



A Compressed Σ -Protocol Theory for Lattices

Thomas Attema^{1,2,3(✉)}, Ronald Cramer^{1,2}, and Lisa Kohl¹

¹ Cryptology Group, CWI, Amsterdam, The Netherlands
{cramer,lisa.kohl}@cwi.nl

² Mathematical Institute, Leiden University, Leiden, The Netherlands
cramer@math.leidenuniv.nl

³ Cyber Security and Robustness, TNO, The Hague, The Netherlands
thomas.attema@tno.nl

Abstract. We show a *lattice-based* solution for commit-and-prove transparent circuit zero-knowledge (ZK) with *polylog-communication*, the *first* not depending on PCPs.

We start from *compressed Σ -protocol theory* (CRYPTO 2020), which is built around basic Σ -protocols for opening an arbitrary linear form on a long secret vector that is compactly committed to. These protocols are first compressed using a recursive “folding-technique” adapted from Bulletproofs, at the expense of logarithmic rounds. Proving in ZK that the secret vector satisfies a given constraint – captured by a circuit – is then by (blackbox) reduction to the linear case, via arithmetic secret-sharing techniques adapted from MPC. Commit-and-prove is also facilitated, i.e., when commitment(s) to the secret vector are created ahead of any circuit-ZK proof. On several platforms (incl. DL) this leads to logarithmic communication. Non-interactive versions follow from Fiat-Shamir.

This abstract modular theory strongly suggests that it should somehow be supported by a lattice-platform *as well*. However, when going through the motions and trying to establish low communication (on a SIS-platform), a certain significant lack in current understanding of multi-round protocols is exposed.

Namely, as opposed to the DL-case, the basic Σ -protocol in question typically has *poly-small challenge* space. Taking into account the compression-step – which yields *non-constant* rounds – and the necessity for parallelization to reduce error, there is no known tight result that the compound protocol admits an efficient knowledge extractor. We resolve the state of affairs here by a combination of two novel results which are fully general and of independent interest. The first gives a tight analysis of efficient knowledge extraction in case of non-constant rounds combined with poly-small challenge space, whereas the second shows that parallel repetition indeed forces rapid decrease of knowledge error.

Moreover, in our present context, arithmetic secret sharing is not defined over a large finite field but over a quotient of a number ring and this forces our careful adaptation of how the linearization techniques are deployed.

We develop our protocols in an abstract framework that is conceptually simple and can be flexibly instantiated. In particular, the framework applies to arbitrary rings and norms.

1 Introduction

Compressed Σ -Protocol Theory [6] is built around basic Σ -protocols for opening an arbitrary linear form on a long secret vector that is compactly committed to. More precisely, these Σ -protocols allow a prover to prove that a committed vector \mathbf{x} satisfies a constraint $L(\mathbf{x}) = y$ captured by a linear form L . They are first compressed using a recursive “folding-technique” adapted from Bulletproofs [14, 16]. Compression reduces the communication complexity from linear down to logarithmic in the dimension of the secret vector \mathbf{x} , at the expense of a logarithmic number of rounds. Proving in ZK that the secret vector satisfies an arbitrary (non-linear) constraint – captured by an arithmetic circuit – is then by (blackbox) reduction to the linear case, via arithmetic secret-sharing techniques adapted from MPC. It was shown how to instantiate this theory from different hardness assumptions, i.e., the Discrete Logarithm (DL), Strong-RSA and Knowledge-of-Exponent (KEA) assumption. The latter assumption even results in *constant* communication, instead of logarithmic. Non-interactive versions follow from the Fiat-Shamir transform [26].

The starting point is always a compact and homomorphic vector commitment scheme, i.e., commitments should have size constant (or logarithmic) in the dimension of the committed vector. After instantiating such a commitment scheme from any of the aforementioned hardness assumption, compressed Σ -protocol theory can be described in an abstract and modular manner. This strongly suggests that the theory should also be supported by a lattice platform. This belief was further strengthened by the recent lattice-based Bulletproof instantiation for proving knowledge of a SIS preimage [15].

However, when going through the motions and trying to establish low communication (on a SIS-platform), a certain significant lack in current understanding of multi-round protocols and several challenges are exposed.

1.1 Challenges for Lattice Instantiations

As opposed to the DL-case, the lattice-based Σ -protocol typically has polynomially small challenge space. Taking into account the compression-step – which yields non-constant rounds – there is no known result from which a tight knowledge soundness property can be derived. In prior works, this lack in understanding was handled by an alternative non-tight security analysis [14]. Recent works, while remaining non-tight, have improved the tightness [3, 21, 30, 31, 41].

The situation is further complicated by the necessity for parallelization to reduce the knowledge error. While parallel repetition of interactive proofs has been studied extensively in the context of decreasing the *soundness error* [18, 19, 28], to the best of our knowledge there does not exist a general parallel repetition theorem for decreasing the *knowledge error*.

Setting aside the knowledge error issues addressed previously, the main difference between the lattice setting and the other settings is a norm bound. Instead of proving knowledge of a preimage for some homomorphism Ψ , we aim to prove

knowledge of a *short* pre-image. More precisely, for some homomorphism Ψ , we aim to construct a protocol for the following relation

$$R_{\Psi,\alpha} = \{(P; x) : P = \Psi(x), \|x\| \leq \alpha\}$$

where $(P; x) \in R_{\Psi,\alpha}$ is a pair of a public statement P and a secret witness x . The DL-based protocols are designed for exactly the same abstract relation, but without the norm-bound. This minor difference introduces a number of challenges that have been dealt with in the context of plain Σ -protocols for some time now. For example, given a preimage x with $\|x\| \leq \beta$, a prover is typically only capable of proving knowledge of a preimage y with $\|y\| \leq \alpha\beta$. The factor $\alpha \geq 1$ is referred to as the *soundness slack*. In multi-round protocols the soundness slack accumulates and a more careful analysis is warranted.

Finally, in our present lattice context, committed vectors typically have coefficients in the quotient of a number ring $\mathcal{R} = \mathbb{Z}[X]/(f(X))$ by a rational prime (p) . However, the structure of the ring \mathcal{R}_p may not readily allow for the large sets with invertible pairwise differences required for Shamir secret sharing.

1.2 Contributions

We show a lattice-based solution for commit-and-prove transparent circuit ZK with polylogarithmic communication, the first not depending on PCPs.

To this end, we resolve the lack in understanding regarding knowledge soundness by a combination of two novel results which are fully general and of independent interest. The first gives a tight analysis of efficient knowledge extraction in case of non-constant rounds, whereas the second shows that parallel repetition indeed forces rapid decrease of knowledge error.

By our extractor analysis, we *tightly* prove that (k_1, \dots, k_μ) -special soundness implies knowledge soundness, without imposing any restrictions on the size of the challenge sets. In a concurrent and independent work this result was deemed out of reach with current techniques [3]. More concretely, they apply the non-tight analysis of [21] and derive a knowledge error $\kappa \leq 8.16 \log n/|\mathcal{C}|$, where n is the size of the input. By contrast, we provide a tight bound and show that $\kappa \leq 2 \log n/|\mathcal{C}|$. This inequality contains a simplified expression and is therefore non-tight, for the tight bound we refer to Theorem 1. Furthermore, our result answers an open question regarding knowledge extractors, recently made explicit [30, Question D.4.], in the affirmative. It is generally applicable to all aforementioned platforms and therefore improves upon the analysis of [3, 14, 21, 30, 31, 41], directly yielding better parameters for multi-round protocols such as Bulletproofs. Towards showing that (k_1, \dots, k_μ) -special soundness tightly implies knowledge soundness, we observe that for the special case of 2-special soundness (where this implication is well-known) we can give a very simple proof that we have not encountered in the literature before. In contrast to standard proof techniques, our extractor can be modeled by a negative hyper geometric distribution. This simplification turns out to be generalizable to the multi-round scenario. Even though the general proof is building on this simplification, its analysis turns out to be quite involved.

By the second result, we show that parallel repetition indeed forces a rapid decrease of knowledge error, explicitly proving a result that is often taken for granted whereas it actually requires a careful analysis. More precisely, it is known that parallel repetition decreases the *soundness* error. However, *knowledge soundness* is a strictly stronger notion than soundness. Nevertheless, by a careful analysis, we prove that prior results also apply to knowledge sound protocols and allow for a rapid decrease of knowledge error. The $(2, 2)$ -special sound signature scheme MQDSS was already presented with a tight knowledge error analysis [17]. However, their analysis crucially depends on the fact that this signature scheme has a *constant* number of rounds and therefore does not apply to our setting. Our techniques are generic and also apply to this protocol, indeed yielding exactly the same knowledge error.

Furthermore, we describe a careful adaptation of the arithmetic secret sharing based linearization strategy from [6]. First, the evaluation points of Shamir’s secret sharing scheme have to be chosen from an *exceptional*, instead of an arbitrary, subset of the ring \mathcal{R}_p , i.e., a subset with invertible differences. In many practical scenarios this minor adaptation suffices. However, some rings do not contain “large enough” exceptional subsets. For this reason, we extend the linearization technique to work for small rings \mathcal{R}_p by defining the secret sharing scheme over an appropriately chosen ring extension. Some care is warranted to prevent dishonest provers from choosing secret elements in the extension ring.

Subsequently, we note that working in a lattice-platform is considerably more tedious. Traditionally the security analysis depends strongly on various protocol design choices. Our approach is less sensitive to these choices. This is very convenient when considering variations. More precisely, we develop our protocols in an abstract framework that is conceptually simple and can be flexibly instantiated. In particular, the framework applies to arbitrary rings, challenge sets and norms. Our framework captures general rejection sampling strategies, gives precise bounds on the introduced soundness slack and generalizes beyond factor-2 per-round compression.

The communication complexity of our protocols, when instantiated from the Module Short Integer Solution (MSIS) assumption and appropriately chosen rings, is polylogarithmic in the input size. Due to the soundness slack it does not achieve the logarithmic communication of a DL-based instantiation. Our protocols are transparent, i.e., no trusted setup, and easily ported to the commit-and-prove paradigm, where commitment(s) to the secret vector have been created ahead of any circuit-ZK proof. Moreover, various efficiency improvements, developed for DL-based (compressed) Σ -protocol theory, almost directly carry over to the lattice-setting.

1.3 Related Work

Circuit ZK with Polylogarithmic Complexity from PCPs. A generic class of (zero-knowledge) proof systems is based on *Probabilistically Checkable Proofs* (PCPs). The security of these protocols only relies on the existence of collision-resistant hash functions and they achieve polylogarithmic communication complexity.

However, large concrete costs have long prevented PCP-based protocols from being deployed in practice. Recent advances have rendered PCP-based protocols practical [5, 11, 12]. Still, for small problem instances, PCP-based protocols are often outperformed by other approaches relying on more structured hardness assumptions. In particular, PCP approaches rely on Merkle-tree commitments and therefore have an implicit lower bound in the order of a hundred kilobytes, whereas protocols relying on the compression mechanism such as Bulletproofs can go down to as much as a few kilobytes. Even though the soundness slack introduced by the compression mechanism is currently somewhat limiting in terms of concrete efficiency, we expect that on the long run the non-PCP lattice-based approach will lead to more succinct proofs.

Circuit ZK with Sublinear Complexity from Lattice Assumptions. The first protocol of this form achieving a sub-linear communication complexity $\tilde{O}(\sqrt{\lambda n})$, where n is the input size and λ the security parameter, was presented in [9]. A key component of their protocol is a compact commitment scheme. In our lattice instantiation we use exactly the same compact commitment scheme. While their approach is inherently limited to communication complexity in the order of $\tilde{O}(\sqrt{\lambda n})$, our approach yields the first lattice-based (non-PCP) protocol that achieves polylogarithmic complexity in the input length. On the other hand, our approach requires a larger number of rounds. Getting a similar communication-complexity/round trade-off as [9] by using a larger per-round compression seems currently out of reach, due to the large soundness slack introduced (which scales exponentially in the compression factor).

Lattice-based proof of knowledge of SIS preimages. The lattice-based Bulletproof instantiation of [15] is most similar to our compressed Σ -protocol. However, in this work the aforementioned knowledge error issues were overlooked. Moreover, their work only considers proving knowledge of a SIS preimage, i.e., it does not consider generic arithmetic circuit relations. Furthermore, it is not zero-knowledge and it is tailored to a specific lattice-instantiation. By contrast, our protocol is a circuit ZK protocol that can be instantiated from a wide variety of lattices. For the specific scenario of proving knowledge of a SIS preimage, we obtain a comparable communication complexity.

