



Shorter Lattice-Based Zero-Knowledge Proofs via One-Time Commitments

Vadim Lyubashevsky^{1(✉)}, Ngoc Khanh Nguyen^{1,2}, and Gregor Seiler^{1,2}

¹ IBM Research - Zurich, Ruschlikon, Switzerland

{vad,nkn}@zurich.ibm.com

² ETH Zurich, Zurich, Switzerland

gseiler@inf.ethz.ch

Abstract. There has been a lot of recent progress in constructing efficient zero-knowledge proofs for showing knowledge of an \vec{s} with small coefficients satisfying $A\vec{s} = \vec{t}$. For typical parameters, the proof sizes have gone down from several megabytes to a bit under 50KB (Esgin et al., Asiacrypt 2020). These are now within an order of magnitude of the sizes of lattice-based signatures, which themselves constitute proof systems which demonstrate knowledge of something weaker than the aforementioned equation. One can therefore see that this line of research is approaching optimality. In this paper, we modify a key component of these proofs, as well as apply several other tweaks, to achieve a further reduction of around 30% in the proof output size. We also show that this savings propagates itself when these proofs are used in a general framework to construct more complex protocols.

1 Introduction

Zero-knowledge proofs and commit-and-prove protocols form the foundations of virtually all privacy-based protocols. In preparing for the (eventual) coming of quantum computing, there has been a lot of focus in recent years on building such protocols based on quantum-safe assumptions. Quantum-safe PCP/IOP schemes whose security is just based on the collision-resistance of cryptographic hash functions have existed for decades, and there has been a lot of recent work around them. The main feature of these constructions is that their outputs are sublinear in the statement size. The main downside is that they can be very slow and memory intensive. Furthermore, there is a lower bound of around 100KB for proofs of statements that have small size. So it's quite likely that they are not the best solution for all scenarios.

In the last few years, new techniques emerged in lattice cryptography that made it a promising foundation for quantum-safe privacy-based protocols. These techniques have been used for blockchain applications [EZS+19], verifiable random functions [EKS+20], and for proofs of arithmetic statements [LNS20a]. In all of these scenarios, lattice schemes appear to be the best solutions available. For example, a commit-and-prove protocol for integer products is significantly smaller and several orders of magnitude faster than other quantum-safe solutions.

These results show that lattices might eventually be very reasonable substitutes for the classical cryptography that is currently embedded in those protocols.

At the core of many of the recent privacy-based protocols is the BDLOP commitment scheme [BDL+18] which is able to commit to an arbitrary message vector \vec{m} modulo q . In [BDL+18], it was also shown how one can prove knowledge of \vec{m} by only proving knowledge of the commitment randomness \vec{r} . The proof technique was essentially using the “Fiat-Shamir with Aborts” [Lyu09] framework to keep the proofs small and avoid leaking any information about \vec{r} . If one uses the Gaussian rejection sampling procedure [Lyu12], then the magnitude of each coefficient of the \vec{r}' output in the proof is around $12 \cdot \kappa \|\vec{r}'\|$; where the κ is some constant that comes from the challenge and the 12 comes from a Gaussian tail bound needed for (statistical) zero-knowledge.

The increased coefficient size means that the proof is noticeably larger than the randomness itself (and gives rise to a vicious cycle of having to increase the modulus and dimension of the underlying lattice). It nevertheless seems necessary because leaking some information about the randomness can be dangerous. For example, if one were to repeatedly perform proofs of knowledge for the same commitment and leaking something about the same randomness each time, eventually the entire randomness could be recovered by even a passive observer.

1.1 One-Time Commitments

If one looks closer at how the BDLOP commitment is being used in many of the privacy-based protocols, one would notice that the scheme is used to commit to some intermediate value, give a proof-of-knowledge of the value (i.e. proof of knowledge of the commitment randomness), and then discards the commitment. So only *one* proof of knowledge is performed. In this case, it’s not immediately clear whether some leakage of the randomness vector is problematic. Still, it seems somewhat risky to only perform a minimal perturbation on the randomness and hope that this is good enough for LWE to remain secure. Instead of relying completely on heuristic security, it would be good to have a technique which lowers the proof size, and at the same time concretely allows one to understand exactly how the LWE problem is affected by the leakage.

For efficiency purposes, the BDLOP commitment instantiation is only computationally hiding (i.e. based on LWE), so one cannot use prior techniques to prove that enough entropy remains in the randomness after the leakage. While there are techniques that show that LWE-based encryption schemes can tolerate some leakage, the results are asymptotic and it’s unclear what the actual practical implication of the leakage is [DGK+10, GKPV10, BD20]. There has also been some work examining the practical aspects of leakage in the Ring-LWE setting [DDGR20], and the security of the scheme is barely affected when the leakage is small.

We show that a particular rejection sampling strategy information-theoretically leaks just one bit of randomness, and allows us to exactly quantify what this leakage is. More specifically, in addition to the public LWE samples, there is also a short public vector \vec{z} , and we require that the whole LWE secret

(i.e. the secret concatenated with the error vector) has a non-negative inner product with it. Because the LWE secret is uniformly distributed around 0, the probability that the inner product will be non-negative is greater than $1/2$, and so this extra restriction loses (less than) one bit of entropy.

We observe that the leakage essentially transforms the LWE problem instance into (a much less leaky version of) an extended-LWE one, which was shown to be equivalent to LWE in [AP12]. The decisional extended-LWE problem asks to distinguish between the distributions $(\mathbf{B}, \mathbf{B}\vec{r}, \vec{z}, \langle \vec{r}, \vec{z} \rangle)$ and $(\mathbf{B}, \vec{u}, \vec{z}, \langle \vec{r}, \vec{z} \rangle)$, where \vec{r} is sampled from the secret/noise domain of the LWE problem, \vec{u} is uniformly random, and \vec{z} is sampled from some efficiently sampleable distribution. One caveat is that for efficiency, we use a structured matrix \mathbf{B} , and the proof in [AP12] does not carry over to our setting.¹

Furthermore, it's clear that there is an equivalence between the *search* versions of extended Ring-LWE and Ring-LWE – and so a separation between the decisional versions would be quite interesting because all the best current lattice algorithms work by trying to solve the search problem. In short, we do not believe that this one bit of leakage has any practical consequences on the hardness of the Ring-LWE problem. It would actually be interesting if this assumption found even more applications, beside the original ones (e.g. [OPW11]) that inspired the Extended-LWE assumption, for improving the efficiency of lattice schemes.²

1.2 Technical Overview

The last prover move of a Schnorr-type Σ -protocol is a value $\vec{z} = \vec{y} + \vec{v}$, where \vec{y} is a “masking” vector that the prover created during the first move and \vec{v} is the combination of the secret and the challenge (usually their product, but this is unimportant for our purposes and we can consider the whole \vec{v} to be secret). If we would like \vec{z} to have small norm, then we cannot choose the coefficients of \vec{y} to be so large that the sum $\vec{y} + \vec{v}$ statistically hides \vec{v} . Instead we use rejection sampling to force the distribution of \vec{z} to be independent of \vec{v} . In particular, if we would like the distribution of the output to be f , while the distribution of \vec{z} is g , then one should output \vec{z} with probability $f(\vec{z})/(M \cdot g(\vec{z}))$, where M is some positive integer set so that this ratio is never larger than 1. Since $1/M$ is the probability that something will be output, we also want to have M as small as possible. Combining these two requirements, it's easy to see that M should be set to $\max_{\vec{z}}(g(\vec{z})/f(\vec{z}))$.

If we follow [Lyu12] where the target distribution f is a discrete Gaussian with standard deviation ε and the distribution g is a shifted Gaussian (by \vec{v}) with the same standard deviation, then via [Lyu12, Lemma 4.5], the value for M is derived as

¹ For the readers familiar with the *sample-preserving* reduction between search and decisional LWE problems [MM11], the underlying obstacles for that reduction and the extended-LWE reduction not carrying over to the Ring-LWE setting are similar.

² Indeed, much of the progress in constructions of practical classical cryptography has come from making stronger, but still plausible, assumptions that stem from discrete log or factoring.

$$\exp\left(\frac{-2\langle \vec{z}, \vec{v} \rangle + \|\vec{v}\|^2}{2\mathfrak{s}^2}\right) \leq \exp\left(\frac{24\mathfrak{s}\|\vec{v}\| + \|\vec{v}\|^2}{2\mathfrak{s}^2}\right) = M.$$

To obtain the inequality, one uses a standard 1-dimensional tail bound for the inner product of a discrete Gaussian with arbitrary vector (c.f. [Lyu12, Lemma 4.3]). And in fact, it is this tail bound that contributes the most to M . For example, if we would like to have $M = \exp(1)$, then we would need to set $\mathfrak{s} \approx 12\|\vec{v}\|$. On the other hand, if we knew that $\langle \vec{z}, \vec{v} \rangle \geq 0$, then we could set $\mathfrak{s} \approx \|\vec{v}\|/\sqrt{2}$, a decrease of around a factor of 17. So the intuition is for the prover to throw away the potential outputs \vec{z} that have the aforementioned inner product be negative-valued. The effect of this is that it leaks information about \vec{z} – in particular, the fact that its inner product with the secret is positive.

Interestingly, the new value of \mathfrak{s} is identical to what one would get by using bimodal rejection sampling, as in the BLISS signature scheme [DDLL13]. We cannot directly apply the technique from that paper because it required that the public value \mathbf{B}_0 be set up such that $\mathbf{B}_0\vec{r} = q \pmod{2q}$. Intuitively, with this setup, leaking the sign of the secret \vec{r} doesn't matter because both \vec{r} and $-\vec{r}$ satisfy the equality. Since we do not have this setup, the bimodal technique from BLISS would leak some information. At this point, we do not see a way to quantify the leakage stemming from the bimodal distribution, unlike the fairly simple leakage that we proposed instead. We should also mention that a recent work [TWZ20] showed how to apply the bimodal technique to the BDLOP commitment scheme without any leakage. Their transformation is not for free, though – it increases the length of the output, and also has a technical subtlety that makes it difficult to apply to our current result. The issue is that a part of our security proof requires that the challenge space has a particular distribution over $\{-1, 0, 1\}$ (it's necessary for Lemma 2.1, which is crucial for the product proof over fully-splitting rings). In the proof in [TWZ20], on the other hand, the challenges are masked by the prover and so their distribution is possibly determined in an adversarial fashion.

1.3 Applications and Paper Overview

A good benchmark for the progress of the development of lattice-based proof techniques is the simple proof of knowledge of a trinary secret vector \vec{s} satisfying $\mathbf{B}\vec{s} = \vec{t}$. Up until a few years ago, for a particular real-world parametrization of the problem (\mathbf{B} being a 1024×2048 dimensional matrix over \mathbb{Z}_q for $q \approx 2^{32}$), the smallest proof used “Stern proofs” [Ste93] ported to lattices [LNSW13], and had proof sizes of several megabytes. Switching the proof strategy to showing that a coefficient s_i is trinary by using a commit-and-prove strategy by proving the algebraic equation $(s_i - 1) \cdot s_i \cdot (s_i + 1) = 0$ lowered the proof sizes to around 300KB [YAZ+19, BLS19]. Then even more recently, by utilizing more efficient product proofs [Ezs+19, ALS20], it was shown in [ENS20] how to obtain proofs of size under 50KB.

In Sect. 3, we apply our new technique to the aforementioned benchmark, as well as to the recent framework in [LNS20a] for proving integer relations.

