast-fourier lattice-based compact signatures over ntru. Technical report, 2020. <https://falcon-sign.info/falcon.pdf>.
- FS86. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
- GK15. 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.
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- HHGP⁺03. Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. Ntrusign: Digital signatures using the ntru lattice. In *CT-RSA*, pages 122–140, 2003.
- LAZ19. 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.
- LMPR08. Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72, 2008.
- LNP22. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. *Cryptology ePrint Archive*, Paper 2022/284, 2022. <https://eprint.iacr.org/2022/284>. To appear at CRYPTO 2022.
- LNPS21. Vadim Lyubashevsky, Ngoc Khanh Nguyen, Maxime Plançon, and Gregor Seiler. Shorter lattice-based group signatures via ”almost free” encryption

- and other optimizations. In *ASIACRYPT (4)*, pages 218–248. Springer, 2021.
- LNS20. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Practical lattice-based zero-knowledge proofs for integer relations. In *CCS*, pages 1051–1070. ACM, 2020.
- LNS21a. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Shorter lattice-based zero-knowledge proofs via one-time commitments. In *Public Key Cryptography (1)*, pages 215–241. Springer, 2021.
- LNS21b. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. SMILE: set membership from ideal lattices with applications to ring signatures and confidential transactions. In *CRYPTO (2)*, volume 12826 of *Lecture Notes in Computer Science*, pages 611–640. Springer, 2021.
- LNSW13. San Ling, Khoa Nguyen, Damien Stehlé, and Huaxiong Wang. Improved zero-knowledge proofs of knowledge for the ISIS problem, and applications. In *PKC*, pages 107–124, 2013.
- LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015.
- Lyu09. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
- Lyu12. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
- NY90. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, pages 427–437, 1990.
- TWZ20. Yang Tao, Xi Wang, and Rui Zhang. Short zero-knowledge proof of knowledge for lattice-based commitment. In *PQCrypto*, volume 12100 of *Lecture Notes in Computer Science*, pages 268–283. Springer, 2020.
- YAZ⁺19. 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)*, pages 147–175. Springer, 2019.
- YEL⁺21. Tsz Hon Yuen, Muhammed F. Esgin, Joseph K. Liu, Man Ho Au, and Zhimin Ding. Dualring: Generic construction of ring signatures with efficient instantiations. In *CRYPTO (1)*, volume 12825 of *Lecture Notes in Computer Science*, pages 251–281. Springer, 2021.