1.4 Roadmap

We start by presenting the general result that (k_1, \dots, k_μ) -special soundness tightly implies knowledge soundness in Sect. 3. We first outline a very simple proof for the special case of 2-special soundness, which is novel to the best of our knowledge. Subsequently, we show how this proof can be generalized to the multi-round setting. Using results from [19], we prove that parallel repetition of multi-round public-coin protocols not only reduces the soundness error, but also the knowledge error (see Sect. 4). In Sect. 5, we give an abstract theory for lattice-based compressed Σ -protocols. In Sect. 6, we show how to instantiate our abstract framework from the Module Short Integer Solution (MSIS) problem.

We further provide an asymptotic parameter analysis for our instantiation and comparison with [15]. In Sect. 7, we briefly explain how to handle non-linear relations and refer to the full version of this paper [1] for a detailed description of our techniques. Moreover, in the full version, we discuss a number of extensions for amortization over many linear forms, reducing the communication complexity and for obtaining commit-and-prove protocols directly.

2 Preliminaries

We say a function is *negligible*, if it vanishes faster than any inverse polynomial. If a function vanishes slower than some inverse polynomial, we say it is *noticeable*. For formal definitions and definitions of *statistical distance* and *statistically close distributions* we refer to the full version of this paper [1].

2.1 Interactive Proofs

Let $R \subset \{0, 1\}^* \times \{0, 1\}^*$ be a binary relation. If $(x; w) \in R$, we say x is a *statement* and w is a *witness* for x . We only consider NP relations, i.e., relations R for which a witness w can be verified in time $\text{poly}(|x|)$ for all $(x; w) \in R$. In particular it follows that $|w| = \text{poly}(|x|)$. The set of statements x that admit a witness w is denoted by L_R , i.e., $L_R = \{x : \exists w \text{ s.t. } (x; w) \in R\}$. The set of witnesses for a statement x is denoted by $R(x)$, i.e., $R(x) = \{w : (x; w) \in R\}$.

In the following we give a brief overview of interactive proof systems. For a more thorough treatment, we refer to the full version of this paper [1].

An *interactive proof* $\Pi = (\mathcal{P}, \mathcal{V})$ for relation R is an interactive protocol between two probabilistic polynomial time machines, a prover \mathcal{P} and a verifier \mathcal{V} . Both \mathcal{P} and \mathcal{V} take as public input a statement x and, additionally, \mathcal{P} takes as private input a witness $w \in R(x)$, which is written as $\Pi(x; w)$ or $\text{INPUT}(x; w)$. If all of the verifier's random coins are made public, Π is said to be *public-coin*.

We say an interactive proof is *complete* if \mathcal{V} accepts after every honest execution that takes as input a public-private pair $(x; w) \in R$.

An interactive proof is said to be *knowledge sound* with knowledge error κ , if from every (potentially dishonest) efficient prover P^* that convinces the verifier with probability $\epsilon(x) > \kappa(|x|)$, one can efficiently extract a witness w with $(x; w) \in R$ with probability at least $\epsilon(x) - \kappa(|x|)$.

An interactive proof that is both complete with completeness error $\gamma: \mathbb{N} \rightarrow [0, 1)$ and knowledge sound with knowledge error $\kappa < 1 - \gamma$ is said to be a *Proof or Knowledge* (PoK). PoKs for which knowledge soundness only holds under computational assumptions are also referred to as *Arguments of Knowledge*.

An interactive protocol Π is said to be *special honest verifier zero-knowledge* (SHVZK) if given the challenge by the verifier, one can efficiently simulate accepting transcripts. If simulation is restricted to non-aborting executions of Π , we refer to the protocol as *non-abort special honest verifier zero knowledge*.

A 3-move public-coin protocol is said to be *special sound* if there exists a polynomial time algorithm that on input a statement x and two accepting transcripts (a, c, z) and (a, c', z') , with $c \neq c'$ and common first message a , outputs

a witness $w \in R(x)$. If the algorithm takes as input k transcripts, with pairwise distinct challenges and a common first message, the protocol is *k-special sound*.

A 3-move protocol that is public-coin, complete, k -special sound and SHVZK is said to be Σ -*protocol*.

A (k_1, \dots, k_μ) -*tree of transcripts* for a $(2\mu + 1)$ -move protocol is a set of $K = \prod_{i=1}^\mu k_i$ transcripts arranged in a tree structure, such that the nodes in this tree correspond to the prover’s messages and the edges correspond to the verifier’s challenges, and that further every node at depth i has precisely k_i children corresponding to k_i pairwise distinct challenges. For a graphic representation we refer to Figure 1 of the full version of this paper [1].

A $(2\mu + 1)$ -move public-coin protocol is (k_1, \dots, k_μ) -special sound if there exists an efficient algorithm that on input a (k_1, \dots, k_μ) -tree of accepting transcripts outputs a witness $w \in R(x)$.

2.2 Lattices

A lattice Λ is a discrete additive subgroup of \mathbb{R}^m . The lattice Λ is said to be q -ary if $q\mathbb{Z}^m \subset \Lambda \subset \mathbb{Z}^m$. Let $A \in \mathbb{Z}_q^{k \times m}$, then $\Lambda_q^\perp(A) = \{\mathbf{x} \in \mathbb{Z}^m : A\mathbf{x} = 0 \pmod q\}$ defines a q -ary lattice in \mathbb{Z}^m .

We also consider lattices defined over a ring $\mathcal{R} = \mathbb{Z}[X]/f(X)$, where $f(X)$ is a monic irreducible polynomial of degree d . Via the coefficient embedding norms on \mathbb{C} -vector spaces extend to vectors of ring elements, i.e., for $\mathbf{x} = (x_1, \dots, x_m) \in \mathcal{R}^m$ with $x_i = \sum_{j=1}^d a_{i,j} X^{j-1} \in \mathcal{R}$ we define

$$\|\mathbf{x}\|_2 = \|(a_{1,1}, \dots, a_{m,d})\|_2, \quad \text{and} \quad \|\mathbf{x}\|_\infty = \max_{i,j} |a_{i,j}|.$$

For a prime $q \in \mathbb{N}$, we write $\mathcal{R}_q = \mathbb{Z}[X]/(q, f(X)) = \mathbb{Z}_q[X]/(f(X))$. Let $A \in \mathcal{R}^{k \times m}$, then $\Lambda_q^\perp(A) = \{\mathbf{x} \in \mathcal{R}^m : A\mathbf{x} = 0 \pmod q\}$ defines a q -ary lattice in \mathbb{Z}^{dm} . Finding a non-zero and short element in a lattice $\Lambda_q^\perp(A)$ is referred to as the Module Short Integer Solution (MSIS) problem [33]. The MSIS problem is assumed to be a computationally hard problem.

Definition 1 (MSIS $_{k,m,\beta}$ Problem). *Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(X)$ and let $q \in \mathbb{N}$ be a prime. The MSIS $_{k,m,\beta}$ problem over \mathcal{R}_q is defined as follows. Given a matrix $A \leftarrow_{\mathcal{R}} \mathcal{R}_q^{k \times m}$ sampled uniformly at random, find a non-zero vector $\mathbf{s} \in \mathcal{R}^m$ such that $A\mathbf{s} = 0 \pmod q$ and $\|\mathbf{s}\|_2 \leq \beta$.*

Micciancio and Regev [38] showed that a MSIS-algorithm is expected to output a MSIS solution with norm

$$\|\mathbf{s}\|_2 \geq \min \left(q, 2^{2\sqrt{dk \log \delta \log q}} \right), \tag{1}$$

where δ is the root Hermite factor of the lattice reduction algorithm that is used. In particular, smaller values of δ require better lattice reduction algorithms. In general, $\delta \approx 1.0045$ is assumed to achieve 128-bit computational security [4, 25].

In this work, we will be interested in vectors that are short with respect to the ℓ_∞ -norm. For this reason we also consider the following variant of the MSIS problem, where “shortness” is defined in terms of the ℓ_∞ -norm. Clearly, the hardness of $\text{MSIS}_{k,m,\beta}^\infty$ is implied by the hardness of $\text{MSIS}_{k,m,\sqrt{dm}\beta}$.

Definition 2 (MSIS $_{k,m,\beta}^\infty$ Problem over \mathcal{R}_q). *Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(X)$ and let $q \in \mathbb{N}$ be a prime. The MSIS $_{k,m,\beta}^\infty$ problem over \mathcal{R}_q is defined as follows. Given a matrix $A \leftarrow_R \mathcal{R}_q^{k \times m}$ sampled uniformly at random, find a non-zero vector $\mathbf{s} \in \mathcal{R}^m$ such that $A\mathbf{s} = 0 \pmod q$ and $\|\mathbf{s}\|_\infty \leq \beta$.*

2.3 Commitment Schemes

A commitment scheme allows a prover to create a commitment P to an element x such that the prover can later open P to the committed element x . Informally, a commitment scheme is required to be *binding*, i.e., a prover cannot open a commitment P to two different elements $x \neq y$, and *hiding*, i.e., the commitment P does not reveal any information about the committed vector x . A commitment scheme consists of a setup algorithm, generating the scheme’s public parameters, and a commitment function COM . The commitment function takes as input an element x and randomness γ (and public parameters pp) and outputs a commitment P , i.e., $\text{COM}(x, \gamma) = P$. To open a commitment a prover reveals (x, γ) such that a verifier can verify that $\text{COM}(x, \gamma) = P$. The commitment scheme is said to be *homomorphic* if the commitment function COM (considered respective to fixed public parameters) is a group homomorphism.

The primary commitment scheme of interest to us, described in Definition 3, was already implicit in Ajtai’s seminal work [2]. It allows a prover to commit to a *short* vector $\mathbf{x} \in S_\eta^n = \{\mathbf{y} \in \mathcal{R}^n : \|\mathbf{y}\|_\infty \leq \eta\}$ by sampling $\gamma \leftarrow_R S_\eta^r$ uniformly at random and evaluating the commitment function $P = \text{COM}(x, \gamma)$. Note that, we consider this commitment scheme for secrets and randomness bounded in the ℓ_∞ -norm. We will typically instantiate this commitment scheme with norm bound $\eta = \lceil (p - 1)/2 \rceil$ for some prime $p < q$. This allows a prover to commit to arbitrary vectors in \mathcal{R}_p^n . The properties of this commitment scheme are summarized in Lemma 1 and Lemma 2. Note in particular that by Eq. 1 it follows that the hardness does not depend on the rank n . It follows that the size of a commitment is constant in the rank $m = n + r$; we say that this commitment scheme is *compact*.

Definition 3 (Compact Lattice-Based Commitment Scheme [2]). *Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(x) \in \mathbb{Z}[X]$ of degree d and let $q \in \mathbb{N}$ be a prime. Let $\eta \in \mathbb{N}$ and let $S_\eta = \{x \in \mathcal{R} : \|x\|_\infty \leq \eta\}$. Then, the following setup and commitment algorithms define a commitment scheme:*

- *Setup:* $A_1 \leftarrow_R \mathcal{R}_q^{k \times r}$, $A_2 \leftarrow_R \mathcal{R}_q^{k \times n}$.
- *Commit:* $\text{COM} : S_\eta^n \times S_\eta^r \rightarrow \mathcal{R}_q^k$, $(\mathbf{x}, \gamma) \mapsto A_1\gamma + A_2\mathbf{x} \pmod q$.

Lemma 1 (Hiding). *The commitment scheme of Definition 3 is statistically hiding with statistical security parameter λ , where $\lambda \in \mathbb{N}$ is such that $r \geq \frac{dk \log q + 2\lambda}{d \log(2\eta + 1)}$.*

Lemma 2 (Binding). *The commitment scheme of Definition 3 is binding, conditioned on the hardness of the $\text{MSIS}_{k,n+r,2\eta}^\infty$ -problem over \mathcal{R}_q .*

It is generally hard to construct efficient protocols for proving knowledge of an opening (\mathbf{x}, γ) for a commitment P , i.e., (\mathbf{x}, γ) such that $\text{COM}(\mathbf{x}, \gamma) = P$ and $\|(\mathbf{x}, \gamma)\|_\infty \leq \eta$. For this reason, we introduce the notion of relaxed openings.