Proving	[ENS20]	[LNS20a]	this work	heuristic masking with no rejection sampling
Knowledge of an LWE Sample	47KB	47KB	33.3 KB	27 KB
Integer Addition	–	24.8KB	16.9 KB	13.8 KB
Integer Multiplication	–	40.2KB	28.2 KB	23 KB

Fig. 1. Proof size comparisons for secure commit-and prove protocols for proving knowledge of a (Module)-LWE sample of total dimension (secret and error vector) 2048, and 128-bit integer addition and multiplication. In the last column we list the proof sizes when only masking the secret vector $\mathbf{c}\vec{r}$ with a uniformly random masking vector that has the same infinity norm bound than the secret vector and not performing any rejection sampling. This is possibly still secure, but lacks a security proof at this point.

[LNS20b, Appendix B] shows how this framework can be further optimised using commitment compression from [DKL+18]. In total, we obtain a size reduction of around 30% for these proofs (see Fig. 1). As a side note, the proofs in [ENS20, LNS20a] used the uniform distribution instead of discrete Gaussians. The reason was that Gaussians did not provide much of a size reduction (a few kilobytes) and are a bit more cumbersome to work with. But in light of our modified rejection sampling, which only appears to apply to the Gaussian technique, we believe that it is now worthwhile to switch in order to take advantage of the additional size reduction.

To see how far our new rejection sampling procedure is from the (almost) best possible, we also compute the proof size for the case in which we don't do *any* rejection sampling when giving a proof of knowledge of the commitment randomness. In particular, we heuristically mask the secret by adding a random masking vector whose coefficients are as big as those of the vector. This undoubtedly leaks something about the randomness, but it may still be reasonable to hope that the message in the commitment remains hidden because there is still enough entropy in the LWE secret (i.e. the randomness of the commitment).³ This strategy leads to proof sizes that are around 20% smaller. We do not see how one can get any sort of security reduction for this approach, but we don't see an attack either. It is therefore an interesting open problem whether this approach can lead to something with a security proof. This would lead to essentially optimal parameters for this approach of creating zero-knowledge proofs.

Since our basic proofs are fairly close to optimal (as far as this line of research goes), we find it worthwhile to give further applications of them. Combined with a few extra tricks we develop along the way, we believe this leads to the shortest current constructions for certain primitives. Section 4 describes how one constructs a verifiable decryption scheme for Kyber [BDK+18], which is a finalist of the NIST PQC competition. This approach can be easily extended to other lattice-based KEM finalists, such as Saber [DKRV18] or NTRU [HPS98].

³ Even smaller sizes would be of course obtained if one does no masking at all, but then the scheme would be clearly insecure.

Complementing the previous section, which requires range proofs in the infinity norm, in [LNS20b, Appendix E] we further consider range proofs with respect to the Euclidean norm. Concretely, we show how to prove that a vector \vec{v} has $\|\vec{v}\| < a$ for some integer a . In [LNS20b, Appendix F], we describe an alternative approach for creating an important ingredient in some proofs – namely, *approximate* range proofs [BL17, BN20, LNS20a] – by using bimodal Gaussians [DDLL13] which further increases the efficiency of the aforementioned protocol. We also remark that the latter two sections make use of the improved framework from [LNS20a] defined in Sect. 3.

2 Preliminaries

2.1 Notation

Denote \mathbb{Z}_p to be the ring of integers modulo p . Let q be an odd prime. We write $\vec{v} \in R^m$ to denote vectors over a ring R and matrices over R will be written as regular capital letters M . Define I_n to be the $n \times n$ identity matrix over \mathbb{Z}_q . For an integer a , we define a vector $\vec{a} = (a, \dots, a)$ unless stated otherwise. By default, all vectors are column vectors. We write $\vec{v}||\vec{w}$ for a usual concatenation of \vec{v} and \vec{w} (which is still a column vector). For $\vec{v}, \vec{w} \in \mathbb{Z}_q^k$, $\vec{v} \circ \vec{w}$ is the usual component-wise multiplication. 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 . We write $[n]$ to denote the set $\{1, \dots, n\}$. Given two functions $f, g : \mathbb{N} \rightarrow [0, 1]$, we write $f(\mu) \approx g(\mu)$ if $|f(\mu) - g(\mu)| < \mu^{-\omega(1)}$. A function f is negligible if $f \approx 0$. We write $\text{negl}(n)$ to denote an unspecified negligible function in n .

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

Modular Reductions. We define $r' = r \bmod^\pm q$ to be the unique element r' in the range $-\frac{q-1}{2} \leq r' \leq \frac{q-1}{2}$ such that $r' \equiv r \pmod{q}$. We also denote $r' = r \bmod^+ q$ to be the unique element r' in the range $0 \leq r' < q$ such that $r' \equiv r \pmod{q}$. When the exact representation is not important, we simply write $r \bmod q$.

Sizes of Elements. For an element $w \in \mathbb{Z}_q$, we write $\|w\|_\infty$ to mean $|w \bmod^\pm q|$. Define the ℓ_∞ and ℓ_p norms for $\mathbf{w} = w_0 + w_1X + \dots + w_{d-1}X^{d-1} \in \mathcal{R}$ as follows:

$$\|\mathbf{w}\|_\infty = \max_j \|w_j\|_\infty, \quad \|\mathbf{w}\|_p = \sqrt[p]{\|w_0\|_\infty^p + \dots + \|w_{d-1}\|_\infty^p}.$$

If $\vec{\mathbf{w}} = (\mathbf{w}_1, \dots, \mathbf{w}_m) \in \mathcal{R}^m$, then

$$\|\vec{\mathbf{w}}\|_\infty = \max_j \|\mathbf{w}_j\|_\infty, \quad \|\vec{\mathbf{w}}\|_p = \sqrt[p]{\|\mathbf{w}_1\|^p + \dots + \|\mathbf{w}_k\|^p}.$$

By default, we denote $\|\vec{\mathbf{w}}\| := \|\vec{\mathbf{w}}\|_2$.

2.2 Cyclotomic Rings

Suppose (q) splits into l prime ideals of degree d/l in \mathcal{R} . This means $X^d + 1 \equiv \varphi_1 \dots \varphi_l \pmod{q}$ with irreducible polynomials φ_j of degree d/l modulo q . We assume that \mathbb{Z}_q contains a primitive $2l$ -th root of unity $\zeta \in \mathbb{Z}_q$ but no elements whose order is a higher power of two, i.e. $q - 1 \equiv 2l \pmod{4l}$. Therefore, we have

$$X^d + 1 \equiv \prod_{j \in \mathbb{Z}_{2l}^\times} \left(X^{\frac{d}{l}} - \zeta^j \right) \pmod{q} \quad (1)$$

where ζ^j ($j \in \mathbb{Z}_{2l}^\times$) ranges over all the l primitive $2l$ -th roots of unity.

Let $\mathcal{M}_q := \{\mathbf{p} \in \mathbb{Z}_q[X] : \deg(\mathbf{p}) < d/l\}$. We define the Number Theoretic Transform (NTT) of a polynomial $\mathbf{p} \in \mathcal{R}_q$ as follows:

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

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

2.3 Automorphisms

The ring \mathcal{R}_q has a group of automorphisms $\text{Aut}(\mathcal{R}_q)$ that is isomorphic to $\mathbb{Z}_{2d}^\times \cong \mathbb{Z}_2 \times \mathbb{Z}_{d/2}$,

$$i \mapsto \sigma_i: \mathbb{Z}_{2d}^\times \rightarrow \text{Aut}(\mathcal{R}_q),$$

where σ_i is defined by $\sigma_i(X) = X^i$. Note that for $i \in \mathbb{Z}_{2d}^\times$ and odd j it holds that $(\sigma_i(X - \zeta^j)) = (X - \zeta^{ji^{-1}})$ in \mathcal{R}_q (as ideals), and for $\mathbf{f} \in \mathcal{R}_q$,

$$\sigma_i(\mathbf{f} \bmod (X - \zeta^j)) = \sigma_i(\mathbf{f}) \bmod (X - \zeta^{ji^{-1}}).$$

Let k be a divisor of l and $\sigma := \sigma_{2l/k+1} \in \text{Aut}(\mathcal{R}_q)$. Then, we can write

$$(X^d + 1) = \prod_{j \in \mathbb{Z}_{2l/k}^\times} \prod_{i=0}^{k-1} \sigma^i \left(X^{\frac{d}{k}} - \zeta^j \right).$$

2.4 Challenge Space

Let $\mathcal{C} := \{-1, 0, 1\}^d \subset \mathcal{R}_q$ be the challenge set of ternary polynomials with coefficients $-1, 0, 1$. We define the following probability distribution $C : \mathcal{C} \rightarrow [0, 1]$. The coefficients of a challenge $\mathbf{c} \leftarrow C$ are independently identically distributed with $P(0) = 1/2$ and $\Pr(1) = \Pr(-1) = 1/4$.

Consider the coefficients of the polynomial $\mathbf{c} \bmod (X^{d/l} - \zeta^j)$ for $\mathbf{c} \leftarrow C$. Then, all coefficients follow the same distribution over \mathbb{Z}_q . Let us write Y for the random variable over \mathbb{Z}_q that follows this distribution. Attema et al. [ALS20] give an upper bound on the maximum probability of Y .

Lemma 2.1. *Let the random variable Y over \mathbb{Z}_q be defined as above. Then for all $x \in \mathbb{Z}_q$,*

$$\Pr(Y = x) \leq \frac{1}{q} + \frac{2l}{q} \sum_{j \in \mathbb{Z}_q^\times / \langle \zeta \rangle} \prod_{i=0}^{l-1} \left| \frac{1}{2} + \frac{1}{2} \cos(2\pi j y \zeta^i / q) \right|. \quad (2)$$

In particular, [ALS20, ENS20] computed that for $q \approx 2^{32}$, the maximum probability for each coefficient of $c \bmod X^{d/l} - \zeta^j$ is around $2^{-31.4}$. In general, we will call this probability \mathfrak{p} .

An immediate consequence of Lemma 2.1 is that polynomial $\mathbf{c} \leftarrow C$ is invertible in \mathcal{R}_q with overwhelming probability as long as parameters q, d, l are selected so that $q^{-d/l}$ is negligible.

Let k be a divisor of d such that $q^{-kd/l}$ is negligible and set $\sigma = \sigma_{2l/k+1}$. Let us define a probability distribution \tilde{C} over \mathcal{R}^k which first samples $\mathbf{c} = \mathbf{c}_0 + \mathbf{c}_1 X + \dots + \mathbf{c}_{k-1} X^{k-1} \leftarrow C$ and outputs $(\mathbf{c}_0, \dots, \mathbf{c}_{k-1})$ where each \mathbf{c}_i is defined as $\mathbf{c}_i = \sum_{j=0}^{d/k-1} c_{jk+i} X^{jk}$. Clearly, we have

$$\mathbf{c} = \sum_{i=0}^{k-1} \mathbf{c}_i X^i.$$

2.5 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 (M-LWE) and Module-SIS (M-SIS) [LS15]. Both problems are defined over \mathcal{R}_q .

Definition 2.2 (M-SIS $_{\kappa, m, B}$). *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 $\vec{z} \in \mathcal{R}_q^m$ such that $\mathbf{A}\vec{z} = \vec{\mathbf{0}}$ over \mathcal{R}_q and $0 < \|\vec{z}\| \leq B$. An algorithm \mathcal{A} is said to have advantage ϵ in solving M-SIS $_{\kappa, m, B}$ if*

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

Definition 2.3 ($\text{M-LWE}_{m,\lambda,\chi}$). *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) $(\mathbf{A}, \mathbf{A}\vec{s} + \vec{e})$ for $\mathbf{A} \leftarrow \mathcal{R}_q^{m \times \lambda}$, a secret vector $\vec{s} \leftarrow \chi^\lambda$ and error vector $\vec{e} \leftarrow \chi^m$, and 2) $(\mathbf{A}, \vec{b}) \leftarrow \mathcal{R}_q^{m \times \lambda} \times \mathcal{R}_q^m$. Then, \mathcal{A} is said to have advantage ϵ in solving $\text{M-LWE}_{m,\lambda,\chi}$ if*

$$\left| \Pr [b = 1 \mid \mathbf{A} \leftarrow \mathcal{R}_q^{m \times \lambda}; \vec{s} \leftarrow \chi^\lambda; \vec{e} \leftarrow \chi^m; b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{A}\vec{s} + \vec{e})] - \Pr [b = 1 \mid \mathbf{A} \leftarrow \mathcal{R}_q^{m \times \lambda}; \vec{b} \leftarrow \mathcal{R}_q^m; b \leftarrow \mathcal{A}(\mathbf{A}, \vec{b})] \right| \geq \epsilon. \tag{3}$$

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

2.6 Probability Distributions

For sampling randomness in the commitment scheme that we use, and to define a variant of the Ring Learning with Errors problem, we need to define an error distribution χ^d on \mathcal{R} . In this paper we sample the coefficients of the random polynomials in the commitment scheme using the distribution χ on $\{-1, 0, 1\}$ where ± 1 both have probability $5/16$ and 0 has probability $6/16$. This distribution is chosen (rather than the more “natural” uniform one) because it is easy to sample given a random bitstring by computing $a_1 + a_2 - b_1 - b_2 \pmod 3$ with uniformly random bits a_i, b_i .

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

Definition 2.4. *The discrete Gaussian distribution on \mathcal{R}^ℓ centered around $\vec{v} \in \mathcal{R}^\ell$ with standard deviation $\mathfrak{s} > 0$ is given by*

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

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

2.7 Approximate Range Proofs

Baum and Lyubashevsky [BL17] showed that if $B\vec{s}$ has small coefficients, for a vector \vec{s} over \mathbb{Z}_q and uniformly random binary matrix B , then with high probability \vec{s} must have small coefficients as well. More recently, Lyubashevsky et al. [LNS20a] generalise their result for $B\vec{s} + \vec{e}$ where \vec{e} is an arbitrary vector over \mathbb{Z}_q . We extend the lemma for the case when B is sampled from a distribution centered at 0 . The main advantage of this approach is that the infinity norm of $B\vec{s}$ decreases significantly, which is essential for the rejection sampling.

Concretely, we define the probability distribution $C : \{-1, 0, 1\} \rightarrow [0, 1]$ such that $\Pr_{c \leftarrow C}[c = 0] = p$ and $\Pr_{c \leftarrow C}[c = 1] = \Pr_{c \leftarrow C}[c = -1] = (1 - p)/2$ for some $p \in [0, 1]$. Then, we have the following lemma.

Lemma 2.5. *Let $\vec{s} \in \mathbb{Z}_q^m$ and $\vec{y} \in \mathbb{Z}_q^n$. Then*

$$\Pr \left[\|B\vec{s} + \vec{y}\|_\infty < \frac{1}{2} \|\vec{s}\|_\infty : B \leftarrow C^{n \times m} \right] \leq \max\{p, 1 - p\}^n.$$

We provide the proof of Lemma 2.5 in [LNS20b, Appendix A].

2.8 Commit-and-Prove

Let R_L be a polynomial-time verifiable relation containing (ck, x, w) . We will call ck the commitment key, x the statement and w the witness. Also, we define a language L_{ck} as the set of statements x for which there exists a witness w such that $(ck, x, w) \in R_L$.

We define the commit-and-prove functionality similarly as in [EG14, CLOS02] for a relation R_L . Roughly speaking, we want to commit to messages m_1, \dots, m_n and prove certain statements about them. Therefore, $w = (m_1, \dots, m_n)$ constitutes a witness for $x \in L_{ck}$.

Formally, a commit-and-prove functionality (CP) consists of four algorithms $CP = (\text{Gen}, \text{Com}, \text{Prove}, \text{Verify})$. We require $\text{Com}, \text{Verify}$ to be deterministic whereas Gen, Prove are probabilistic.

- $\text{Gen}(1^\mu)$: Given a security parameter μ , generates a commitment key ck . The commitment key specifies a message space \mathcal{M}_{ck} a randomness space \mathcal{R}_{ck} and commitment space \mathcal{C}_{ck} .
- $\text{Com}_{ck}(m; r)$: Given a commitment key ck , a message $m \in \mathcal{M}_{ck}$ and randomness $r \in \mathcal{R}_{ck}$ returns a commitment $c \in \mathcal{C}_{ck}$.
- $\text{Prove}_{ck}(x, ((m_1, r_1), \dots, (m_n, r_n)))$: Given a commitment key ck , statement x and commitment openings $m_i \in \mathcal{M}_{ck}, r_i \in \mathcal{R}_{ck}$ and $(ck, x, (m_1, \dots, m_n)) \in R_L$ returns a proof π .
- $\text{Verify}_{ck}(x, c_1, \dots, c_n, \pi)$: Given a commitment key ck , a statement x , a proof π and commitments $c_i \in \mathcal{C}_{ck}$, outputs 1 (accept) or 0 (reject).

Definition 2.6 (Correctness). *The commit-and-prove functionality CP has statistical correctness with correctness error $\rho : \mathbb{N} \rightarrow [0, 1]$ if for all adversaries \mathcal{A} :*

$$\Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, m_1, r_1, \dots, m_n, r_n) \leftarrow \mathcal{A}(ck); c_i = \text{Com}_{ck}(m_i; r_i); \left[\pi \leftarrow \text{Prove}_{ck}(x, ((m_1, r_1), \dots, (m_n, r_n))) : \text{Verify}_{ck}(x, c_1, \dots, c_n, \pi) = 0 \right] \leq \rho(\mu) \right] \quad (4)$$

where \mathcal{A} outputs $m_i \in \mathcal{M}_{ck}, r_i \in \mathcal{R}_{ck}$ so that $(ck, x, (m_1, \dots, m_n)) \in R_L$.

Definition 2.7 (Knowledge Soundness). *The commit-and-prove functionality CP is knowledge sound with knowledge error $\epsilon : \mathbb{N} \rightarrow [0, 1]$ if for all PPT \mathcal{A} there exists an expected polynomial time extractor \mathcal{E} so that :*

$$\Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, c_1, \dots, c_n, \pi) \leftarrow \mathcal{A}(ck); ((m_1^*, r_1^*) \dots, (m_n^*, r_n^*)) \leftarrow \mathcal{E}(c_1, \dots, c_n) : \left[\text{Verify}_{ck}(x, c_1, \dots, c_n, \pi) = 1 \wedge ((ck, x, (m_1^*, \dots, m_n^*)) \notin R_L \vee \exists i, \text{Com}(m_i^*; r_i^*) \neq c_i) \right] \right] \quad (5)$$

is less or equal to $\epsilon(\mu)$, where \mathcal{E} outputs $m_i^* \in \mathcal{M}_{ck}$ and $r_i^* \in \mathcal{R}_{ck}$.

In lattice-based zero-knowledge proofs, it is sometimes useful to relax the definition of knowledge soundness by only requiring $r_1^*, \dots, r_n^* \in \bar{\mathcal{R}}_{ck}$ where $\mathcal{R}_{ck} \subseteq \bar{\mathcal{R}}_{ck}$. However, the definition still makes sense if one can argue that the extracted commitments are still binding. This is what we do by additionally defining notions of *weak opening* (Definition 2.9).

The next property is a new notion called *simulatability*. Informally, it means that there exists an efficient simulator \mathcal{S} which can simulate both the commitment generation and the proof at the same time.

Definition 2.8 (Simulatability). *The commit-and-prove functionality CP is simulatable if there exist PPT simulators SimCom and SimProve such that for all PPT adversaries \mathcal{A} :*

$$\Pr \left[\begin{array}{l} ck \leftarrow \text{Gen}(1^\mu); (x, m_1, \dots, m_n) \leftarrow \mathcal{A}(ck); r_1, \dots, r_n \leftarrow \xi; \forall i, c_i = \text{Com}_{ck}(m_i, r_i); \\ \pi \leftarrow \text{Prove}_{ck}(x, (m_1, r_1), \dots, (m_n, r_n)) : (ck, x, (m_1, \dots, m_n)) \in R_L \wedge \mathcal{A}(c_1, \dots, c_n, \pi) = 1 \end{array} \right] \quad (6)$$

$$\approx \Pr \left[\begin{array}{l} ck \leftarrow \text{Gen}(1^\mu); (x, m_1, \dots, m_n) \leftarrow \mathcal{A}(ck); c_1, \dots, c_n \leftarrow \text{SimCom}_{ck}(x); \\ \pi \leftarrow \text{SimProve}_{ck}(x, c_1, \dots, c_n) : (ck, x, (m_1, \dots, m_n)) \in R_L \wedge \mathcal{A}(c_1, \dots, c_n, \pi) = 1 \end{array} \right]$$

where ξ is a probability distribution on \mathcal{R}_{ck} .

The difference between simulatability and zero-knowledge is that randomness r_1, \dots, r_n is directly generated from ξ as it would in the real-world protocol rather than chosen from adversary. This property becomes crucial when using the BDLOP commitments [BDL+18].

2.9 Commitment Scheme

We recall the BDLOP commitment scheme from [BDL+18] used in our constructions as well as previous works [ALS20, ENS20, LNS20a]. Suppose that we want to commit to a message vector $\vec{m} = (m_1, \dots, m_n) \in \mathcal{R}_q^n$ for $n \geq 1$ and that module ranks of κ and λ are required for M-SIS and M-LWE security, respectively. Then, in the key generation, a matrix $B_0 \leftarrow \mathcal{R}_q^{\kappa \times (\kappa + \lambda + n)}$ and vectors $\vec{b}_1, \dots, \vec{b}_n \leftarrow \mathcal{R}_q^{\kappa + \lambda + n}$ are generated and output as public parameters. Note that one could choose to generate $B_0, \vec{b}_1, \dots, \vec{b}_n$ in a more structured way as in [BDL+18] since it saves some computation. However, for readability, we write the commitment matrices in the “Knapsack” form as above. In our case, the hiding property of the commitment scheme is established via the duality between the Knapsack and MLWE problems. We refer to [Ezs+19, Appendix C] for a more detailed discussion.

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

$$\vec{t}_0 = B_0 \vec{r} \bmod q,$$

$$t_i = \langle \vec{b}_i, \vec{r} \rangle + m_i \bmod q,$$

for $i \in [n]$, where \vec{t}_0 forms the binding part and each t_i encodes a message polynomial m_i .

We recall the notion of a weak opening [ALS20].

Definition 2.9. A weak opening for the commitment $\vec{t} = \vec{t}_0 \parallel t_1 \parallel \dots \parallel t_n$ consists of d polynomials $\bar{c}_i \in \mathcal{R}_q$, a randomness vector \vec{r}^* over \mathcal{R}_q and messages $m_1^*, \dots, m_n^* \in \mathcal{R}_q$ such that

$$\begin{aligned} \|\bar{c}_i\|_1 &\leq 2k \text{ and } \bar{c}_i \bmod X - \zeta^{2i+1} \neq 0 \text{ for all } 0 \leq i < d, \\ \|\bar{c}_i \vec{r}^*\|_2 &\leq 2\beta \text{ for all } 0 \leq i < d, \\ B_0 \vec{r}^* &= \vec{t}_0, \\ \langle \vec{b}_i, \vec{r}^* \rangle + m_i^* &= t_i \text{ for } i \in [n] \end{aligned}$$

Attema et al. show that the commitment scheme is still binding with respect to weak openings.