Definition 4 ((β, ζ)-Relaxed Commitment Opening). *Let $\beta \in \mathbb{N}$ and $\zeta \in \mathcal{R}$. A (β, ζ) -relaxed opening of a commitment P is a tuple $(\mathbf{x}, \gamma) \in \mathcal{R}^{n+r}$, such that $\text{COM}(\mathbf{x}, \gamma) = \zeta P$ and $\|(\mathbf{x}, \gamma)\|_\infty \leq \beta$.*

Hence, a relaxed opening differs in two ways from a standard commitment opening. First, a relaxed opening for P contains an approximation factor ζ , such that the opening gives a short preimage for ζP instead of the commitment P . Second, the norm-bound β of relaxed openings can be different from the norm bound η on honestly committed vectors (typically $\beta > \eta$).

As long as it is infeasible to find two distinct relaxed openings (\mathbf{x}, γ) and (\mathbf{x}', γ') of a commitment P with $(\mathbf{x}, \gamma) \neq (\mathbf{x}', \gamma')$, proving knowledge of relaxed opening is sufficient in most practical scenarios. In this case, we say the commitment scheme is binding with respect to relaxed openings.

Lemma 3 (Binding with respect to (β, ζ) -Relaxed Openings). *Let $\beta \in \mathbb{N}$ and $\zeta \in \mathcal{R}$. The commitment scheme of Definition 3 is binding with respect to (β, ζ) -relaxed openings, conditioned on the hardness of the $\text{MSIS}_{k,n+r,2\beta}^\infty$ -problem over \mathcal{R}_q .*

3 Multi-round Special Soundness Tightly Implies Knowledge Soundness

In this section we prove that a (k_1, \dots, k_μ) -special sound protocol is *knowledge sound* and give a concrete and tight knowledge error. More precisely, we show the existence of an efficient knowledge extractor. From this it follows that Bulletproofs [14, 16] and Compressed Σ -Protocols [6] are *Proofs/Arguments of Knowledge* (PoKs). We are the first to prove a *tight* bound on the knowledge error. Prior works mainly relied on the asymptotic extractor analysis of [14]. This asymptotic analysis results in conservative concrete security estimates. Moreover, the analysis of [14] is restricted to protocols with exponentially large challenge sets. When the challenge sets are small, such as in lattice based protocols, a refined analysis is required. Our result solves both problems. It gives tight security guarantees resulting in optimal concrete parameters for (k_1, \dots, k_μ) -special sound protocols and it is applicable to protocols with small challenge sets. The main result of this section is summarized in Theorem 1.

Theorem 1 ((k_1, \dots, k_μ) -Special Soundness implies Knowledge Soundness). *Let $\mu, k_1, \dots, k_\mu \in \mathbb{N}$ be such that $K = \prod_{i=1}^\mu k_i$ can be upper bounded by a polynomial. Let $(\mathcal{P}, \mathcal{V})$ be a (k_1, \dots, k_μ) -special sound $(2\mu + 1)$ -move interactive protocol for relation R , where \mathcal{V} samples each challenge uniformly at random from a challenge set of size $N \geq \max_i(k_i)$. Then $(\mathcal{P}, \mathcal{V})$ is knowledge sound with knowledge error*

$$\kappa = \frac{N^\mu - \prod_{i=1}^\mu (N - k_i + 1)}{N^\mu} \leq \frac{\sum_{i=1}^\mu (k_i - 1)}{N}. \tag{2}$$

First, in Sect. 3.1, we considers the special case of 2-special soundness (for which the above implication is well-known). We give a very simple proof that we have not encountered in literature before. In contrast to standard proof techniques, this simplification turns out to be generalizable to the multi-round scenario. Second, in Sect. 3.2, we prove Theorem 1 in its full generality.

3.1 2-Special Soundness

This section is a warm up in which we present a novel proof for the well-known result that 2-special soundness implies knowledge soundness. Later we show that our techniques generalize to prove a similar result for $2\mu + 1$ -move protocols that are (k_1, \dots, k_μ) -special sound. We make a minor modification to the “collision-game” defined in [20]. The knowledge extractor essentially plays this game in order to extract a collision of two accepting transcripts (a, c, z) and (a, c', z') with common first message a . By the special soundness property a witness can be computed efficiently given this collision. Our modification increases the success probability of the knowledge extractor of [20] from $(\epsilon(x) - \kappa(|x|))^2$ to $\epsilon(x) - \kappa(|x|)$, where $\kappa(|x|)$ is the knowledge error and $\epsilon(x)$ the success probability of the prover for a statement x . In contrast to the extractor of [20], which runs in *strict* polynomial time, our extractor runs in *expected* polynomial time. However, this is sufficient for proving knowledge soundness.

If the input x is clear from context, we simply write ϵ to denote $\epsilon(x)$. All other parameters will implicitly depend on $|x|$ (e.g., we denote $\kappa(|x|)$ by κ).

A similar result can be found in [29]. However, our approach significantly simplifies the knowledge extractor and its analysis. For instance, the extractor of [29] is composed of two algorithms considering different scenarios, whereas this case distinction is not required in our knowledge extractor. This simplification will allow for a generalization to the (k_1, \dots, k_μ) -special sound case.

The collision game. Let us now describe the game. We consider a binary matrix $H \in \{0, 1\}^{R \times N}$. The R rows correspond to the prover’s randomness and the N columns correspond to the verifier’s randomness, i.e., the verifier samples a challenge uniformly at random from a challenge set of size N . An entry of H equals 1 if and only if the corresponding protocol transcript is accepting.

The idea of the knowledge extractor is to sample elements from H until two 1-entries in the same row are found. The ij -th entry of H can be obtained

by executing the prover with fixed randomness corresponding to the i -th row and verifier’s challenge corresponding to the j -th column, and checking if the resulting transcript would be accepted. As the prover’s randomness is fixed along one row, finding two 1-entries in the same row corresponds to two finding two accepting transcripts (a, c, e) and (a, c', e') , which by the 2-special soundness allows to extract a witness. The difference to the knowledge extractor of [29] is the following:

1. Our knowledge extractor checks one entry of H (for position ij sampled at random), *and aborts if this is not a 1-entry.*
2. If the first entry was a 1-entry, our knowledge extractor then samples along row i *without replacement.*

More precisely, the knowledge extractor will play the following collision-game. An entry of H is selected uniformly at random. If this entry equals 1, continue sampling different elements from this row (without replacement) until a second 1-entry is found or until the row has been exhausted. If the first entry does not equal 1, the game aborts. The collision game outputs success if and only if two 1-entries in the same row have been found.

In contrast the above collision-game, the collision-game of [20] simply checks 2 random entries of H and outputs success if both of them are 1-entries.

Lemma 4 (Collision-Game). *Let $H \in \{0, 1\}^{R \times N}$ and let ϵ denote the fraction of 1-entries in H . The expected number of H -entries queried in the collision-game defined above is at most 2. Moreover, the success probability of the collision-game is greater than or equal to $\epsilon - 1/N$.*

Proof. Expected Number of Queries. Let ϵ_i be the fraction of 1-entries in row i . Assuming that the first entry lies in row i and equals 1, the remainder of the collision game can be modeled by a negative hypergeometric distribution. Elements from a population of size $N - 1$, containing $\epsilon_i N - 1$ 1-entries, are drawn (without replacement) until a second 1-entry has been found. The expected number of draws equals $(N - 1 + 1)/(\epsilon_i N - 1 + 1) = 1/\epsilon_i$ if $\epsilon_i > 1/N$ (see the full version of this paper [1]). If there is no second 1-entry in the row, then the number of draws is always equal to $N - 1$. Hence, the expected number of draws can be upper bounded by $1/\epsilon_i$. The expected number of H -entries queried is therefore at most

$$\frac{1}{R} \sum_{i=1}^R \left(1 + \epsilon_i \frac{1}{\epsilon_i} \right) = 2.$$

Success Probability. The collision-game succeeds if the first entry is a 1 that lies in a row containing at least two 1-entries. For $0 \leq k \leq N$, let δ_k be the fraction of rows with exactly k 1-entries. Then the success probability equals

$$\sum_{k=2}^N \frac{k}{N} \delta_k = \left(\sum_{k=0}^N \frac{k}{N} \delta_k \right) - \frac{\delta_1}{N} \geq \epsilon - 1/N,$$

which proves the second part of the lemma. □




Fig. 1. We say a (k_1, \dots, k_μ) -tree as depicted above is a (k_1, \dots, k_μ) -tree of 1-entries in H , if $H(a, c_1^1, c_2^{1,1}, \dots, c_\mu^{1, \dots, 1}) = H(a, c_1^1, c_2^{1,1}, \dots, c_\mu^{1, \dots, 2}) = \dots = H(a, c_1^{k_1}, c_2^{k_1, k_2}, \dots, c_\mu^{k_1, \dots, k_\mu}) = 1$.

From Lemma 4 it immediately follows that 2-special soundness implies knowledge soundness with knowledge error $1/N$.

Corollary 1. *Let $(\mathcal{P}, \mathcal{V})$ be a special sound 3-move interactive protocol for relation R , where \mathcal{V} samples each challenge uniformly at random from a challenge set of size $N \geq k$. Then $(\mathcal{P}, \mathcal{V})$ is knowledge sound with knowledge error $\kappa = 1/N$.*

Remark 1. Lemma 4 has a straightforward generalization to the k -special soundness scenario. In this generalization the collision game draws until it has obtained k , instead of 2, 1-entries in the same row. Hence, it again involves a negative hypergeometric distribution, but now with different parameters. In this case, the expected number of queries is at most k and the success probability is greater than or equal to $\epsilon - (k - 1)/N$.

3.2 (k_1, \dots, k_μ) -Special Soundness

In this section, we generalize the collision-game of Sect. 3.1 to the (k_1, \dots, k_μ) -special soundness scenario.

The (k_1, \dots, k_μ) -collision game. To define the (k_1, \dots, k_μ) -collision-game, let $H \in \{0, 1\}^{R \times N \times \dots \times N}$ be a $(\mu + 1)$ -dimensional binary matrix. For $a \in \{1, \dots, R\}$ and $c_1, \dots, c_i \in \{1, \dots, N\}$, we let $H(a, c_1, \dots, c_i) \in \{0, 1\}^{N \times \dots \times N}$ be the $(\mu - i)$ dimensional submatrix of H that contains all entries of H for which the first $i + 1$ coordinates are equal to (a, c_1, \dots, c_i) . The first dimension corresponds to

the prover’s randomness and the other dimensions correspond to the verifier’s random choices, i.e., we consider protocols in which the verifier samples all μ challenges uniformly at random from a challenge set of size N . For a fixed public input x , we define the matrix H such that $H(a, c_1, \dots, c_\mu) = 1$ if and only if a transcript with prover’s randomness a and verifier’s challenges c_1, \dots, c_μ will lead to an accepting transcript.

In Sect. 2, we have defined (k_1, \dots, k_μ) -trees of accepting transcripts for $(2\mu + 1)$ -move protocols. Similarly, we define (k_1, \dots, k_μ) -trees of 1-entries in matrix H (Fig. 1). Such trees can be defined recursively as follows. For $\mu = 0$, a tree of 1-entries is simply a 1-entry in H . For arbitrary μ , a (k_1, \dots, k_μ) -tree is the union of k_1 (k_2, \dots, k_μ) -trees in $H(a, c_1), \dots, H(a, c_{k_1})$, respectively, for a fixed a and pairwise distinct c_i . Hence, a (k_1, \dots, k_μ) -tree of 1-entries in matrix H is a set of $K = \prod_{i=1}^\mu k_i$ 1-entries that are in a (k_1, \dots, k_μ) -tree structure.

We define TREE to be the algorithm playing the (k_1, \dots, k_μ) -collision-game. By playing this game TREE aims to find a (k_1, \dots, k_μ) -tree of 1-entries in matrix H . The algorithm TREE is defined recursively as follows. On input $a \in \{1, \dots, R\}$ and $c_1, \dots, c_\mu \in \{1, \dots, N\}$, $\text{TREE}_\mu(a, c_1, \dots, c_\mu)$ successfully outputs $H(a, c_1, \dots, c_\mu)$ if this entry equals 1 and it aborts otherwise. For $0 \leq i \leq \mu - 1$ and on input $a \in \{1, \dots, R\}$ and $c_1, \dots, c_i \in \{1, \dots, N\}$, $\text{TREE}_i(a, c_1, \dots, c_i)$ aims to find a (k_{i+1}, \dots, k_μ) -tree of 1-entries in matrix $H(a, c_1, \dots, c_i)$. The algorithm $\text{TREE}_i(a, c_1, \dots, c_i)$ proceeds by sampling $c_{i+1} \in \{1, \dots, N\}$ uniformly at random and running $\text{TREE}_{i+1}(a, c_1, \dots, c_{i+1})$. If this instantiation of TREE_{i+1} aborts the algorithm $\text{TREE}_i(a, c_1, \dots, c_i)$ aborts. Otherwise it continues sampling different c_{i+1} ’s (i.e., without replacement) until it has found k_{i+1} (k_{i+2}, \dots, k_μ) -trees of 1-entries or until it has exhausted all possible c_{i+1} ’s. In the latter case $\text{TREE}_i(a, c_1, \dots, c_i)$ aborts, in the former case $\text{TREE}_i(a, c_1, \dots, c_i)$ outputs a (k_{i+1}, \dots, k_μ) -tree of 1-entries in matrix $H(a, c_1, \dots, c_i)$.

The (k_1, \dots, k_μ) -collision-game samples $a \in \{1, \dots, R\}$ uniformly at random and runs $\text{TREE}_0(a)$. If $\text{TREE}_0(a) = \perp$ it aborts and otherwise it outputs a (k_1, \dots, k_μ) -tree of 1-entries in $H(a)$. The following lemma gives the expected run-time and success probability of the tree finding algorithm TREE. For a proof of the following lemma, we refer to the full version of this paper [1].

Lemma 5 ((k_1, \dots, k_μ) -Tree Finding Algorithm). *Let $H \in \{0, 1\}^{R \times N \times \dots \times N}$ be a $(\mu + 1)$ -dimensional matrix and let ϵ denote the fraction of 1-entries in H . The expected number of entries queried by the (k_1, \dots, k_μ) -tree finding algorithm TREE defined above is at most $K = \prod_{i=1}^\mu k_i$. Moreover, TREE successfully outputs a (k_1, \dots, k_μ) -tree of 1-entries in H with probability at least*

$$\epsilon - \frac{N^\mu - \prod_{i=1}^\mu (N - k_i + 1)}{N^\mu} \geq \epsilon - \frac{\sum_{i=1}^\mu (k_i - 1)}{N}.$$

A knowledge extractor, with rewindable black-box access to a possible dishonest prover \mathcal{P}^* , essentially runs this tree finding algorithm to obtain a (k_1, \dots, k_μ) -tree of accepting transcripts. It evaluates one protocol interaction with \mathcal{P}^* and

recursively rewinds \mathcal{P}^* , fixing its internal randomness and following the tree finding strategy of TREE. By the (k_1, \dots, k_μ) -special soundness property a witness can then be extracted efficiently from the obtained (k_1, \dots, k_μ) -tree of accepting transcripts. Hence, from Lemma 5 it immediately follows that a (k_1, \dots, k_μ) -special sound protocol is knowledge sound with knowledge error κ , where

$$\kappa = \frac{N^\mu - \prod_{i=1}^\mu (N - k_i + 1)}{N^\mu} \leq \frac{\sum_{i=1}^\mu (k_i - 1)}{N}.$$

The latter inequality follows since we have $N \geq \max_i(k_i)$ and thus $\prod_{i=1}^\mu (N - k_i + 1) \leq N^\mu - N^{\mu-1} \sum_{i=1}^\mu (k_i - 1)$. This proves Theorem 1.

3.3 Tightness of Our Extraction Analysis

The knowledge error κ of Theorem 1 is optimal, i.e., there exists a dishonest prover that succeeds in cheating with probability κ . Typically a dishonest prover can cheat in a k -special sound protocol by guessing a set of $k - 1$ challenges and hoping that the verifier selects one of these challenges. The success probability of this attack is equal to $(k - 1)/N$, where N is the size of the challenge set. More generally, a cheating strategy for a (k_1, \dots, k_μ) -special sound $(2\mu + 1)$ -move protocol goes as follows. For every round i , the cheating prover guesses a set of $k_i - 1$ challenges. The cheating prover succeeds if there exists a round i for which the verifier chooses one of the $k_i - 1$ challenges guessed by the prover. The success probability of this attack is easily seen to be equal to the knowledge error κ . Hence, this knowledge error is optimal. Alternatively, we observe that there exist matrices H with $\epsilon = \kappa$, i.e., for which the fraction of 1-entries equals κ , that do not contain a (k_1, \dots, k_μ) -tree of 1-entries.

Moreover, the tree finding algorithm is optimal in the following sense. The expected number of H -entries that are queried is exactly equal to the number of entries in a tree. Hence, we can not hope to find a tree faster than this. Moreover, taking a closer look at the proof of Lemma 5 shows that the success probability actually has the following lower bound

$$f(\epsilon) = \left(\prod_{j=1}^\mu \frac{N}{N - k_j + 1} \right) (\epsilon - \kappa).$$

Hence, if $\epsilon = 1$ the success probability of TREE is at least $f(1) = 1$, which is what we would expect.

3.4 A Note on Witness Extended Emulation

Lindell showed that a technical issue arises when using Proofs of Knowledge as subprotocols in larger cryptographic protocols [34]. To prove security of the compound protocol, a simulator is typically required to run the extractor of the PoK. However, the naive simulation approach does not necessarily run in polynomial time. To this end, Lindell defined the notion of *witness-extended*

emulation (WEE), capturing precisely the properties required when using PoKs as subprotocols. Moreover, he showed that any PoK has WEE, thereby solving this technical issue for all PoKs at once. Hence, from our extraction analysis it follows that any (k_1, \dots, k_μ) -special sound protocol has WEE.

Previously, there was no proof showing that a (k_1, \dots, k_μ) -special sound protocol is knowledge sound. For this reason prior works (e.g., [14]) resorted to proving witness-extended emulation directly. However, these results are non-tight and only apply to protocols with exponentially large challenge sets.

4 Decreasing the Knowledge Error of Public-Coin Interactive Protocols

In this section, we establish a novel parallel repetition theorem showing that the knowledge error can be decreased by repeating the protocol in parallel.

We want the knowledge error of a PoK to be negligible in the security parameter. If this is not the case the protocol is typically repeated, say t times. The verifier of the composed protocol only accepts if all t instances of the basic protocol are accepted. Ideally, and perhaps intuitively, this approach reduces the knowledge error from κ down to κ^t . This is indeed the case if the repetitions are executed sequentially [27]. However, sequential repetition increases the round complexity. Since the security loss due to the Fiat-Shamir transformation increases exponentially in the number of rounds [23], this is unacceptable when considering the non-interactive instantiations of our protocols (see the full version of this paper [1]). Further, also in the interactive setting we would like to avoid the additional round complexity introduced by sequential composition.

For this reason, we aim to repeat the protocol in parallel. We write $(\mathcal{P}^t, \mathcal{V}^t)$ for the t -fold parallel repetition of an interactive argument $(\mathcal{P}, \mathcal{V})$. However, it is not true in general that parallel repetition decreases the knowledge error exponentially. There even exist interactive protocols for which parallel repetition does not decrease the success probability of a dishonest prover at all [10, 39]. Analyzing parallel repetitions is significantly more complicated than analyzing sequential repetitions, because a dishonest prover does not have to treat all t parallel instances independently, i.e., a message corresponding to a specific instance may depend on the messages and challenges of the other parallel instances.

If $(\mathcal{P}, \mathcal{V})$ is a 2-special sound 3-move protocol, then $(\mathcal{P}^t, \mathcal{V}^t)$ is 2-special sound too. It therefore follows that the knowledge error of a 2-special sound protocol decreases exponentially in the number of parallel repetitions. However, a similar result does not hold in general, i.e., in general special-soundness is not preserved by parallel repetition. For example, it is easily seen that the parallel repetition of a k -special sound protocol for $k \neq 2$ is not k -special-sound.

Several parallel repetition results, considering multi-round public-coin interactive arguments, have been established [18, 19, 28], showing that parallel repetition reduces the soundness error. However, “soundness” is a weaker notion than “knowledge soundness”. Informally the soundness error is the success probability of a cheating prover and soundness does not require the existence of a knowledge extractor.

To the best of our knowledge a parallel repetition result for decreasing the *knowledge error* has not been established yet, even though the lattice-based Bulletproof protocols of [15] implicitly rely on such a parallel repetition result. In Theorem 3, we show that the knowledge error of a public-coin argument decreases close to exponentially in the number of parallel repetitions. Our proof uses the following result from [19]. This theorem shows that, given oracle access to a (possibly dishonest) prover \mathcal{P}^* that, for statements x , succeeds in convincing \mathcal{V}^t with probability $\epsilon(x)$, a prover $\mathfrak{P}^{(\mathcal{P}^*)}$ that succeeds in convincing \mathcal{V} with probability $\approx \epsilon(x)^{1/t}$ can be constructed.

Theorem 2 (Theorem 2 of [19]). *Let $(\mathcal{P}, \mathcal{V})$ be a public-coin interactive argument for a language L . Let $t: \mathbb{N} \rightarrow \mathbb{N}$, and let $(\mathcal{P}^t, \mathcal{V}^t)$ be the t -fold parallel repetition of $(\mathcal{P}, \mathcal{V})$. There exists an oracle machine $\mathfrak{P}^{(\cdot)}$ such that for every $\xi: \mathbb{N} \rightarrow (0, 1)$, every $\delta: \{0, 1\}^* \rightarrow (0, 1)$, every $x \in \{0, 1\}^*$, and every PPT prover \mathcal{P}^* , it holds that if*

$$\Pr((\mathcal{P}^*, \mathcal{V}^t)(x) = 1) \geq \underbrace{(1 + \xi(|x|))\delta(x)^{t(|x|)}}_{\epsilon(x):=},$$

then

$$\Pr((\mathfrak{P}^{(\mathcal{P}^*)}, \mathcal{V})(x) = 1) \geq \delta(x).$$

Furthermore, $\mathfrak{P}^{(\mathcal{P}^*)}$ runs in time $\text{poly}(|x|, t(|x|), \xi(|x|)^{-1}, \epsilon(x)^{-1}, (1 - \delta(x))^{-1})$.

Theorem 3 now shows that the t -fold parallel repetition of knowledge sound interactive argument is knowledge sound and that the knowledge error decreases close to exponential in t . More precisely, the theorem shows that if $(\mathcal{P}, \mathcal{V})$ has knowledge error κ , then $(\mathcal{P}^t, \mathcal{V}^t)$ has knowledge error $\kappa^t + \nu$, for arbitrary noticeable ν . Therefore, by choosing t large enough, we can show that $(\mathcal{P}^t, \mathcal{V}^t)$ has knowledge error $1/|x|^c$ for any $c \in \mathbb{N}$. Note though that we cannot show that $(\mathcal{P}^t, \mathcal{V}^t)$ has negligible knowledge error $\text{negl}(\lambda)$, because the running time of $\mathfrak{P}^{(\mathcal{P}^*)}$ scales with the inverse success probability of \mathcal{P}^* .

While it might seem that this barrier is rather an artifact of the proof technique of [19] on which we build, it was shown by [22] that Theorem 2 is tight when considering soundness amplification of protocols in general. More precisely, based on some cryptographic assumptions they showed that parallel repetition does not amplify security beyond negligible, meaning that for any negligible function negl one can find an instantiation that when starting with non-negligible soundness error, the protocol can always be broken with probability $\text{negl}(|x|)$, no matter how many parallel repetitions one runs.

For a proof of the theorem we refer to the full version of this paper [1].

Theorem 3. *Let $(\mathcal{P}, \mathcal{V})$ be a public-coin interactive argument for a relation R that is knowledge sound with knowledge error $\kappa: \mathbb{N} \rightarrow (0, 1)$. Let $t: \mathbb{N} \rightarrow \mathbb{N}$ be upper bounded by a polynomial. Let $\nu: \mathbb{N} \rightarrow (0, 1)$ be an arbitrary noticeable function. Then, $(\mathcal{P}^t, \mathcal{V}^t)$ is knowledge sound with knowledge error $\kappa' = \kappa^t + \nu$.*

Remark 2. The properties *completeness* and *special honest verifier zero-knowledge* are easily seen to be preserved by parallel repetition, although the completeness error increases in the number parallel repetitions.