2.10 Framework by Lyubashevsky et al. [LNS20a]

Recently, Lyubashevsky et al. [LNS20a] proposed a general framework for proving linear and multiplicative relations between committed messages. Concretely, suppose that the prover \mathcal{P} has a vector of n messages $\vec{m} = (m_1, \dots, m_n) \in \mathbb{Z}_q^{nl}$. Let \mathbf{pp} be a public set of polynomials $P : (\mathbb{Z}_q^l)^n \rightarrow \mathbb{Z}_q^l$ of n variables over \mathbb{Z}_q^l with standard component-wise addition and multiplication and define $\alpha = \max_{P \in \mathbf{pp}} \deg(P)$ ⁴. For readability, we will often “concatenate” polynomials in \mathbf{pp} , i.e. we will write $\mathbf{pp} = \{P_1, \dots, P_t\}$ where each $P_i : (\mathbb{Z}_q^l)^n \rightarrow \mathbb{Z}_q^{u_i l}$ and $u_1, \dots, u_t > 0$.

For the linear relations, let us set $\mathbf{ulp} = (A, \vec{u}) \in \mathbb{Z}_q^{vl \times nl} \times \mathbb{Z}_q^{vl}$. In practice, A can have arbitrary number of rows but then we would have to pad rows with zeroes in order to get a multiple of l .

Overall, [LNS20a] provides a protocol $\pi = (\text{Com}_{n,\alpha}, \Pi_n^\alpha(\mathbf{pp}, \mathbf{ulp}))$ where the prover \mathcal{P} first generates the BDLOP commitments to \vec{m} (sub-protocol $\text{Com}_{n,\alpha}$) and then wants to prove that $\vec{m} \in \mathcal{L}_n(\mathbf{pp}, \mathbf{ulp})$ (sub-protocol $\Pi_n^\alpha(\mathbf{pp}, \mathbf{ulp})$) where

$$\mathcal{L}_n(\mathbf{pp}, \mathbf{ulp}) := \{\vec{m} \in \mathbb{Z}_q^{nl} : \forall P \in \mathbf{pp}, P(\vec{m}) = \vec{0} \text{ and } A\vec{m} = \vec{u}\}. \tag{7}$$

In this paper we are only interested in applying the LNS framework only for the case $l = d$, i.e. $X^d + 1$ splits completely into linear factors modulo q , unless stated otherwise.

Let us formulate the protocol π in terms of the commit-and-prove functionality from Sect. 2.8. First, the relation R_{LNS} we are interested in here is

$$(ck, (\mathbf{pp}, \mathbf{ulp}), \vec{m}) \in R_{LNS} \iff \vec{m} \in \mathcal{L}_n(\mathbf{pp}, \mathbf{ulp}).$$

We define a CP functionality $LNS = (\text{LNSGen}, \text{LNSCom}, \text{LNSProve}^U, \text{LNSVerify}^U)$ for the relation R_{LNS} as follows:

⁴ Although Lyubashevsky et al. only consider the case $\alpha \leq 3$, it can be easily generalised by sending more garbage commitments.

- $\text{LNSGen}(1^\mu)$: It outputs a commitment key ck specifies the message space $\mathcal{M}_{ck} = \mathbb{Z}_q^{nd}$, randomness space $\mathcal{R}_{ck} = \{-1, 0, 1\}^{d \cdot (\kappa + \lambda + n + \alpha)}$ ⁵ and the commitment space $\mathcal{C}_{ck} = \mathcal{R}_q^{\kappa + n}$. It also generates the matrix $\mathbf{B}_0 \leftarrow \mathcal{R}_q^{\kappa \times (\kappa + \lambda + n + \alpha)}$ and vectors $\vec{\mathbf{b}}_1, \dots, \vec{\mathbf{b}}_{n+\alpha} \leftarrow \mathcal{R}_q^{\kappa + \lambda + n + \alpha}$.

- $\text{LNSCom}_{ck}((\vec{\mathbf{m}}_1, \dots, \vec{\mathbf{m}}_n); \vec{\mathbf{r}})$ outputs $(\vec{\mathbf{t}}_0, \mathbf{t}_1, \dots, \mathbf{t}_n)$ where

$$\vec{\mathbf{t}}_0 = \mathbf{B}_0 \vec{\mathbf{r}} \bmod q,$$

$$\mathbf{t}_i = \langle \vec{\mathbf{b}}_i, \vec{\mathbf{r}} \rangle + \text{NTT}^{-1}(\vec{\mathbf{m}}_i) \bmod q \text{ for } i \in [n].$$

- $\text{LNSProve}_{ck}^U((\text{pp}, \text{ulp}), \vec{\mathbf{m}}_1, \dots, \vec{\mathbf{m}}_n, \vec{\mathbf{r}})$: It runs the non-interactive version of the protocol $\Pi_n^\alpha(\text{pp}, \text{ulp})$ (e.g. [LNS20a, Fig. 8]), using the Fiat-Shamir transform, and outputs the proof π . Letter U denotes that the algorithm uses uniform rejection sampling.

- $\text{LNSVerify}_{ck}^U((\text{pp}, \text{ulp}), \vec{\mathbf{t}}_0, \mathbf{t}_1, \dots, \mathbf{t}_n, \pi)$: Check the verification equations of $\Pi_n^\alpha(\text{pp}, \text{ulp})$ (e.g. [LNS20a, Fig. 9]).

Eventually, Lyubashevsky et al. show that the commit-and-prove functionality LNS defined above for the relation R_{LNS} is both knowledge sound with negligible knowledge error under the Module-SIS assumption and simulatable under the Module-LWE assumption where the randomness $\vec{\mathbf{r}}$ distribution is defined over $\xi = \chi^{d \cdot (\kappa + \lambda + n + \alpha)}$.

3 Opening Proof with Improved Rejection Sampling

In lattice-based zero-knowledge proofs, e.g. [BLS19, ALS20], the prover will want to output a vector $\vec{\mathbf{z}}$ whose distribution should be independent of a secret randomness vector $\vec{\mathbf{r}}$, so that $\vec{\mathbf{z}}$ cannot be used to gain any information on the prover’s secret. During the protocol, the prover computes $\vec{\mathbf{z}} = \vec{\mathbf{y}} + \mathbf{c}\vec{\mathbf{r}}$ where $\vec{\mathbf{r}}$ is the randomness used to commit to the prover’s secret, $\mathbf{c} \leftarrow C$ is a challenge polynomial, and $\vec{\mathbf{y}}$ is a “masking” vector. In order to remove the dependency of $\vec{\mathbf{z}}$ on $\vec{\mathbf{r}}$, one applies the *rejection sampling* technique.

Lemma 3.1 ([BLS19]). *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. Also, write $\mathfrak{s} = 11T$ and $M = 3$. Now, sample $\vec{\mathbf{v}} \leftarrow \rho$ and $\vec{\mathbf{y}} \leftarrow D_{\mathfrak{s}}^{\ell d}$, set $\vec{\mathbf{z}} = \vec{\mathbf{y}} + \vec{\mathbf{v}}$, and run $b \leftarrow \text{Rej}_0(\vec{\mathbf{z}}, \vec{\mathbf{v}}, \mathfrak{s})$ as defined in Fig. 2. Then, the probability that $b = 0$ is at least $(1 - 2^{-100})/M$ and the distribution of $(\vec{\mathbf{v}}, \vec{\mathbf{z}})$, conditioned on $b = 0$, is within statistical distance of $2^{-100}/M$ of the product distribution $\rho \times D_{\mathfrak{s}}^{\ell d}$.*

Let us recall how parameters \mathfrak{s} and M are usually selected. Namely, the repetition rate M is chosen to be an upper-bound on:

$$\frac{D_{\mathfrak{s}}^{\ell d}(\vec{\mathbf{z}})}{D_{\vec{\mathbf{v}}, \mathfrak{s}}^{\ell d}(\vec{\mathbf{z}})} = \exp\left(\frac{-2\langle \vec{\mathbf{z}}, \vec{\mathbf{v}} \rangle + \|\vec{\mathbf{v}}\|^2}{2\mathfrak{s}^2}\right) \leq \exp\left(\frac{24\mathfrak{s}\|\vec{\mathbf{v}}\| + \|\vec{\mathbf{v}}\|^2}{2\mathfrak{s}^2}\right) = M. \quad (8)$$

⁵ Note that the length of $\vec{\mathbf{r}}$ is not $\kappa + \lambda + n$ as in Sect. 2.9 since the prover will later in the protocol commit to α garbage polynomials using the same $\vec{\mathbf{r}}$.

$\text{Rej}_0(\vec{z}, \vec{v}, \mathfrak{s})$	$\text{Rej}_1(\vec{z}, \vec{v}, \mathfrak{s})$
01 $u \leftarrow [0, 1)$	01 If $\langle \vec{z}, \vec{v} \rangle < 0$
02 If $u > \frac{1}{M} \cdot \exp\left(\frac{-2\langle \vec{z}, \vec{v} \rangle + \ \vec{v}\ ^2}{2\mathfrak{s}^2}\right)$	02 return 1
03 return 1	03 $u \leftarrow [0, 1)$
04 Else	04 If $u > \frac{1}{M} \cdot \exp\left(\frac{-2\langle \vec{z}, \vec{v} \rangle + \ \vec{v}\ ^2}{2\mathfrak{s}^2}\right)$
05 return 0	05 return 1
	06 Else
	07 return 0

Fig. 2. Two rejection sampling algorithms: the one used in previous works (left) and the one proposed in this section (right).

For the inequality we used the tail bound which says that with probability at least $1 - 2^{-100}$ we have $|\langle \vec{z}, \vec{v} \rangle| < 12\mathfrak{s}\|\vec{v}\|$ for $\vec{z} \leftarrow D_{\mathfrak{s}}^{\text{td}}$ [Ban93, Lyu12]. Hence, by setting $\mathfrak{s} = 11\|\vec{v}\|$ we obtain $M \approx 3$.

In this section we propose a new way to apply rejection sampling. Namely, we force \vec{z} to satisfy $\langle \vec{z}, \vec{v} \rangle \geq 0$, otherwise we abort. With this additional assumption, we can set M in the following way:

$$\exp\left(\frac{-2\langle \vec{z}, \vec{v} \rangle + \|\vec{v}\|^2}{2\mathfrak{s}^2}\right) \leq \exp\left(\frac{\|\vec{v}\|^2}{2\mathfrak{s}^2}\right) = M. \tag{9}$$

Hence, for $M \approx 3$ one would select $\mathfrak{s} = 0.675 \cdot \|\vec{v}\|$. Note that the probability for $\vec{z} \leftarrow D_{\sigma}^{\text{td}}$ that $\langle \vec{z}, \vec{v} \rangle \geq 0$ is at least $1/2$. Hence, the expected number of rejections would be at most $2M = 6$. On the other hand, if one aims for $M = 6$ repetitions using (8), then $\mathfrak{s} = 6.74 \cdot \|\vec{v}\|$. Thus, we manage to reduce the standard deviation by around a factor of 10.

Subset Rejection Sampling. In order to prove security of our new rejection sampling algorithm, we need the following modification of the rejection sampling lemma by Lyubashevsky [Lyu12].