5 A General Framework for Compressed Σ -Protocols over Lattices

The main pivot of compressed Σ -protocol theory [6] is a basic Σ -protocol for proving that a committed vector satisfies some linear constraint. Subsequently, a compression mechanism is applied (recursively) to reduce the communication complexity from linear down to polylogarithmic in the input size. The composition of these protocols is referred to as a compressed Σ -protocol. In this section we present a natural abstraction similar to the one presented in [7, Appendix A] extended to the lattice setting. This requires a number of non-trivial adaptations that are explained in the following. Subsequently, we show how to instantiate this abstraction from a concrete lattice assumption.

In the following we first give an abstraction of the standard Σ -protocol to the lattice setting and then explain how the compression mechanism extends to this setting. Note that we give both protocols in a very abstract fashion, with the goal of allowing to instantiate them from a broad variety of lattice-based assumptions. Note that our abstraction is not restricted to instantiations based on lattices, but is tailored to this setting.

5.1 Standard Σ -Protocol

In this section we recall what we will refer to as *standard Σ -protocol* for proving knowledge of a preimage of some given module homomorphism Ψ .¹ This protocol can be viewed as the abstraction of the protocol of Schnorr [40] to arbitrary module homomorphisms, where we have to build in several relaxations in order to make it compatible with the lattice setting.

First, in the lattice setting the witness is required to be small, we therefore define a pair $(Y; y)$ to be in the target relation if $Y = \Psi(y)$ and $\|y\| \leq \alpha$, for some $\alpha \in \mathbb{N}$. Note that this requires to define a norm in the preimage space, we therefore in the following restrict to modules with norm. If the preimage is not required to be small (as, e.g., is the case in the discrete log setting), one does not have to require a norm on the module and can simply ignore the corresponding requirements in the protocols. The requirement of the witness y to have small norm is also where the main difficulty stems from, because one now has to transform a witness y into a witness x , such that

1. the norm of x is not much larger than y (as otherwise the statement becomes meaningless), but
2. x still hides y .

¹ For an introduction into modules and module homomorphisms we refer to [32].

In order to ensure the second without a too large knowledge error, the relation that one can prove knowledge of does not correspond to the target relation R , but some relaxed relation R' . In this case, we say the protocol is a protocol for the pair of relations (R, R') , i.e., an honest prover knows a witness for R but can only prove knowledge of a witness for R' .

In fact, there are two sources introducing “soundness slack”: First, x itself will in general already have larger norm than y (in order to ensure hiding). Second, even worse, extracting a witness \tilde{y} from two accepting transcripts, introduces additional slack. This slack is more difficult to control, as it depends on the inverse of challenge differences. As challenge differences will not necessarily be invertible over the underlying ring, we introduce an additional relaxation on the relation. Namely, for some fixed element ζ (in our examples, we will typically have that ζ is a power of two) we will consider relations R' , such that $(X; x) \in R'$ if $\Psi(x) = \zeta \cdot X$ and $\|x\| \leq \beta$. We refer to ζ as an approximation factor.

More formally, let $\mathcal{R} = \{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of rings, let $M = \{M_\lambda\}_{\lambda \in \mathbb{N}}, N = \{N_\lambda\}_{\lambda \in \mathbb{N}}$ be ensembles of \mathcal{R} -modules, let $\Psi = \{\Psi_\lambda: M_\lambda \rightarrow N_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of efficiently computable \mathcal{R} -module homomorphisms and let $\zeta = \{\zeta_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of approximation factors (i.e., $\zeta_\lambda \in \mathcal{R}_\lambda$ for all λ). Let further $\|\cdot\|$ be a norm on M , let $\alpha, \beta: \mathbb{N} \rightarrow \mathbb{N}$ with $\alpha \leq \beta$. Then, we define the relations $R(\Psi, \alpha) = \{R_\lambda(\Psi, \alpha)\}_{\lambda \in \mathbb{N}}$ and $R(\Psi, \beta, \zeta) = \{R_\lambda(\Psi, \beta, \zeta)\}_{\lambda \in \mathbb{N}}$ via

$$R_\lambda(\Psi, \alpha) = \left\{ (Y; y) : y \in M_\lambda, Y = \Psi_\lambda(y), \|y\| \leq \alpha(\lambda) \right\},$$

$$R_\lambda(\Psi, \beta, \zeta) = \left\{ (Y; y) : y \in M_\lambda, \zeta_\lambda \cdot Y = \Psi_\lambda(y), \|y\| \leq \beta(\lambda) \right\}.$$

In the following we abstract the notion of *rejection sampling* [35,36], which is used in lattice based cryptography to sample a value, such that

1. the sample algorithm is somewhat norm-preserving, i.e., the norm of the sampled value is not too much larger than the norm of the witness,
2. adding this value to the witness statistically hides the witness or the rejection sampling strategy aborts, and, finally,
3. the abort probability is essentially independent of the witness.

Definition 5 (V-Hiding and β -Bounded Sampling). *Let $\mathcal{R} = \{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of rings and let $M = \{M_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of \mathcal{R} -modules. Let $V = \{V_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of sets with $V_\lambda \subseteq M_\lambda$ for all λ . Let $(\mathcal{D}, \mathcal{F})$ such that \mathcal{D} is an ensemble of efficiently sampleable distributions $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ over M , and \mathcal{F} a PPT algorithm. We say $(\mathcal{D}, \mathcal{F})$ -is V -hiding, if there exists a PPT algorithm \mathcal{F}' such that for each $\lambda \in \mathbb{N}$:*

- \mathcal{F} on input $r \in M_\lambda$ and $v \in V_\lambda$, outputs $r + v$ or \perp ,
- \mathcal{F}' on input 1^λ , outputs an element $z \in M_\lambda$ or \perp ,

such that the output distributions of $(\mathcal{D}, \mathcal{F})$ and \mathcal{F}' are statistically close. More precisely, there exists a negligible function $\text{negl}: \mathbb{N} \rightarrow \mathbb{N}$ such that for all $\lambda \in \mathbb{N}$ and for all $v \in V_\lambda$ we have

$$\Delta(\{\mathcal{F}(r, v) \mid r \leftarrow \mathcal{D}_\lambda\}, \{\mathcal{F}'(1^\lambda)\}) \leq \text{negl}(\lambda),$$

where the probability is taken over the randomness of \mathcal{D}_λ and the random coins of $\mathcal{F}, \mathcal{F}'$. If the distribution of $(\mathcal{D}, \mathcal{F})$ and \mathcal{F}' are equal, we say $(\mathcal{D}, \mathcal{F})$ -is perfectly V -hiding.

Note that by the above considerations we can upper bound the abort probability of $(\mathcal{D}, \mathcal{F})$ by

$$\delta(\lambda) = \Pr[\mathcal{F}'(1^\lambda) = \perp] + \text{negl}(\lambda),$$

for all $\lambda \in \mathbb{N}$.

Let further $\beta: \mathbb{N} \rightarrow \mathbb{N}$. We say that $(\mathcal{D}, \mathcal{F})$ is β -bounded if for all $\lambda \in \mathbb{N}$, $v \in V_\lambda$ and r in the support of \mathcal{D}_λ it holds $\|\mathcal{F}(r, v)\| \leq \beta(\lambda)$ whenever $\mathcal{F}(r, v) \neq \perp$.

To improve readability, we will in the following omit the security parameter, and, e.g., simply say “Let \mathcal{R} be a ring...”, or “Let $\alpha \in \mathbb{N}...$ ”, even though we assume all variables to be parametrized by the security parameter.

Before stating the Σ -protocol, we introduce the notion of an ζ -exceptional subset, which will ensure that the protocol satisfies special soundness.

Definition 6 (ζ -Exceptional Subset). Let \mathcal{R} be a ring, $\zeta \in \mathcal{R}$ and $\mathcal{C} \subseteq \mathcal{R}$ be a set. We say \mathcal{C} is an ζ -exceptional subset of \mathcal{R} , if for all pairs of distinct elements $c, c' \in \mathcal{C}$ there exists a non-zero element $a \in \mathcal{R}$ such that $a(c - c') = \zeta$. If \mathcal{C} is a 1-exceptional subset of \mathcal{R} , we simply say that \mathcal{C} is an exceptional subset.

We further need to give bounds on the soundness slack introduced by extraction. To this end, for ζ -exceptional subsets $\mathcal{C} \subset \mathcal{R}$ we define $w(\mathcal{C})$ and $\bar{w}(\mathcal{C}, \zeta)$:

$$\begin{aligned} w(\mathcal{C}) &= \max_{c \in \mathcal{C}, x \in \mathcal{R} \setminus \{0\}} \frac{\|cx\|}{\|x\|}, \\ \bar{w}(\mathcal{C}, \zeta) &= \max_{c \neq c' \in \mathcal{C}, x \in \mathcal{R} \setminus \{0\}} \max_{a \in \mathcal{R}: a(c-c') = \zeta} \frac{\|ax\|}{\|x\|}. \end{aligned} \tag{3}$$

The value $w(\mathcal{C})$ gives an upper bound on how much the norm of an element in \mathcal{R} increases when multiplied by an element in \mathcal{C} , i.e., $w(\mathcal{C})$ is such that $\|cx\| \leq w(\mathcal{C})\|x\|$ for all $c \in \mathcal{C}$ and $x \in \mathcal{R}$. Note that if $\mathcal{R} = \mathbb{Z}$ and with absolute value $|\cdot|$, we simply have $w(\mathcal{C}) = \max\{|c|: c \in \mathcal{C}\}$.

The value $\bar{w}(\mathcal{C}, 1)$ gives an upper bound on how much the norm of an element in \mathcal{R} increases when multiplied with the inverse of challenge differences, i.e., $\bar{w}(\mathcal{C}, 1)$ is such that $\|(c-c')^{-1}x\| \leq \bar{w}(\mathcal{C}, 1)\|x\|$ for all $x \in \mathcal{R}$ and distinct $c, c' \in \mathcal{C}$. In general, the value $\bar{w}(\mathcal{C}, \zeta)$ gives an upper bound on how much the norm of an element in \mathcal{R} increases when multiplied with an a such that $a(c - c') = \zeta$ for challenges $c \neq c'$. Note that $\bar{w}(\mathcal{C}, \zeta)$ is only well-defined if \mathcal{C} is ζ -exceptional.

The maximum over $a \in \mathcal{R}$ in Eq. 3 can be replaced by a minimum, potentially resulting in tighter norm bounds. More precisely, the extractor can choose the element a that minimizes $\|ax\|/\|x\|$. However, this requires the minimum to be efficiently computable. To avoid this additional assumption we take the maximum over all a . Moreover, in most practical applications \mathcal{R} does not have zero-divisors and $a \in R$ is uniquely defined.

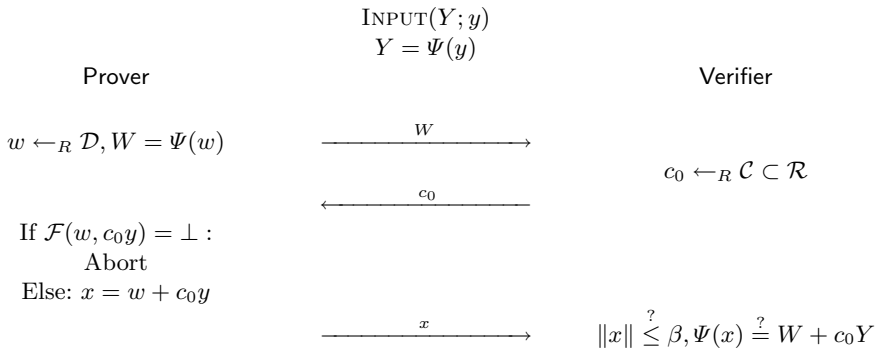
For a module M over \mathcal{R} with norm $\|\cdot\|$, similarly we define

$$w_M(\mathcal{C}) = \max_{c \in \mathcal{C}, x \in M \setminus \{0\}} \frac{\|cx\|}{\|x\|} \text{ and } \bar{w}_M(\mathcal{C}, \zeta) = \max_{c \neq c' \in \mathcal{C}, x \in M \setminus \{0\}} \max_{a \in \mathcal{R}: a(c-c') = \zeta} \frac{\|ax\|}{\|x\|}.$$

Note that for $M = \mathcal{R}^n$ and $\|\cdot\|$ over M defined as ℓ_p -norm (for $p \in \mathbb{N} \cup \{\infty\}$), we have $w_M(\mathcal{C}) = w(\mathcal{C})$ and $\bar{w}_M(\mathcal{C}, \zeta) = \bar{w}(\mathcal{C}, \zeta)$.