Lemma 3.2 (Subset Rejection Sampling). *Let V be an arbitrary set and $h : V \rightarrow \mathbb{R}, f : \mathbb{Z}^m \rightarrow \mathbb{R}$ be probability distributions. Also define a family of set $S_v \subseteq \mathbb{Z}^m$ for $v \in V$ and $S = \{(z, v) \subseteq V \times \mathbb{Z}^m : z \in S_v\}$. Suppose $g_v : \mathbb{Z}^m \rightarrow \mathbb{R}$ is a family distributions indexed by all $v \in V$ and there exist $M, \gamma \geq 0$ which satisfy:*

$$\begin{aligned} \forall v \in V, z \in S_v : M g_v(z) &\geq f(z) \\ \forall v \in V : \sum_{z \in S_v} f(z) &\geq \gamma. \end{aligned} \tag{10}$$

Then the distributions of the output of \mathcal{A} and \mathcal{F} , defined in Fig. 3, are identical. Moreover, the probability that \mathcal{A} and \mathcal{F} output something is at least $\frac{\gamma}{M}$.

\mathcal{A}	\mathcal{F}
01 $v \leftarrow h$	01 $v \leftarrow h$
02 $z \leftarrow g_v$	02 $z \leftarrow f$
03 if $z \notin S_v$ then abort	03 if $z \notin S_v$ then abort
04 output (z, v) with probability $\frac{f(z)}{Mg_v(z)}$	04 output (z, v) with probability $\frac{1}{M}$.

Fig. 3. Algorithms \mathcal{A} and \mathcal{F} for Lemma 3.2.

Proof. Let $v \in V$. If $z \in S_v$, the probability that \mathcal{A} outputs z is equal to $g_v(z) \cdot \frac{f(z)}{Mg_v(z)} = \frac{f(z)}{M}$. Otherwise, the probability that \mathcal{A} outputs $z \notin S_v$ is exactly 0. Hence

$$\Pr[\mathcal{A} \text{ outputs something}] = \sum_{v \in V} h(v) \sum_{z \in S_v} \frac{f(z)}{M} \geq \frac{\gamma}{M}.$$

Moreover, the probability that \mathcal{F} outputs something is also $\frac{\sum_{(z,v) \in S} h(v)f(z)}{M}$. Hence:

$$\begin{aligned} \Delta(\mathcal{A}, \mathcal{F}) &= \frac{1}{2} \left(\sum_{(z,v) \in S} |\mathcal{A}(z, v) - \mathcal{F}(z, v)| \right) \\ &= \frac{1}{2} \sum_{v \in V} h(v) \left(\sum_{z \in S_v} \left| g_v(z) \cdot \frac{f(z)}{Mg_v(z)} - \frac{f(z)}{M} \right| \right) \\ &= \frac{1}{2} \sum_{v \in V} h(v) \left(\sum_{z \in S_v} \left| \frac{f(z)}{M} - \frac{f(z)}{M} \right| \right) = 0. \end{aligned}$$

□

Later on, we will consider the special case when $f := D_{\mathbf{s}}^m, g_{\vec{v}} := D_{\vec{v}, \mathbf{s}}^m$ for $\vec{v} \in V \subseteq \mathbb{Z}^m$ and

$$S_v := \{\vec{z} \in \mathbb{Z}^m : \langle \vec{v}, \vec{z} \rangle \geq 0\}.$$

Then, the probability that \mathcal{A} outputs something is at least:

$$\frac{1}{M} \sum_{\vec{v} \in V} h(\vec{v}) \Pr_{\vec{z} \leftarrow D_{\mathbf{s}}^m} [\langle \vec{v}, \vec{z} \rangle \geq 0] \geq \frac{1}{2M}.$$

Therefore, we can set $\gamma = 1/2$. Here, we used the fact that $\Pr_{\vec{z} \leftarrow D_{\mathbf{s}}^m} [\langle \vec{v}, \vec{z} \rangle > 0] = \Pr_{\vec{z} \leftarrow D_{\mathbf{s}}^m} [\langle \vec{v}, \vec{z} \rangle < 0]$. We also highlight that the value M we chose in (9) indeed satisfies Eq. 10.

Extended M-LWE. One observes that with the new approach, the verifier learns some new information about the secret. Indeed, if a prover \mathcal{P} returns \vec{z} then the verifier \mathcal{V} knows that $\langle \vec{z}, \vec{v} \rangle \geq 0$. However, later on we will show

that the opening proof from [ALS20] using the new rejection sampling is still simulatable assuming that a new problem, which we call *Extended M-LWE*, is computationally hard. For readability, we will describe it in a “knapsack” form.

Definition 3.3 (Extended M-LWE $_{m,k,\lambda,\chi,\xi_1,\xi_2}$). *The Extended Module-LWE problem with parameters $m, \lambda > 0$ and error distributions χ, ξ_1 and ξ_2 over \mathcal{R} and \mathcal{R}^k respectively, asks the adversary \mathcal{A} to distinguish between the following two cases:*

1. $\left(\mathbf{B}, \mathbf{B}\vec{r}, \mathbf{c}_1, \dots, \mathbf{c}_k, \vec{z}, \text{sign} \left(\left\langle \vec{z}, \begin{pmatrix} \mathbf{c}_1 \vec{r} \\ \vdots \\ \mathbf{c}_k \vec{r} \end{pmatrix} \right\rangle \right) \right)$ for $\mathbf{B} \leftarrow \mathcal{R}_q^{m \times (m+\lambda)}$, a secret vector $\vec{r} \leftarrow \chi^{m+\lambda}$ and $(\vec{z}, \mathbf{c}_1, \dots, \mathbf{c}_k) \leftarrow \xi_1^{k(m+\lambda)} \times \xi_2$
2. $\left(\mathbf{B}, \vec{u}, \mathbf{c}_1, \dots, \mathbf{c}_k, \vec{z}, \text{sign} \left(\left\langle \vec{z}, \begin{pmatrix} \mathbf{c}_1 \vec{r} \\ \vdots \\ \mathbf{c}_k \vec{r} \end{pmatrix} \right\rangle \right) \right)$ for $\mathbf{B} \leftarrow \mathcal{R}_q^{m \times (m+\lambda)}$, $\vec{u} \leftarrow \mathcal{R}_q^m$ and $(\vec{z}, \mathbf{c}_1, \dots, \mathbf{c}_k) \leftarrow \xi_1^{k(m+\lambda)} \times \xi_2$

where $\text{sign}(a) = 1$ if $a \geq 0$ and 0 otherwise. Then, \mathcal{A} is said to have advantage ϵ in solving Extended M-LWE $_{m,k,\lambda,\chi,\xi_1,\xi_2}$ if

$$\left| \Pr \left[b = 1 \mid \mathbf{B} \leftarrow \mathcal{R}_q^{m \times (m+\lambda)}; \vec{r} \leftarrow \chi^{m+\lambda}; \vec{z} \leftarrow \xi_1^{k(m+\lambda)}; \vec{c} \leftarrow \xi_2; b \leftarrow \mathcal{A}(\mathbf{B}, \mathbf{B}\vec{r}, \vec{z}, \vec{c}, s) \right] - \Pr \left[b = 1 \mid \mathbf{B} \leftarrow \mathcal{R}_q^{m \times \lambda}; \vec{u} \leftarrow \mathcal{R}_q^m; \vec{z} \leftarrow \xi_1^{k(m+\lambda)}; \vec{c} \leftarrow \xi_2; b \leftarrow \mathcal{A}(\mathbf{B}, \vec{u}, \vec{z}, \vec{c}, s) \right] \right| \geq \epsilon.$$

where

$$s = \text{sign} \left(\left\langle \vec{z}, \begin{pmatrix} \mathbf{c}_1 \vec{r} \\ \vdots \\ \mathbf{c}_k \vec{r} \end{pmatrix} \right\rangle \right) \text{ and } \vec{c} = (\mathbf{c}_1, \dots, \mathbf{c}_k).$$

To simplify notation, we will write Extended M-LWE $_{m,\lambda}$ to denote Extended M-LWE $_{m,1,\lambda,\chi^d, D_{\mathfrak{g}}^d, C}$ where χ is defined in Sect. 2.6.

We note that the LWE problem with various side information has already been discussed in e.g. [DGK+10, AP12, DDGR20]. As far as we are aware, our new variant of M-LWE is the closest to the Extended LWE problem defined by Alperin-Sheriff and Peikert [AP12]⁶. Indeed, in [LNS20b, Appendix D] we show that hardness of the non-algebraic version of our Extended M-LWE (i.e. without any polynomial ring structure) can be reduced to plain LWE using similar techniques as in [AP12].

Hardness of Ring/Module-LWE is often analysed as an LWE problem since, so far, the best known attacks do not make use of the algebraic structure of the polynomial ring [ADPS15]. Then, the only two attacks which would be relevant to our Module-LWE problem are the primal and dual attacks

⁶ In [AP12] the hint is the inner product $\langle \vec{r}, \vec{z} \rangle$ of the secret vector \vec{r} and some \vec{z} sampled from a given distribution D .

[Alb17, AGVW17, APS15]. Interestingly, the concrete analysis suggests that solving search-LWE (using the primal attack) is more efficient than solving the decisional version of LWE (using the dual attack). Thus, we believe that the search Extended-MLWE problem should still be hard with one bit, i.e. s , leaked since the vector \vec{r} has high enough entropy.

3.1 Concrete Instantiation

We describe how to apply our rejection sampling in the opening proof by Attema et al. [ALS20]. For readability, we consider the simple version of the protocol without any commitment compression [DKL+18] or Galois automorphisms [LNS20a, Appendix A.3] for boosting soundness. We discuss how to apply all those improvements in Sect. 3.2 and [LNS20b, Appendix B].

One observes that the protocol presented in Fig. 4 is not very meaningful since it only shows that prover \mathcal{P} has a polynomial $\mathbf{m} \in \mathcal{R}_q$. However, it is a key ingredient to prove linear [ENS20] and multiplicative [ALS20] relations between committed messages.

Formally, let us define the following the commit-and-prove functionality $CP = (\text{Gen}, \text{Com}, \text{Prove}, \text{Verify})$:

- $\text{Gen}(1^\mu)$: Given a security parameter μ , generates a commitment key ck which specifies a message space $\mathcal{M}_{ck} = R_q$, a randomness space $\mathcal{R}_{ck} = \{-1, 0, 1\}^{(\lambda+\kappa+1)d}$ and commitment space $\mathcal{C}_{ck} = \mathcal{R}_q^{\kappa+1}$. It also generates the matrix $\mathbf{B}_0 \leftarrow \mathcal{R}_q^{\kappa \times (\kappa+\lambda+1)}$ and the vector $\vec{b}_1 \leftarrow \mathcal{R}_q^{\kappa+\lambda+1}$.
- $\text{Com}_{ck}(\mathbf{m}; \vec{r})$: Given a commitment key ck , a message $\mathbf{m} \in \mathcal{M}_{ck}$ and randomness $\vec{r} \in \mathcal{R}_{ck}$ returns a commitment $(\vec{t}_0, \mathbf{t}_1) = (\mathbf{B}_0 \vec{r}, \langle \vec{b}_1, \vec{r} \rangle + \mathbf{m}) \in \mathcal{C}_{ck}$.
- $\text{Prove}_{ck}(x, \mathbf{m}, \vec{r})$: It first generates $\vec{y} \leftarrow D_s^{\kappa+\lambda+1}$ and computes $\mathbf{c} = H(\mathbf{B}_0 \vec{y})$. Then, it computes $\vec{z} = \vec{y} + \mathbf{c} \vec{r}$ and gets $b \leftarrow \text{Rej}_1(\vec{z}, \mathbf{c} \vec{r}, \mathfrak{s})$. If $b = 0$, it outputs $\pi = (\mathbf{c}, \vec{z})$.
- $\text{Verify}_{ck}(x, \vec{t}_0, \mathbf{t}_1, \pi)$: Parse $\pi = (\mathbf{c}, \vec{z})$. If $\|\vec{z}\| \leq \mathfrak{s} \sqrt{2(\lambda + \kappa + 1)d}$ and $\mathbf{c} = H(\vec{B}_0 \vec{z} - \mathbf{c} \vec{t}_0)$, return 1. Otherwise, return 0.