We now state the standard Σ -protocol Π_0 for the pair of relations $(R(\Psi, \alpha), R(\Psi, 2\beta\sigma, \zeta))$ in Protocol 1. Further, we summarize its properties in Theorem 4. For a proof we refer to the full version of this paper [1].

Protocol 1 Standard Σ -Protocol Π_0 for the pair of relations $(R(\Psi, \alpha), R(\Psi, 2\beta\sigma, \zeta))$, where $\sigma = \bar{w}_M(\mathcal{C}, \zeta)$. Here, $(\mathcal{D}, \mathcal{F})$ is V -hiding and β -bounded, where $V = \{cy \mid y \in M, \|y\| \leq \alpha, c \in \mathcal{C}\}$.



Theorem 4 (Standard Σ -Protocol). *Let \mathcal{R} be a ring, let M, N be \mathcal{R} -modules and let $\Psi: M \rightarrow N$ be an efficiently computable \mathcal{R} -module homomorphism.*

Further, let $\zeta \in \mathcal{R}$ and $\mathcal{C} \subset \mathcal{R}$ be a finite ζ -exceptional subset of \mathcal{R} , let $\alpha, \beta \in \mathbb{N}$ and $\delta \in [0, 1)$, let $V = \{cy \mid y \in M, \|y\| \leq \alpha, c \in \mathcal{C}\}$ and let $(\mathcal{D}, \mathcal{F})$ be a β -bounded V -hiding distribution with abort probability δ .

Then, the protocol Π_0 (as defined in Protocol 1) is a 3-move protocol for relations $(R(\Psi, \alpha), R(\Psi, 2\beta\sigma, \zeta))$ defined via

$$R(\Psi, \alpha) = \left\{ (Y; y) : y \in M, Y = \Psi(y), \|y\| \leq \alpha \right\},$$

$$R(\Psi, 2\beta\sigma, \zeta) = \left\{ (Y; y) : y \in M, \zeta \cdot Y = \Psi(y), \|y\| \leq 2\beta\sigma \right\},$$

where $\sigma = \bar{w}_M(\mathcal{C}, \zeta)$.

It is complete with completeness error δ , unconditionally 2-special sound and statistical non-abort special honest verifier zero-knowledge.

Remark 3. In some settings it is beneficial to introduce another relaxation. For example, if $\zeta = 1$ (i.e., if challenge difference are invertible), the aforementioned approach requires *inverses* of challenge differences to be of small norm. The following relaxed relation only requires challenge differences, and not necessarily their inverses, to be of small norm. It introduces an adapted approximation factor $\bar{c} \in \bar{\mathcal{C}} = \{c - c'; c, c' \in \mathcal{C}, c \neq c'\}$ and is defined as follows

$$R(\Psi, \beta, \bar{\mathcal{C}}) = \left\{ (Y; y, \bar{c}) : y \in M, \bar{c} \cdot Y = \Psi(y), \|y\| \leq \beta, \bar{c} \in \bar{\mathcal{C}} \right\}.$$

The approximation factor \bar{c} is not fixed and part of the secret witness. This relaxation allows for more efficient Σ -protocols. However, when composed with other protocols the fact that the approximation factors are not fixed introduces additional difficulties. These can be handled, but in most settings the required adjustments negate the benefits of this relaxed relation, we therefore do not consider it further.

For a generic transformation from non-abort SHVZK to SHVZK (or even standard zero-knowledge) we refer to the full version of this paper [1].

5.2 Compression Mechanism

Observe that the final message x of protocol Π_0 is a witness for statement $X := W + c_0Y$, i.e., the final message can be viewed as a trivial proof of knowledge for $X \in L_{R(\Psi, \beta)}$. In the following, we will present a general view on the compression mechanism that allows to replace this trivial PoK by a more efficient one, using Bulletproof’s folding mechanism [14, 16]. This protocol does not need to be SHVZK, since it is a replacement for the trivial PoK.

Compression function. The Bulletproof folding mechanism relies on an compression function that allows to compress the witness iteratively. In the following, we outline the properties the compression function has to satisfy. The main purpose of giving this abstraction is to improve readability of the protocols. In the full version of this paper [1], we further give an abstraction generalizing to larger compression rate and the corresponding compression mechanism.

Definition 7 (Extractable compression function). *Let M, M' be \mathcal{R} -modules, such that M is of even rank n and M' of rank $n/2$. Let $\mathcal{C} \subset \mathcal{R}$ be an exceptional subset of \mathcal{R} . Let $\text{Comp} = \{\text{Comp}_c : M \rightarrow M' : c \in \mathcal{C}\}$ and $\Phi = \{\Phi_c : M' \rightarrow M : c \in \mathcal{C}\}$, where Φ_c is an \mathcal{R} -module homomorphism for each $c \in \mathcal{C}$. Then, we say (Comp, Φ) is an extractable compression function for \mathcal{C} , if the following holds: There exist maps $\pi_L, \pi_R : M \rightarrow M$, such that for all $c \in \mathcal{C}$:*

$$\Phi_c(\text{Comp}_c(x)) = \pi_L(x) + c \cdot x + c^2 \cdot \pi_R(x).$$

We further say that (Comp, Φ) is (τ, τ') -norm preserving, if for all $c \in \mathcal{C}, x \in M, z \in M'$:

$$\|\text{Comp}_c(x)\| \leq \tau \cdot \|x\| \text{ and } \|\Phi_c(z)\| \leq \tau' \cdot \|z\|.$$

The reason why $\Phi_c \circ \text{Comp}_c$ has to be of this specific form is to allow *extractability* even if the maps π_L, π_R are not evaluated honestly. More precisely, let $\Psi: M \rightarrow N$. Then, given pairwise distinct $c_1, c_2, c_3 \in \mathcal{C}$ and $z_1, z_2, z_3 \in M'$ such that $\Psi \circ \Phi_{c_i}(z_i) = A + c_i X + c_i^2 B$ for $i \in [3]$ (for arbitrary $A, B \in N$), it is possible to extract an $x \in M$ with $\Psi(x) = X$ (resulting in 3-special soundness of the compression mechanism). In the lattice setting it is further crucial that we can give a meaningful bound on the norm of the extracted x . In the proof of Theorem 5 we will show that this is indeed the case.

Example 1 (Bulletproof compression function [14,16]). Let $M = \mathcal{R}^n$ and $M' = \mathcal{R}^{n/2}$. Then, the Bulletproof compression function is obtained as

$$\begin{aligned} \text{Comp}_c((x_L, x_R)) &= x_L + c \cdot x_R, \\ \Phi_c(z) &= (cz, z), \end{aligned}$$

and

$$\begin{aligned} \pi_L((x_L, x_R)) &= (0, x_L), \\ \pi_R((x_L, x_R)) &= (x_R, 0). \end{aligned}$$

Recall that $w(\mathcal{C}) = \max_{c \in \mathcal{C}, x \in \mathcal{R} \setminus \{0\}} \|cx\|_\infty / \|x\|_\infty$. The Bulletproof compression function is $(1 + w(\mathcal{C}), w(\mathcal{C}))$ -norm preserving, as for all $c \in \mathcal{C}, x \in M$

$$\begin{aligned} \|x_L + c \cdot x_R\|_\infty &\leq \|x\|_\infty + w(\mathcal{C})\|x\|_\infty, \\ \|(cz, z)\|_\infty &\leq w(\mathcal{C})\|z\|_\infty, \end{aligned}$$

whenever $w(\mathcal{C}) \geq 1$ (which will be the case for our instantiations).

Using the Bulletproof compression function with the p -norm $\|\cdot\|_p$ for arbitrary $p \in \mathbb{N} \cup \{\infty\}$ instead of restricting to the infinity norm, we obtain that the Bulletproof compression function is $(1 + w_p(\mathcal{C}), 1 + w_p(\mathcal{C}))$ -norm preserving, because in general we can only guarantee

$$\|(cz, z)\|_p \leq w_p(\mathcal{C})\|z\|_p + \|z\|_p,$$


where now $w_p(\mathcal{C}) = \max_{c \in \mathcal{C}, x \in \mathcal{R} \setminus \{0\}} \|cx\|_p / \|x\|_p$.

The idea of the compression mechanism is as follows: First the prover commits to $A = \Psi(\pi_L(x))$ and $B = \Psi(\pi_R(x))$. Next, the verifier sends a challenge $c \in \mathcal{C}$. Using the compression mechanism, the prover then compresses x as $z = \text{Comp}_c(x)$. Now, the verifier can check if indeed $\Psi(\Phi_c(z)) = A + cX + c^2B$. As $\text{Comp}_c(x)$ is 2-compressing, this strategy reduces communication complexity by roughly a factor 2. Note that this factor 2 reduction comes at the cost of sending two elements $A, B \in N$. Hence, in practice the reduction of the communication cost depends on the size of the \mathcal{R} -module N . Finally, by extractability it follows that the compression mechanism is 3-special sound.

The compression mechanism is graphically displayed in Protocol 2 and its properties are summarized in Theorem 5. For a formal proof we refer to the full version of this paper [1].

Protocol 2 Generic Compression Mechanism Π_1 for relations $(R(\Psi, \beta), R(\Psi, \beta\sigma, \zeta^3))$, where $\sigma = 6\tau\tau'w_M(\mathcal{C})^2\bar{w}_M(\mathcal{C}, \zeta)^3$. Recall that (Comp, Φ) is a (τ, τ') -norm preserving extractable compression map, i.e. for all $c \in \mathcal{C}$:

$$\Phi_c(\text{Comp}_c(x)) = \pi_L(x) + cx + c^2\pi_R(x).$$



Theorem 5 (Compression Mechanism). *Let M, M', N be \mathcal{R} -modules, such that M has even rank n and M' has rank $n/2$ over \mathcal{R} , and let $\Psi: M \rightarrow N$ be an \mathcal{R} -module homomorphism. Further, let $\zeta \in \mathcal{R}$ and let \mathcal{C} be a finite ζ -exceptional subset of \mathcal{R} , let (Comp, Φ) be a (τ, τ') -norm preserving extractable compression function for \mathcal{C} with projection maps π_L, π_R , and let $\sigma = 6\tau\tau'w_M(\mathcal{C})^2\bar{w}_M(\mathcal{C}, \zeta)^3$. Then, Π_1 as given in Protocol 2 is a 3-move protocol for relations $(R(\Psi, \beta), R(\Psi, \beta\sigma, \zeta^3))$ which satisfies perfect completeness and unconditional 3-special soundness.*

5.3 Compressed Σ -Protocol

In this setting we build on the previous sections in order to present the compressed Σ -Protocol Π_{comp} , allowing to reduce complexity to polylogarithmic in the input length (when choosing a suitable instantiation).

The introduced soundness slack makes concatenating protocols a bit more involved than in the plain setting. For more details and a formal treatment of this issue we refer to the full version of this paper [1]. Informally

$$\Pi_{\text{comp}} = \Pi_1 \diamond \dots \diamond \Pi_1 \diamond \Pi_0,$$

for the appropriate instantiations of Π_0 and Π_1 . Recall, that in the composition $\Pi_b \diamond \Pi_a$, the final message of protocol Π_a is replaced by an execution of Π_b .

Building on the composition theorem and the results of the previous sections, where the compression function is instantiated with the Bulletproof compression function, we obtain the following corollary.

Corollary 2 (Generic Compressed Σ -Protocol). *Let $\mu \in \mathbb{N}$. Let $M = \mathcal{R}^{2\mu}$ and $\|\cdot\|_\infty$ the infinity norm on M (for some underlying norm on \mathcal{R}). Let $\Psi: M \rightarrow N$ be an \mathcal{R} -module homomorphism, let $\zeta \in \mathcal{R}$ and let \mathcal{C} be a finite ζ -exceptional subset of \mathcal{R} . Let $\alpha, \beta \in \mathbb{N}$ and $\delta \in [0, 1)$, let $V = \{cy \mid y \in M, \|y\|_\infty \leq \alpha, c \in \mathcal{C}\}$ and let $(\mathcal{D}, \mathcal{F})$ be a β -bounded V -hiding distribution with abort probability δ . Then, there exists a $(2\mu + 3)$ -move public-coin protocol Π_{comp} for the pair of relations*

$$(R(\Psi, \alpha), R(\Psi, 2\beta \cdot \bar{w}(\mathcal{C}, \zeta) \cdot \sigma^\mu, \zeta^{3\mu+1})),$$

where $\sigma = 6 \cdot w(\mathcal{C})^3 \cdot (1 + w(\mathcal{C})) \cdot \bar{w}(\mathcal{C}, \zeta)^3$.