Here, $H : \{0, 1\}^* \rightarrow \{-1, 0, 1\}^d \subset \mathcal{R}_q$ is a random oracle which generates output from the distribution C (see Sect. 2.4). The language R_L , for which CP is defined, is trivial: $(ck, x, \mathbf{m}) \in R_L \iff \mathbf{m} \in \mathcal{R}_q$.

Correctness and knowledge soundness of CP can be proven almost identically as in [ALS20, Theorem 4.4]. Hence, we will only focus on simulatability.

Theorem 3.4. *Suppose Extended M-LWE $_{\kappa+1, \lambda}$ is computationally hard. Then, the commit-and-prove functionality $CP = (\text{Gen}, \text{Com}, \text{Prove}, \text{Verify})$ defined above for the language R_L is simulatable in the random oracle model H .*

Proof. Let us consider the following hybrid algorithms.

- $\text{Prove}_{ck}^1(x, \mathbf{m}, \vec{r})$: It first generates $\vec{y} \leftarrow D_s^{\kappa+\lambda+1}$, $\mathbf{c} \leftarrow C$. Then, it computes $\vec{z} = \vec{y} + \mathbf{c} \vec{r}$ and gets $b \leftarrow \text{Rej}_1(\vec{z}, \mathbf{c} \vec{r}, \mathfrak{s})$. If $b = 1$, it outputs $\pi = (\mathbf{c}, \vec{z})$ and programs $\mathbf{c} = H(\vec{B}_0 \vec{z} - \mathbf{c} \vec{t}_0)$.




Fig. 4. Opening proof for the commitment scheme. If $b = 0$ then the protocol is identical to the one described in [ALS20] and uses Rej_0 defined in Fig. 2. On the other hand, if $b = 1$ then we apply the new rejection sampling algorithm Rej_1 in Fig. 2.

- $\text{Prove}_{ck}^2(x, \mathbf{m}, \vec{\mathbf{r}})$: It first generates $\vec{\mathbf{z}} \leftarrow D_s^{\kappa + \lambda + 1}, \mathbf{c} \leftarrow C$. If $\langle \vec{\mathbf{z}}, \mathbf{c}\vec{\mathbf{r}} \rangle \geq 0$ then with probability $1/M$ it outputs $\pi = (\mathbf{c}, \vec{\mathbf{z}})$ and programs $\mathbf{c} = H(\mathbf{B}_0 \vec{\mathbf{z}} - \mathbf{c}\vec{\mathbf{t}}_0)$.

It is easy to see that the difference between Prove and Prove^1 is that the algorithm programs the random oracle at one particular value $\mathbf{B}\vec{\mathbf{y}}$ (without checking whether it was already set). Hence, by arguing similarly as for zero-knowledge in [KLS18, DKL+18] we have that for all PPT adversaries \mathcal{A} :

$$\begin{aligned}
 & \Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); \vec{\mathbf{r}} \leftarrow \chi^{(\lambda + \kappa + 1)d}; (\vec{\mathbf{t}}_0, \mathbf{t}_1) = \text{Com}_{ck}(\mathbf{m}; \vec{\mathbf{r}}); \right. \\
 & \quad \left. \pi \leftarrow \text{Prove}_{ck}(x, \mathbf{m}, \vec{\mathbf{r}}) : (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{\mathbf{t}}_0, \mathbf{t}_1, \pi) = 1 \right] \\
 & \approx \Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); \vec{\mathbf{r}} \leftarrow \chi^{(\lambda + \kappa + 1)d}; (\vec{\mathbf{t}}_0, \mathbf{t}_1) = \text{Com}_{ck}(\mathbf{m}; \vec{\mathbf{r}}); \right. \\
 & \quad \left. \pi \leftarrow \text{Prove}_{ck}^1(x, \mathbf{m}, \vec{\mathbf{r}}) : (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{\mathbf{t}}_0, \mathbf{t}_1, \pi) = 1 \right].
 \end{aligned} \tag{11}$$

For the next hybrid, we apply Lemma 3.2. Namely, let $m = (\lambda + \kappa + 1)d$ and $V = \{\mathbf{c}\vec{\mathbf{r}} : \mathbf{c} \in \{-1, 0, 1\}^d, \vec{\mathbf{r}} \in \{-1, 0, 1\}^m\}$. Set the probability distribution $h : V \rightarrow \mathbb{R}$ as

$$h(\vec{\mathbf{v}}) = \Pr[\mathbf{c}\vec{\mathbf{r}} = \vec{\mathbf{v}} : \mathbf{c} \leftarrow C, \vec{\mathbf{r}} \leftarrow \chi^m].$$

Next, we set $f := D_s^m$, $g_{\vec{v}} := D_{\vec{v},s}^m$ for $\vec{v} \in V$ and

$$S_{\vec{v}} := \{\vec{z} \in \mathcal{R}_q^{\lambda+\kappa+1} : \langle \vec{v}, \vec{z} \rangle \geq 0\}.$$

Then, by Lemma 3.2 we have:

$$\begin{aligned} & \Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); \vec{r} \leftarrow \chi^{(\lambda+\kappa+1)d}; (\vec{t}_0, \mathbf{t}_1) = \text{Com}_{ck}(\mathbf{m}; \vec{r}); \right. \\ & \left. \pi \leftarrow \text{Prove}_{ck}^1(x, \mathbf{m}, \vec{r}) : (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{t}_0, \mathbf{t}_1, \pi) = 1 \right] \\ &= \Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); \vec{r} \leftarrow \chi^{(\lambda+\kappa+1)d}; (\vec{t}_0, \mathbf{t}_1) = \text{Com}_{ck}(\mathbf{m}; \vec{r}); \right. \\ & \left. \pi \leftarrow \text{Prove}_{ck}^2(x, \mathbf{m}, \vec{r}) : (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{t}_0, \mathbf{t}_1, \pi) = 1 \right] \end{aligned} \quad (12)$$

for all adversaries \mathcal{A} . Now we define two algorithms SimCom and SimProve responsible for the simulation.

- $\text{SimCom}_{ck}(x)$: It samples \vec{t}_0 and \mathbf{t}_1 uniformly at random from \mathcal{R}_q^κ and \mathcal{R}_q respectively and returns $(\vec{t}_0, \mathbf{t}_1)$.
- $\text{SimProve}_{ck}(x, \vec{t}_0, \mathbf{t}_1)$: It first generates $\vec{z} \leftarrow D_s^{\kappa+\lambda+1}$, $\vec{r}^* \leftarrow \chi^{\lambda+\kappa+1}$ and $\mathbf{c} \leftarrow C$. Then, if $\langle \vec{z}, \mathbf{c}\vec{r}^* \rangle \geq 0$ then with probability $1/M$ it outputs $\pi = (\mathbf{c}, \vec{z})$ and programs $\mathbf{c} = H(\mathbf{B}_0\vec{z} - \mathbf{c}\vec{t}_0)$.

For the sake of contradiction suppose there exists a PPT adversary \mathcal{A} such that

$$\begin{aligned} & \left| \Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); \vec{r} \leftarrow \chi^{(\lambda+\kappa+1)d}; (\vec{t}_0, \mathbf{t}_1) = \text{Com}_{ck}(\mathbf{m}; \vec{r}); \right. \right. \\ & \left. \left. \pi \leftarrow \text{Prove}_{ck}^2(x, \mathbf{m}, \vec{r}) : (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{t}_0, \mathbf{t}_1, \pi) = 1 \right] \right. \\ & \left. - \Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); (\vec{t}_0, \mathbf{t}_1) \leftarrow \text{SimCom}_{ck}(x); \right. \right. \right. \\ & \left. \left. \pi \leftarrow \text{SimProve}_{ck}(x, \vec{t}_0, \mathbf{t}_1); (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{t}_0, \mathbf{t}_1, \pi) = 1 \right] \right| = \epsilon. \end{aligned} \quad (13)$$

Let us construct an adversary \mathcal{B} which solves the Extended M-LWE $_{\kappa+1,\lambda}$ also with probability ϵ using the algorithm \mathcal{A} . Concretely, suppose that \mathcal{B} is given a tuple $((\mathbf{B}_0 \parallel \vec{\mathbf{b}}_1), (\vec{t}_0 \parallel \mathbf{u}_1), \vec{z}, \mathbf{c}, \text{sign}(\langle \vec{z}, \mathbf{c}\vec{r}^* \rangle))$ for $\vec{z} \leftarrow D_s^{(\lambda+\kappa+1)d}$, $\mathbf{c} \leftarrow C$ and $\vec{r} \leftarrow \chi^{(\lambda+\kappa+1)d}$. Firstly, \mathcal{A} outputs a pair (x, \mathbf{m}) . Then, \mathcal{B} sets $\mathbf{t}_1 = \mathbf{u}_1 + \mathbf{m}$. Finally, if $\text{sign}(\langle \vec{z}, \mathbf{c}\vec{r}^* \rangle) \geq 0$ then \mathcal{B} sets $\pi = (\mathbf{c}, \vec{z})$ and with probability $1/M$ sends $(\vec{t}_0, \mathbf{t}_1, \pi)$. Otherwise, it aborts. At the end, \mathcal{B} outputs the bit sent from \mathcal{A} .

First, suppose that $\vec{t}_0 = \mathbf{B}\vec{r}$ and $\mathbf{u}_1 = \langle \vec{\mathbf{b}}_1, \vec{r} \rangle$. Then, $(\mathbf{t}_0, \mathbf{t}_1)$ constructed by \mathcal{B} is indeed equal to $\text{Com}_{ck}(\mathbf{m}; \vec{r})$. Then, the way π is built is identical as in Prove^2 . In this case, the probability that \mathcal{B} outputs bit 1 is equal to

$$\Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); \vec{r} \leftarrow \chi^{(\lambda+\kappa+1)d}; (\vec{t}_0, \mathbf{t}_1) = \text{Com}_{ck}(\mathbf{m}; \vec{r}); \right. \\ \left. \pi \leftarrow \text{Prove}_{ck}^2(x, \mathbf{m}, \vec{r}) : (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{t}_0, \mathbf{t}_1, \pi) = 1 \right].$$

On the other hand, assume that \vec{t}_0 and \mathbf{u}_1 are chosen uniformly at random and also independently of \vec{r} . Then, \mathbf{t}_1 is random as well. Hence, the probability that \mathcal{B} outputs 1 is indeed equal to

$$\Pr \left[ck \leftarrow \text{Gen}(1^\mu); (x, \mathbf{m}) \leftarrow \mathcal{A}(ck); (\vec{t}_0, \mathbf{t}_1) \leftarrow \text{SimCom}_{ck}(x); \right. \\ \left. \pi \leftarrow \text{SimProve}_{ck}(x, \vec{t}_0, \mathbf{t}_1); (ck, x, \mathbf{m}) \in R_L \wedge \mathcal{A}(\vec{t}_0, \mathbf{t}_1, \pi) = 1 \right].$$

Thus, \mathcal{B} can efficiently distinguish between the two Extended M-LWE cases with probability ϵ . Since, we assumed that the problem is computationally hard, this implies that ϵ is negligible. Then, the statement holds by the hybrid argument. \square

3.2 Boosting Soundness and Decreasing Standard Deviation

The protocol in Fig. 4 has soundness error around $q^{-d/l}$ which is not necessarily negligible. In order to boost soundness, Attema et al. [ALS20] apply Galois automorphisms. We recall the extended opening proof protocol below.

Prover \mathcal{P} first generates $\vec{y}_0, \dots, \vec{y}_{k-1} \leftarrow D_s^{(\lambda+\kappa+1)d}$ and $\vec{r}, \vec{t}_0, \vec{t}_1$ as before. Next, it outputs $(\vec{t}_0, \vec{t}_1, \vec{w}_0, \dots, \vec{w}_{k-1})$ where $\vec{w}_i = \mathbf{B}_0 \vec{y}_i$. After receiving a challenge $\mathbf{c} \leftarrow C$ from the verifier, \mathcal{P} computes

$$\vec{z}_i = \vec{y}_i + \sigma^i(\mathbf{c})\vec{r} \text{ for } i = 0, \dots, k - 1$$

where $\sigma := \sigma_{2d/k+1}$ where k is a divisor of d . Then, the prover applies rejection sampling $\text{Rej}_1(\vec{z}, \vec{v}, \mathfrak{s})$ where $\vec{z} = \vec{z}_0 \parallel \dots \parallel \vec{z}_{k-1}$ and $\vec{v} = \sigma^0(\mathbf{c})\vec{r} \parallel \dots \parallel \sigma^{k-1}(\mathbf{c})\vec{r}$. If it does not abort, then \mathcal{P} outputs \vec{z} . Finally, the verifier checks that \vec{z} is small and

$$\mathbf{B}_0 \vec{z}_i = \vec{w}_i + \sigma^i(\mathbf{c})\vec{t}_0$$

for $i = 0, \dots, k - 1$. As argued by Attema et al., this protocol has soundness around $q^{-dk/l}$.

More recently, Lyubashevsky et al. [LNS20a, Appendix A.6] (also mentioned in [ENS20]) improved this opening proof by applying a simple modification. Suppose $X^n + 1$ splits completely modulo q , i.e. $l = d$. Let us write the challenge $\mathbf{c} = \mathbf{c}_0 + \mathbf{c}_1 X + \dots + \mathbf{c}_{k-1} X^{k-1} \leftarrow C$ where

$$\mathbf{c}_i = \sum_{j=0}^{d/k-1} c_{jk+i} X^{jk}.$$

By definition of $\sigma = \sigma_{2d/k+1}$, we have that $\sigma(\mathbf{c}_i) = \mathbf{c}_i$ for each i . Therefore, we have:

$$\sigma^i(\mathbf{c}) = \sum_{j=0}^{k-1} \sigma^i(X^j) \mathbf{c}_j.$$

The modified opening proof protocol is presented as follows. Prover \mathcal{P} samples $\vec{y}'_0, \dots, \vec{y}'_{k-1}$ from $D_s^{(\lambda+\kappa+1)d}$ as before. Then, the prover sends $\vec{w}'_i = \mathbf{B}_0 \vec{y}'_i$. After getting a challenge $\mathbf{c} \leftarrow C$, it computes $\mathbf{c}_0, \dots, \mathbf{c}_{k-1}$ as above and calculates \vec{z}'_i as:

$$\begin{pmatrix} \vec{z}'_0 \\ \vec{z}'_1 \\ \vdots \\ \vec{z}'_{k-1} \end{pmatrix} = \begin{pmatrix} \vec{y}'_0 \\ \vec{y}'_1 \\ \vdots \\ \vec{y}'_{k-1} \end{pmatrix} + \begin{pmatrix} \mathbf{c}_0 \vec{r} \\ \mathbf{c}_1 \vec{r} \\ \vdots \\ \mathbf{c}_{k-1} \vec{r} \end{pmatrix}.$$

Since each \mathbf{c}_i has only at most d/k non-zero coefficients, we manage to decrease the standard deviation possibly by a factor of k (in practice the improvement is smaller if one upper-bounds $\|\sigma^i(\mathbf{c})\vec{\mathbf{r}}\|$ more cleverly).

Eventually, the prover applies rejection sampling $\text{Rej}_1(\vec{\mathbf{z}}, \vec{\mathbf{v}}, \mathfrak{s})$ where $\vec{\mathbf{z}} = \vec{\mathbf{z}}'_0 \|\cdots\| \vec{\mathbf{z}}'_{k-1}$ and $\vec{\mathbf{v}} = \mathbf{c}_0\vec{\mathbf{r}} \|\cdots\| \mathbf{c}_{k-1}\vec{\mathbf{r}}$. After receiving vectors $\vec{\mathbf{z}}'_j$, \mathcal{V} first checks whether

$$\mathbf{B}_0\vec{\mathbf{z}}'_i \stackrel{?}{=} \vec{\mathbf{w}}'_i + \mathbf{c}_i\vec{\mathbf{t}}_0.$$

for $i = 0, \dots, k - 1$ and that each $\vec{\mathbf{z}}'_i$ is small.

Note that by computing

$$\vec{\mathbf{y}}_i = \sum_{j=0}^{k-1} \sigma^i(X^j)\vec{\mathbf{y}}'_j, \quad \vec{\mathbf{z}}_i = \sum_{j=0}^{k-1} \sigma^i(X^j)\vec{\mathbf{z}}'_j$$

and $\vec{\mathbf{w}}_i = \mathbf{B}_0\vec{\mathbf{y}}_i$ for $i = 0, \dots, k - 1$, we have:

$$\mathbf{B}_0\vec{\mathbf{z}}_i = \vec{\mathbf{w}}_i + \sigma^i(\mathbf{c})\vec{\mathbf{t}}_0$$

which is the exact verification equation as in [ALS20]. This observation is crucial in [LNS20a] in order to still be able to prove linear and multiplicative relations using techniques from [ALS20, ENS20].

Lyubashevsky et al. bound $\|\vec{\mathbf{v}}\|$ by first finding α so that

$$\Pr \left[\exists i, \|\mathbf{c}_i\vec{\mathbf{r}}\|_\infty > \alpha : \mathbf{c} \leftarrow C, \vec{\mathbf{r}} \leftarrow \chi^{(\lambda+\kappa+1)d} \right] < 2^{-128}$$

and then setting the bound $\|\vec{\mathbf{v}}\| \leq \alpha\sqrt{\lambda + \kappa + 1}$. In [LNS20b, Appendix C] we describe a more optimal way to compute an upper-bound on $\|\vec{\mathbf{v}}\|$ using almost identical methods as in [DDLL13].

3.3 Applications

We apply the new rejection sampling technique in the protocol by Esgin et al. [ENS20, Appendix B] to prove knowledge of secrets in LWE samples. Concretely, for $n = 2048$, we want to prove knowledge of a ternary vector $\vec{\mathbf{s}} \in \{-1, 0, 1\}^n$ such that

$$\vec{\mathbf{u}} = (A' \parallel I_m) \cdot \vec{\mathbf{s}} \pmod{q},$$

where I_m is the m -dimensional identity matrix, $A' \in \mathbb{Z}_q^{m \times (n-m)}$ is a public matrix chosen uniformly at random and q is a modulus of about 32 bits (i.e., $\log q \approx 32$). Note that $\vec{\mathbf{s}}$ here corresponds to the concatenation of a secret vector and an error vector of 1024 dimension each in the usual LWE setting. For fair comparison, we will use the protocol described in [ENS20, Fig. 3] with the following two modifications (i) we do the Gaussian rejection sampling according to Rej_1 instead of the uniform one and (ii) we apply the commitment compression techniques as in [LNS20b, Appendix B].

We set parameters $(q, d, l, k) = (\approx 2^{32}, 128, 128, 4)$ similarly as in [ENS20]. Esgin et al. choose the expected number of repetitions to be 18.87. Since sampling

from a discrete Gaussians is much less efficient than from a uniform distribution, for fairness we set $\mathfrak{s} = T$ and $M \approx 3.3$ where T is the upper-bound on $\|\vec{v}\|$ where $\vec{v} = \mathbf{c}_0 \vec{r} \parallel \cdots \parallel \mathbf{c}_{k-1} \vec{r}$. Esgin et al. use the fact that $\|\vec{v}\|_\infty \leq d/k = 32$ and thus they set $T = 32\sqrt{(\lambda + \kappa + 3 + 16)d}$. We, however, apply the bound described in [LNS20b, Appendix C] and observe that for (κ, λ) selected in the next paragraph, our bound on $\|\vec{v}\|$ is around five times smaller than in [ENS20].

Now we set λ and κ such that M-LWE and M-SIS are hard against known attacks. We measure the hardness with the root Hermite factor δ and aim for $\delta \approx 1.0043$ [ENS20, EZS+19, BLS19]. By assuming that the Extended M-LWE is almost equally hard as M-LWE, we set $\lambda = 10$ as in [ENS20]. On the other hand, for the M-SIS hardness we manage to set $\kappa = 8$ due to having smaller standard deviation \mathfrak{s} . Hence, without using the compression techniques, we obtain the proof of size 41.35KB compared to 47KB by [ENS20]. After applying the additional improvements described in [LNS20b, Appendix B] we reduce the proof size to 33.6KB.

Similarly, we can consider the *LNS* functionality defined in Sect. 2.10 where LNSProve^U use our new Gaussian rejection sampling in Fig. 2 instead of a uniform one. We will denote this variant of the protocol as LNSProve^D and the corresponding CP functionality for the language R_L as

$$\text{LNS}^D = (\text{LNSGen}, \text{LNSCom}, \text{LNSProve}^D, \text{LNSVerify}^D).$$

The total proof size (when using Discrete Gaussians and without any Dilithium compression techniques) is about

$$(n + \kappa + \alpha + 1)d \log q + k(\lambda + \kappa + n + \alpha)d \cdot \log(12\mathfrak{s}) \quad \text{bits.} \quad (14)$$

Here, n, α are defined in Sect. 2.10, k is a divisor of d such that q^{-k} is negligible and \mathfrak{s} is the standard deviation used for the rejection sampling. For efficiency, we also apply Dilithium compression methods described in [LNS20b, Appendix B].

We present the proof sizes for proving n -bit integer addition and multiplication using LNS^D in Fig. 1.

4 Verifiable Decryption

In this section we apply the improved LNS framework with the rejection sampling from Sect. 3 to the problem of constructing a verifiable decryption scheme. We restrict our attention the Kyber key encapsulation scheme [BDK+18] and its NIST level 1 parameter set Kyber512. Kyber is a finalist in the NIST PQC standardization effort. Our techniques work equally well for any of the other lattice-based KEMs in round 3 of the NIST process, i.e. Saber [DKRV18] and NTRU [HPS98]. Kyber512 uses module rank 2 over the ring $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^{256} + 1)$ with modulus $q = 3329$. The public key is given by an MLWE vector $\vec{t} = \mathbf{A}\vec{s} + \vec{e}$ where $\mathbf{A} \in \mathcal{R}_q^{2 \times 2}$ is a uniform public matrix and \vec{s} and \vec{e} are short secret vectors with coefficients in the interval $[-2, 2]$. The encryption of a binary

polynomial $\mathbf{m} \in \{0, 1\}^{256}$ encoding a 256-bit key consists of the rounded vector and polynomial

$$\begin{aligned} \vec{u} &= \text{Compress}_{10}(\mathbf{A}^T \vec{s}' + \vec{e}_u) = \mathbf{A}^T \vec{s}' + \vec{e}'_u \\ \mathbf{v} &= \text{Compress}_4\left(\langle \vec{t}, \vec{s}' \rangle + e_v + \left\lceil \frac{q}{2} \right\rceil \mathbf{m}\right) = \langle \vec{t}, \vec{s}' \rangle + e'_v + \left\lceil \frac{q}{2} \right\rceil \mathbf{m} \end{aligned}$$

where \vec{s}' , \vec{e}_u and e_v are again short, the functions Compress_{10} and Compress_4 compress to 10 and 4 bits per coefficient, and \vec{e}'_u , e'_v include the errors coming from the compression. Finally, decryption uses the observation

$$\mathbf{v} - \langle \vec{s}, \vec{u} \rangle = \langle \vec{e}, \vec{s}' \rangle - \langle \vec{s}, \vec{e}'_u \rangle + e'_v + \left\lceil \frac{q}{2} \right\rceil \mathbf{m},$$

which implies

$$\left\| \mathbf{v} - \langle \vec{s}, \vec{u} \rangle - \left\lceil \frac{q}{2} \right\rceil \mathbf{m} \right\|_\infty < \frac{q}{4}$$

with overwhelming probability. In fact, the decryption algorithm will recover \mathbf{m} precisely if this norm bound is true. In the scheme there is no guarantee for this bound and encrypted keys can fail to be decryptable with probability around 2^{-139} .