It is complete with completeness error δ , unconditionally $(2, 3, \dots, 3)$ -special sound and non-abort special honest-verifier zero-knowledge. Moreover, the communication costs are:

- $\mathcal{P} \rightarrow \mathcal{V}$: $2\mu + 1$ elements of N and 1 element of \mathcal{R} .
- $\mathcal{V} \rightarrow \mathcal{P}$: $\mu + 1$ elements of \mathcal{C} .

In the full version of this paper [1], we outline how the abstract Σ -protocol theory yields a proof of knowledge with knowledge error $\kappa \leq (2\mu + 1)/|\mathcal{C}|$, which can be decreased to $1/\lambda^d$ for arbitrary constant $d \in \mathbb{N}$ by applying the parallel repetition theorem (Theorem 3). Moreover, there we discuss the issues that arise when applying the Fiat-Shamir transform to our protocol in order to transform it into a non-interactive PoK. We further give details on how to use our compressed Σ -protocols non-interactively via the Fiat-Shamir transform.

6 Compressed Σ -Protocols from the MSIS Assumption

The compressed Σ -protocol Π_{comp} of Corollary 2 is typically instantiated with $\Psi(\mathbf{x}, \gamma) = (\text{COM}(\mathbf{x}, \gamma), L(\mathbf{x}))$ for a commitment scheme COM and a linear form L , where γ is the commitment randomness. This allows a prover to show that a committed vector \mathbf{x} satisfies a *linear* constraint. When instantiated with a compact or compressing commitment scheme, for which the size of a commitment is at most polylogarithmic in the size of the secret vector, protocol Π_{comp} achieves communication complexity polylogarithmic in the input size. In the full version of this paper [1], we show how to linearize non-linear constraints and thereby prove that committed vectors satisfy arbitrary *non-linear* constraints. Therefore compressed Σ -protocol Π_{comp} is only required to handle linear instances.

The generalizations of Sect. 5 were introduced to handle *lattice-based* commitment schemes. In this section, we instantiate compressed Σ -protocol Π_{comp} for the following lattice-based commitment function (Definition 3)

$$\text{COM}: \mathcal{R}^n \times \mathcal{R}^r \rightarrow \mathcal{R}_q^k, \quad (\mathbf{x}, \gamma) \mapsto A_1\gamma + A_2\mathbf{x} \pmod{q}.$$

Recall that, $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic irreducible polynomial $f(X)$, $\mathcal{R}_q = \mathcal{R}/(q)$ for a rational prime q , and $A_1 \in \mathcal{R}_q^{k \times r}$ and $A_2 \in \mathcal{R}_q^{k \times n}$ are sampled uniformly at random in the setup phase. This commitment scheme allows a prover to commit to “short” ring elements. We use it to commit to secret vectors of

\mathcal{R}_p^n via their unique representation in $\{x \in \mathcal{R} : \|x\|_\infty \leq \lceil (p-1)/2 \rceil\}$. Subsequently, we aim to prove that a committed vector $\mathbf{x} \in \mathcal{R}_p^n$ satisfies an \mathcal{R}_p -linear constraint $L(\mathbf{x}) = y$ for a linear form $L : \mathcal{R}_p^n \rightarrow \mathcal{R}_p$. To this end, we instantiate protocol Π_{comp} with $\alpha = \lceil (p-1)/2 \rceil$ for the \mathcal{R} -module homomorphism

$$\Psi : \mathcal{R}^n \times \mathcal{R}^r \rightarrow \mathcal{R}_q^k \times \mathcal{R}_p, \quad (\mathbf{x}, \gamma) \mapsto (\text{COM}(\mathbf{x}, \gamma), L(\mathbf{x}) \pmod p).$$

Note that the protocol of Corollary 2 contains an approximation factor $\zeta^{3\mu+1}$. This means that, in the instantiation of this section, a prover claims to know an exact opening (\mathbf{x}, γ) of a commitment P satisfying $L(\mathbf{x}) = y$, but is only capable of proving knowledge of a relaxed opening (\mathbf{x}', γ') such that $\text{COM}(\mathbf{x}', \gamma') = \zeta^{3\mu+1} \cdot P$ and $L(\mathbf{x}) = \zeta^{3\mu+1} \cdot y \in \mathcal{R}_p$. For this reason, we require the approximation factor ζ to be invertible in \mathcal{R}_p . In this case, a commitment to a vector $\mathbf{x}' \in \mathcal{R}_p^n$ is also a commitment to the vector $\tilde{\mathbf{x}} = \zeta^{-3\mu-1} \mathbf{x}' \in \mathcal{R}_p^n$ satisfying the linear constraint $L(\tilde{\mathbf{x}}) = y$. Hence, if $\zeta \in \mathcal{R}_p^*$, we derive precisely the desired functionality of proving that a committed vector satisfies a linear constraint.

The lattice instantiation requires a distribution-algorithm pair $(\mathcal{D}, \mathcal{F})$ that is V -hiding, for $V = \{cy \mid y \in M, \|y\|_\infty \leq \alpha, c \in \mathcal{C}\}$, and β -bounded for some reasonably small $\beta \in \mathbb{N}$. We let \mathcal{D} be a uniform distribution over an appropriate subset of \mathcal{R}^{n+r} . The following lemma shows that this approach gives the required properties. The smallest lattice-based signatures take \mathcal{D} to be a Gaussian distribution. Namely, when the secrets have a bounded ℓ_2 -norm, the Gaussian distribution results in better protocol parameters. In our scenario this is not the case; our secrets are bounded in the ℓ_∞ -norm. Additionally, uniform sampling is less prone to side-channel attacks. For this reason, the digital signature scheme Dilithium also deploys a uniform rejection sampling approach [24].

Lemma 6 (Uniform Rejection Sampling). *Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(X) \in \mathbb{Z}[X]$ of degree d , $\mathcal{C} \subset \mathcal{R}$ and $m, \eta \in \mathbb{N}$. Let $\|z\|_\infty$ be the ℓ_∞ -norm of the coefficient vector of $z \in \mathcal{R}^m$ and let $w(\mathcal{C}) = \max_{c \in \mathcal{R}, x \in \mathcal{R} \setminus \{0\}} \|cx\|_\infty / \|x\|_\infty$. Let $V = \{cx \in \mathcal{R}^m : c \in \mathcal{C} \subset \mathcal{R}, \|x\|_\infty \leq \lceil (p-1)/2 \rceil\}$. Let \mathcal{D} be the uniform distribution over $\{x \in \mathcal{R}^m : \|x\|_\infty \leq \eta\}$ and let*

$$\mathcal{F}(r, v) = \begin{cases} \perp, & \text{if } \|v + r\|_\infty > \eta - w(\mathcal{C}) \lceil (p-1)/2 \rceil, \\ v + r, & \text{otherwise.} \end{cases}$$

Then $(\mathcal{D}, \mathcal{F})$ is perfectly V -hiding and $(\eta - w(\mathcal{C}) \lceil (p-1)/2 \rceil)$ -bounded, with abort probability $\delta \leq 1 - e^{-\frac{w(\mathcal{C})pmd}{2\eta+1}}$.

Proof. Note that, for all $v \in V$, it holds that $\|v\|_\infty \leq w(\mathcal{C}) \lceil (p-1)/2 \rceil$. Hence, the abort probability of the probabilistic algorithm $\{\mathcal{F}(r, v) \mid r \leftarrow \mathcal{D}\}$ equals

$$\begin{aligned} \delta &= 1 - \left(1 - \frac{2w(\mathcal{C}) \lceil (p-1)/2 \rceil}{2\eta + 1} \right)^{md}, \\ &\leq 1 - e^{md \log\left(1 - \frac{w(\mathcal{C})p}{2\eta+1}\right)} \leq 1 - e^{-\frac{w(\mathcal{C})pmd}{2\eta+1}}. \end{aligned}$$

Now let \mathcal{F}' be the algorithm that aborts with probability δ and otherwise outputs a $z \in \{x \in \mathcal{R}^m : \|x\|_\infty \leq \eta - w(\mathcal{C}) \lceil (p-1)/2 \rceil\}$ sampled uniformly at random. Then it is easily seen that $\{\mathcal{F}(r, v) \mid r \leftarrow \mathcal{D}\}$ and $\{\mathcal{F}'\}$ have exactly the same output distributions, i.e., $(\mathcal{D}, \mathcal{F})$ is V -hiding.

Finally, $(\mathcal{D}, \mathcal{F})$ is clearly $(\eta - w(\mathcal{C}) \lceil (p-1)/2 \rceil)$ -bounded. □

The resulting instantiation of Π_{comp} , denoted by $\Lambda_{\text{comp}}(\eta)$, is parameterized by $\eta \in \mathbb{N}$ allowing for a trade-off between the abort probability and communication complexity of the protocol. Its properties are summarized in Corollary 3.

Corollary 3 (Lattice-Based Compressed Σ -Protocol). *Let $n, r, \mu, \eta \in \mathbb{N}$ such that $n+r = 2^\mu$ and let $p, q \in \mathbb{N}$ be primes. Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(X) \in \mathbb{Z}[X]$ of degree d . Let $\zeta \in \mathcal{R}$ such that $\zeta \in \mathcal{R}_p^*$ and let \mathcal{C} be a ζ -exceptional subset of \mathcal{R} . Let $A_1 \in \mathcal{R}_q^{k \times r}$, $A_2 \in \mathcal{R}_q^{k \times n}$ and*

$$\Psi: \mathcal{R}^n \times \mathcal{R}^r \rightarrow \mathcal{R}_q^k \times \mathcal{R}_p, \quad (\mathbf{x}, \gamma) \mapsto (A_1\gamma + A_2\mathbf{x} \pmod q, L(\mathbf{x}) \pmod p).$$

Then, there exists a $(2\mu + 3)$ -move public-coin protocol $\Lambda_{\text{comp}}(\eta)$ for the pair of relations

$$R = \left\{ (P; x) : P = \Psi(x), \|x\|_\infty \leq \lceil (p-1)/2 \rceil \right\},$$

$$R' = \left\{ (P; x) : \zeta^{3\mu+1} \cdot P = \Psi(x), \|x\|_\infty \leq 2\sigma^\mu \bar{w}(\mathcal{C}, \zeta)(\eta - w(\mathcal{C}) \lceil (p-1)/2 \rceil) \right\},$$

where $\sigma = 6 \cdot w(\mathcal{C})^3 \cdot (1 + w(\mathcal{C})) \cdot \bar{w}(\mathcal{C}, \zeta)^3$ with $w(\cdot)$ and $\bar{w}(\cdot)$ defined as in Eq. 3.

It is unconditionally $(2, 3, \dots, 3)$ -special sound, non-abort special honest-verifier zero-knowledge and complete with completeness error

$$\delta \leq 1 - e^{-\frac{w(\mathcal{C})p(n+r)d}{2\eta+1}}.$$

Moreover, the communication costs are:

- $\mathcal{P} \rightarrow \mathcal{V}$: $2\mu + 1$ elements of \mathcal{R}_q^k , $2\mu + 1$ elements of \mathcal{R}_p and 1 element of \mathcal{R} .
- $\mathcal{V} \rightarrow \mathcal{P}$: $\mu + 1$ elements of \mathcal{C} .

Remark 4. Corollary 3 does not require ζ to be invertible in \mathcal{R}_p . In particular, this result is still valid for $\zeta = 0$. However, in this case 0 is a witness for all statements $P \in L_{R'}$ and thereby the claim that is being proven becomes vacuous. For this reason, in most practical scenarios we assume that $\zeta \in \mathcal{R}_p^*$.

6.1 Parameters

In this section, we consider compressed Σ -protocol $\Lambda_{\text{comp}}(\eta)$ defined over the cyclotomic number ring $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ with d a power of two and with challenge set $\mathcal{C} = \{0, \pm 1, \pm X, \dots, \pm X^{d-1}\}$. We show that this protocol has communication complexity *polylogarithmic* in the input size. We only consider the simplified scenario of proving knowledge of a commitment opening.