Now, for a verifiable decryption scheme we need to be able to prove knowledge of a vector \vec{s} and polynomials \mathbf{m} , \mathbf{x} such that

$$\mathbf{v} - \langle \vec{s}, \vec{u} \rangle - \left\lceil \frac{q}{2} \right\rceil \mathbf{m} = \mathbf{x} \tag{15}$$

$$\vec{s} \in \{-2, -1, 0, 1, 2\}^{512} \tag{16}$$

$$\mathbf{m} \in \{0, 1\}^{256} \tag{17}$$

$$\|\mathbf{x}\|_\infty < \frac{q}{4} \tag{18}$$

The first three properties (15), (16), (17) can in principle directly be proven with the LNS framework. For the fourth one we can use a small-base decomposition approach for doing range proofs. A small problem is posed by the magnitude of the Kyber modulus $q = 3329$. While it is possible to instantiate the LNS framework in a way that allows to directly prove linear equations modulo such small primes, this results in quite large soundness error and many repetitions in the opening proof. To circumvent this problem and arrive at a more efficient protocol, we use the framework with a much larger modulus q' and lift Eq. (15) to the integers. This means that we instead prove

$$\mathbf{v} - \langle \vec{s}, \vec{u} \rangle - \left\lceil \frac{q}{2} \right\rceil \mathbf{m} + d\mathbf{q} \equiv \mathbf{x} \pmod{q'} \tag{19}$$

for another secret polynomial $\mathbf{d} \in \mathbb{Z}[X]/(X^{256} + 1)$ whose range $\|\mathbf{d}\|_\infty \leq 2^{10}$ we also need to prove. Note that Eq. (19) for $q' > 2^{23}$ together with the range proofs implies Eq. (15) since from the ranges of the individual terms we know that the equation must hold over the integers, which in turn implies (15) since the additional term, which is a multiple of q , vanishes modulo q .

4.1 Range Proofs

In the protocol sketched above there are the two range proofs $\|\mathbf{d}\|_\infty \leq 2^{10}$ and $\|\mathbf{x}\|_\infty < q/4$. For the first range proof it is actually enough to prove that $q\|\mathbf{d}\|_\infty < q'/4$. We use a 64-bit prime q' for the LNS framework protocol. Then, there is enough head-room between the actual range 2^{10} of \mathbf{d} and $q'/(4q) > 2^{50}$ so that we can use the approximate shortness proof. On the other hand, for proving the range of \mathbf{x} it is important to not have any slack in the proof. So here we decompose \mathbf{x} in base 5. We choose base 5, since for proving the coefficients of \vec{s} to lie in the interval $[-2, 2]$ we already have degree-5 product relations and hence can prove the base-5 decomposition without any additional garbage commitments. The interval of the coefficients of \mathbf{x} has length $q/2 = 1664$ and hence we need 5 base-5 digits for each coefficient. Now, since 1664 is not a power of 5 we write each coefficient of \mathbf{x} in the form

$$a_0 + a_15 + a_25^2 + a_35^3 + a_4260$$

with $a_4 \in 0, \dots, 4$. This decomposition is not unique but precisely maps to the integers in the interval $[0, 1664]$.

4.2 Proof Size

We compute the size of the verifiable decryption scheme for Kyber512 from above. As commitment messages we have the vector \mathbf{s} of \mathbb{Z} -dimension 512, the polynomial \mathbf{m} of dimension 256, the masking polynomial for the approximate range proof for \mathbf{d} , and the expansion of \mathbf{x} in base-5, which has \mathbb{Z} -dimension $5 \cdot 256 = 1280$. This amounts to a total message dimension of $n = 2176$. We then computed that the full LNS protocol with 64-bit modulus, MLWE rank 20 and MSIS rank 5 has a proof size of 43.6KB.

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

References

- ADPS15. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. *IACR Cryptol. ePrint Arch.* **2015**, 1092 (2015)
- AGVW17. Albrecht, M.R., Göpfert, F., Virdia, F., Wunderer, T.: Revisiting the expected cost of solving uSVP and applications to LWE. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10624, pp. 297–322. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_11
- Alb17. Albrecht, M.R.: On dual lattice attacks against small-secret LWE and parameter choices in HELIB and SEAL. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10211, pp. 103–129. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_4

- ALS20. Attema, T., Lyubashevsky, V., Seiler, G.: Practical product proofs for lattice commitments. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12171, pp. 470–499. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56880-1_17
- AP12. Alperin-Sheriff, J., Peikert, C.: Circular and KDM security for identity-based encryption. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 334–352. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30057-8_20
- APS15. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. Cryptology ePrint Archive, Report 2015/046 (2015). <https://eprint.iacr.org/2015/046>
- Ban93. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. *Math. Ann.* **296**(1), 625–635 (1993)
- BD20. Brakerski, Z., Döttling, N.: Hardness of LWE on general entropic distributions. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 551–575. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_19
- BDK+18. Bos, J.W. et al.: CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In: 2018 IEEE European Symposium on Security and Privacy, EuroS&P, pp. 353–367 (2018)
- BDL+18. Baum, C., Damgård, I., Lyubashevsky, V., Oechsner, S., Peikert, C.: More efficient commitments from structured lattice assumptions. In: Catalano, D., De Prisco, R. (eds.) SCN 2018. LNCS, vol. 11035, pp. 368–385. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98113-0_20
- BL17. Baum, C., Lyubashevsky, V.: Simple amortized proofs of shortness for linear relations over polynomial rings. *IACR Cryptology ePrint Archive* **2017**, 759 (2017)
- BLS19. Bootle, J., Lyubashevsky, V., Seiler, G.: Algebraic techniques for Short(er) exact lattice-based zero-knowledge proofs. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 176–202. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_7
- BN20. Baum, C., Nof, A.: Concretely-efficient zero-knowledge arguments for arithmetic circuits and their application to lattice-based cryptography. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020. LNCS, vol. 12110, pp. 495–526. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45374-9_17
- CLOS02. Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: STOC, pp. 494–503. ACM (2002)
- DDGR20. Dachman-Soled, D., Ducas, L., Gong, H., Rossi, M.: LWE with side information: attacks and concrete security estimation. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12171, pp. 329–358. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56880-1_12
- DDLL13. Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal gaussians. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 40–56. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40041-4_3
- DGK+10. Dodis, Y., Goldwasser, S., Tauman Kalai, Y., Peikert, C., Vaikuntanathan, V.: Public-key encryption schemes with auxiliary inputs. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 361–381. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11799-2_22

- DKL+18. Ducas, L., et al.: Crystals-dilithium: a lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2018**(1), 238–268 (2018)
- DKRV18. D’Anvers, J.-P., Karmakar, A., Sinha Roy, S., Vercauteren, F.: Saber: module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. In: Joux, A., Nitaj, A., Rachidi, T. (eds.) *AFRICACRYPT 2018*. LNCS, vol. 10831, pp. 282–305. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89339-6_16
- EG14. Escala, A., Groth, J.: Fine-tuning groth-sahai proofs. In: Krawczyk, H. (ed.) *PKC 2014*. LNCS, vol. 8383, pp. 630–649. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54631-0_36
- EKS+20. Esgin, M.F., et al.: Practical post-quantum few-time verifiable random function with applications to algorand. *Cryptology ePrint Archive*, Report 2020/1222 (2020). <https://eprint.iacr.org/2020/1222>
- ENS20. Esgin, M.F., Nguyen, N.K., Seiler, G.: Practical exact proofs from lattices: new techniques to exploit fully-splitting rings. In: Moriai, S., Wang, H. (eds.) *ASIACRYPT 2020*. LNCS, vol. 12492, pp. 259–288. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64834-3_9
- EZS+19. Esgin, M.F., Zhao, R.K., Steinfeld, R., Liu, J.K., Liu, D.: Matricot: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In: *CCS*, pp. 567–584. ACM (2019)
- GKPV10. Goldwasser, S., Kalai, Y.T., Peikert, C., Vaikuntanathan, V.: Robustness of the learning with errors assumption. In: *ICS*, pp. 230–240. Tsinghua University Press (2010)
- HPS98. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: a ring-based public key cryptosystem. In: *ANTS*, pp. 267–288 (1998)
- KLS18. Kiltz, E., Lyubashevsky, V., Schaffner, C.: A concrete treatment of fiat-Shamir signatures in the quantum random-oracle model. In: Nielsen, J.B., Rijmen, V. (eds.) *EUROCRYPT 2018*. LNCS, vol. 10822, pp. 552–586. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78372-7_18
- LNS20a. Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Practical lattice-based zero-knowledge proofs for integer relations. In: *IACR Cryptology ePrint Archive*, 2020. ACM CCS (2020). <http://eprint.iacr.org/2020/1183>
- LNS20b. Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Shorter lattice-based zero-knowledge proofs via one-time commitments. *Cryptology ePrint Archive*, Report 2020/1448 (2020). <https://eprint.iacr.org/2020/1448>
- LNSW13. Ling, S., Nguyen, K., Stehlé, D., Wang, H.: Improved zero-knowledge proofs of knowledge for the ISIS problem, and applications. In: *PKC*, pp. 107–124 (2013)
- LS15. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. *Des. Codes Crypt.* **75**(3), 565–599 (2014). <https://doi.org/10.1007/s10623-014-9938-4>
- Lyu09. Lyubashevsky, V.: Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In: *ASIACRYPT*, pp. 598–616 (2009)
- Lyu12. Lyubashevsky, V.: Lattice signatures without trapdoors. In: *EUROCRYPT*, pp. 738–755 (2012)
- MM11. Micciancio, D., Mol, P.: Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In: Rogaway, P. (ed.) *CRYPTO 2011*. LNCS, vol. 6841, pp. 465–484. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_26

- OPW11. O’Neill, A., Peikert, C., Waters, B.: Bi-deniable public-key encryption. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 525–542. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_30
- Ste93. Stern, J.: A new identification scheme based on syndrome decoding. In: CRYPTO, pp. 13–21 (1993)
- TWZ20. Tao, Y., Wang, X., Zhang, R.: Short zero-knowledge proof of knowledge for lattice-based commitment. In: Ding, J., Tillich, J.-P. (eds.) PQCrypto 2020. LNCS, vol. 12100, pp. 268–283. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-44223-1_15
- YAZ+19. Yang, R., Au, M.H., Zhang, Z., Xu, Q., Yu, Z., Whyte, W.: Efficient lattice-based zero-knowledge arguments with standard soundness: construction and applications. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 147–175. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_6