Power-of-two cyclotomic number rings \mathcal{R} and their monomial challenge set \mathcal{C} have certain convenient properties. In particular, $w(\mathcal{C}) = 1$ and \mathcal{C} is a 2-exceptional subset of \mathcal{R} . More precisely, $2/(c - c') \in \mathcal{R}$ is a polynomial with coefficients in $\{-1, 0, 1\}$ for all distinct $c, c' \in \mathcal{C}$ [13]. From this it follows that $\bar{w}(\mathcal{C}, 2) \leq d$. For a more detailed discussion on optimal challenge sets see [8, 37].

Let us now determine the asymptotic communication complexity. First note that, by Theorem 1, $\Lambda_{\text{comp}}(\eta)$ has knowledge error $\kappa \leq (2 \log(n + r) + 1)/(2d + 1) \leq \log(n + r)/d$ (assuming that $\log(n + r) < d$). For this reason $t = \Theta(\lambda/(\log d - \log \log(n + r)))$ parallel repetitions are required, where λ is the security parameter. Note that, in the analysis of the lattice-based Bulletproof folding technique it is incorrectly claimed that their protocol achieves $\mathcal{O}(1/d)$ knowledge error [15, p. 20].² However, similar to our protocol, it achieves a $\mathcal{O}(\log(n + r)/d)$ knowledge error.

Moreover, we assume $\eta = \Theta(tdp(n + r))$, which by Corollary 3 is enough to achieve a constant completeness error. From Corollary 3 it now follows that the extractor outputs a $(B, 2^{3\mu+1})$ -relaxed commitment opening, where

$$B = 2d \cdot (12d^3)^\mu \left(\eta - \left\lfloor \frac{p-1}{2} \right\rfloor \right) = \Theta(d^2 t p(n + r)^{3+\log 3+3 \log d}).$$

Hence, the commitment scheme must be instantiated to be binding with respect to $(B, 2^{3\mu+1})$ -relaxed commitment openings, i.e., the $\text{MSIS}_{k,n+r,2B}^\infty$ problem over \mathcal{R}_q must be computationally infeasible (Lemma 3). Recall that commitments are vectors in \mathcal{R}_q^k . From the Micciancio-Regev bound (Eq. 1) it follows that this problem is hard if

$$dk \log q \geq \frac{\log^2(2B\sqrt{n+r})}{4 \log \delta} = \Theta \left(\frac{\log^2 d \log^2 tdp(n+r)}{\log \delta} \right), \tag{4}$$

where δ is the root Hermite factor. Note that we derive an additional $\sqrt{n+r}$ factor because we reduce the MSIS-problem from the ℓ_∞ -norm to the ℓ_2 -norm. When these commitments are considered stand-alone their size is independent of the input rank n , i.e., they are compact. However, the soundness slack of our protocols depends (polynomially) on n . Hence, the commitment scheme must be instantiated such that the bit size $dk \log q$ of commitments is polylogarithmic.

By Lemma 1 it now follows that r is polylogarithmic in the input size. Together with Corollary 3 and the fact that $t = \Theta(\lambda/(\log d - \log \log(n + r)))$, this shows that the prover has to send

$$\mathcal{O} \left(\frac{\lambda \log^2 d \log n \log^2 \lambda d p n}{\log \delta (\log d - \log \log n)} \right)$$

bits of information to the verifier. Hence, this instantiation of $\Lambda_{\text{comp}}(\alpha, \eta)$ indeed achieves communication complexity polylogarithmic in the input size.

² This was confirmed to us by the authors in personal communication and also observed in [3].

Remark 5. The lattice based Bulletproof instantiation of [15] considers the case $k = 1$ and they derive a communication complexity of $\mathcal{O}(d\lambda \log n \log pn / \log \delta)$ (using our notation) under the assumption that $\log q = \Theta(\log d \log pn)$. However, to ensure that the underlying commitment scheme is binding they must choose $d = \Theta(\log q)$. Moreover, they incorrectly estimate their knowledge error to be $\mathcal{O}(1/d)$ instead of $\mathcal{O}(\log n/d)$. Taking these two issues into account gives their protocol a communication complexity of

$$\mathcal{O}\left(\frac{\lambda \log^2 d \log n \log^2 pn}{\log \delta (\log d - \log \log n)}\right).$$

The additional factor λd inside the logarithm of our communication complexity can be explained by the fact that, in contrast to [15], our protocol is zero-knowledge. Besides this factor, our communication complexity is the same.

Remark 6. Because the security loss of the Fiat-Shamir transform is exponential in the number of rounds, the non-interactive variant of the t -fold parallel repetition of protocol $\Lambda_{\text{comp}}(\eta)$ requires a factor $\mathcal{O}(\mu) = \mathcal{O}(\log n)$ more parallel repetitions than the interactive variant. Therefore, the communication complexity of the non-interactive variant is a factor $\mathcal{O}(\log n)$ larger. This issue has been overlooked in prior works.

7 Proving Non-linear Relations

Thus far, we have shown how to prove that committed vectors satisfy *linear* constraints. To handle non-linear constraints, we deploy an adaptation of the strategy from [6] that uses secret sharing to linearize non-linearities.

The techniques from [6] are not directly applicable to the lattice setting, since their relations and arithmetic secret sharing are defined over a large field. In our adaptation the arithmetic secret sharing is not defined over a field but over a quotient of a number ring. This introduces two challenges: (1) the ring may be small and (2) not all ring elements have a multiplicative inverse. In our adaptation, these challenges are handled by defining the secret sharing scheme over an appropriately chosen ring extension. For more details we refer to the full version of this paper [1].

Acknowledgements. We thank Jelle Don, Serge Fehr, Michael Kloof, Vadim Lyubashevsky and Gregor Seiler for the helpful and insightful discussions. Furthermore, we thank Andrej Bogdanov for pointing out an oversight regarding the composition theorem in the first version of this work.

Thomas Attema has been supported by EU H2020 project No 780701 (PROMETHEUS). Ronald Cramer has been supported by ERC ADG project No 74079 (ALGSTRONGCRYPTO) and by the NWO Gravitation project QSC. Lisa Kohl has been supported by the NWO Gravitation project QSC.

References

1. Full version of this paper. IACR ePrint 2021/307
2. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: STOC, pp. 99–108. ACM (1996)
3. Albrecht, M.R., Lai, R.W.: Subtractive sets over cyclotomic rings: limits of Schnorr-like arguments over lattices. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12826, pp. 519–548 (2021, to appear)
4. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. *J. Math. Cryptol.* **9**, 169–203 (2015)
5. Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Ligerio: lightweight sub-linear arguments without a trusted setup. In: CCS, pp. 2087–2104 (2017)
6. Attema, T., Cramer, R.: Compressed Σ -protocol theory and practical application to plug & play secure algorithmics. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12172, pp. 513–543. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_18
7. Attema, T., Cramer, R., Fehr, S.: Compressing proofs of k-out-of-n partial knowledge. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12828, pp. 65–91 (2021, to appear)
8. Attema, T., Cramer, R., Xing, C.: A note on short invertible ring elements and applications to cyclotomic and trinomials number fields. *Math. Cryptol.* **1**, 45–70 (2021)
9. Baum, C., Bootle, J., Cerulli, A., del Pino, R., Groth, J., Lyubashevsky, V.: Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10992, pp. 669–699. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_23
10. Bellare, M., Impagliazzo, R., Naor, M.: Does parallel repetition lower the error in computationally sound protocols? In: FOCS, pp. 374–383 (1997)
11. Ben-Sasson, E., et al.: Computational integrity with a public random string from quasi-linear PCPs. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10212, pp. 551–579. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56617-7_19
12. Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., Ward, N.P.: Aurora: transparent succinct arguments for R1CS. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11476, pp. 103–128. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17653-2_4
13. Benhamouda, F., Camenisch, J., Krenn, S., Lyubashevsky, V., Neven, G.: Better zero-knowledge proofs for lattice encryption and their application to group signatures. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8873, pp. 551–572. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45611-8_29
14. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 327–357. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_12
15. Bootle, J., Lyubashevsky, V., Nguyen, N.K., Seiler, G.: A non-PCP approach to succinct quantum-safe zero-knowledge. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12171, pp. 441–469. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56880-1_16

16. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: IEEE S&P, pp. 315–334 (2018)
17. Chen, M.-S., Hülsing, A., Rijneveld, J., Samardjiska, S., Schwabe, P.: From 5-Pass \mathcal{MQ} -based identification to \mathcal{MQ} -based signatures. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10032, pp. 135–165. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_5
18. Chung, K.-M., Liu, F.-H.: Parallel repetition theorems for interactive arguments. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 19–36. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11799-2_2
19. Chung, K.-M., Pass, R.: Tight parallel repetition theorems for public-coin arguments using KL-divergence. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9015, pp. 229–246. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46497-7_9
20. Cramer, R.: Modular design of secure yet practical cryptographic protocols. Ph.D. thesis, CWI and University of Amsterdam (1996)
21. del Pino, R., Lyubashevsky, V., Seiler, G.: Short discrete log proofs for FHE and ring-LWE ciphertexts. In: Lin, D., Sako, K. (eds.) PKC 2019. LNCS, vol. 11442, pp. 344–373. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17253-4_12
22. Dodis, Y., Jain, A., Moran, T., Wichs, D.: Counterexamples to hardness amplification beyond negligible. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 476–493. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28914-9_27
23. Don, J., Fehr, S., Majenz, C.: The measure-and-reprogram technique 2.0: multi-round Fiat-Shamir and more. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12172, pp. 602–631. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_21
24. Ducas, L., et al.: CRYSTALS-Dilithium: A lattice-based digital signature scheme. TCHES pp. 238–268 (2018)
25. Esgin, M.F., Steinfeld, R., Sakzad, A., Liu, J.K., Liu, D.: Short lattice-based one-out-of-many proofs and applications to ring signatures. In: Deng, R.H., Gauthier-Umaña, V., Ochoa, M., Yung, M. (eds.) ACNS 2019. LNCS, vol. 11464, pp. 67–88. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-21568-2_4
26. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_12
27. Goldreich, O.: Foundations of Cryptography: Basic Tools, vol. 1. Cambridge University Press, Cambridge (2001)
28. Hästad, J., Pass, R., Wikström, D., Pietrzak, K.: An efficient parallel repetition theorem. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 1–18. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11799-2_1
29. Hazay, C., Lindell, Y.: Efficient Secure Two-Party Protocols - Techniques and Constructions. Springer, Information Security and Cryptography (2010)
30. Hoffmann, M., Kloof, M., Rupp, A.: Efficient zero-knowledge arguments in the discrete log setting, revisited. In: CCS, pp. 2093–2110 (2019)
31. Jaeger, J., Tessaro, S.: Expected-time cryptography: generic techniques and applications to concrete soundness. In: Pass, R., Pietrzak, K. (eds.) TCC 2020. LNCS, vol. 12552, pp. 414–443. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64381-2_15
32. Lang, S.: Algebra, Graduate Texts in Mathematics, vol. 211, 3rd edn. Springer, New York (2002). <https://doi.org/10.1007/978-1-4613-0041-0>

33. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.* **75**, 565–599 (2015)
34. Lindell, Y.: Parallel coin-tossing and constant-round secure two-party computation. *J. Cryptol.* **16**, 143–184 (2003). <https://doi.org/10.1007/s00145-002-0143-7>
35. Lyubashevsky, V.: Fiat-Shamir with aborts: applications to lattice and factoring-based signatures. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_35
36. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_43
37. Lyubashevsky, V., Seiler, G.: Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10820, pp. 204–224. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_8
38. Micciancio, D., Regev, O.: Lattice-based cryptography. In: Bernstein, D.J., Buchmann, J., Dahmen, E. (eds.) Post-Quantum Cryptography, pp. 147–191. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-540-88702-7_5
39. Pietrzak, K., Wikström, D.: Parallel repetition of computationally sound protocols revisited. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 86–102. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_5
40. Schnorr, C.P.: Efficient identification and signatures for smart cards. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 239–252. Springer, New York (1990). https://doi.org/10.1007/0-387-34805-0_22
41. Wikström, D.: Special soundness revisited. IACR ePrint 2018/1157