

Lattice-Based Polynomial Commitments: Towards Asymptotic and Concrete Efficiency

Giacomo Fenzi Ngoc Khanh Nguyen
giacomo.fenzi@epfl.ch khanh.nguyen@epfl.ch
EPFL EPFL

Abstract. Polynomial commitments schemes are a powerful tool that enables one party to commit to a polynomial p of degree d , and prove that the committed function evaluates to a certain value z at a specified point u , i.e. $p(u) = z$, without revealing any additional information about the polynomial. Recently, polynomial commitments have been extensively used as a cryptographic building block to transform polynomial interactive oracle proofs (PIOPs) into efficient succinct arguments.

In this paper, we propose a lattice-based polynomial commitment that achieves succinct proof size and verification time in the degree d of the polynomial. Extractability of our scheme holds in the random oracle model under a natural ring version of the BASIS assumption introduced by Wee and Wu (EUROCRYPT 2023). Unlike recent constructions of polynomial commitments by Albrecht et al. (CRYPTO 2022), and by Wee and Wu, we do not require any expensive preprocessing steps, which makes our scheme particularly attractive as an ingredient of a PIOP compiler for succinct arguments. We further instantiate our polynomial commitment, together with the Marlin PIOP (Eurocrypt 2020), to obtain a publicly-verifiable trusted-setup succinct argument for Rank-1 Constraint System (R1CS). Performance-wise, we achieve 26MB proof size for 2^{20} constraints, which is 10X smaller than currently the only publicly-verifiable lattice-based SNARK proposed by Albrecht et al.

1 Introduction

Due to the significant progress in building quantum computers by various industry leaders, e.g. IBM and Google, there has been a tremendous amount of interest in post-quantum cryptography. This is highly evidenced by the NIST PQC Competition for standardising quantum-safe key encapsulation mechanisms and signatures, where the vast majority of the selected algorithms are based on algebraic lattices. Indeed, not only do the lattice-based constructions offer relatively small key and signature sizes [BDK+18; DKL+18; FHK+20], but they are also renowned for their very fast implementation [LS19; Sei18]. Consequently, lattices seem to be a natural candidate to build more complex quantum-safe primitives, such as non-interactive zero-knowledge proofs (NIZKs).

The last several years have seen enormous progress in constructing practically efficient NIZKs for lattice relations [ALS20; ENS20; LNP22] which can produce proofs of size a few dozen kilobytes. This has led to rather compact and practical constructions of privacy-preserving primitives, such as ring signatures [LN22], blind signatures [AKSY22] and anonymous credentials [JRLS22; BLNS23]. Unfortunately, the aforementioned protocols suffer the following limitations – both the proof size and verification time are linear in the length of the witness. Hence, for proving more complex statements, efficient NIZKs with succinct proof size and verification complexity are desired, i.e. zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs).

Polynomial commitment schemes [KZG10] have been getting more and more spotlight in the SNARKs community. The main reason is that, in combination with Polynomial Interactive Oracle Proofs (PIOPs) [BFS20; CHM+20], this cryptographic primitive can be used to obtain succinct arguments with concrete efficiency (see e.g. [Set20; BCHO22; GLS+21]). In a polynomial commitment

scheme, one can commit to any polynomial $f := \sum_{i=0}^d f_i X^i$ of bounded degree d over a ring R , and then later prove that f evaluated at some public point $u \in R$ is equal to a public image $z \in R$, i.e.

$$f(u) = z . \tag{1}$$

In the context of PIOPs, we require both the proof π and the verification time to be succinct (i.e. polylogarithmic in the degree d), even if the evaluation point is chosen adaptively by a verifier. Further, to obtain a SNARK, we need π to be a proof of *knowledge*; thus we call such a polynomial commitment *extractable*.

Recently, various lattice-based polynomial commitments [ACL+22; WW23; CP22; PPS21; BCFL22] were introduced¹, mainly as a direct application of functional commitments [LRY16] over standard cyclotomic rings $R := \mathbb{Z}_q[X]/(X^N + 1)$ where N is a power-of-two. Indeed, (1) can be seen as a degree-one multivariate polynomial

$$[1 \ u \ u^2 \ \dots \ u^d] \begin{bmatrix} f_0 \\ f_1 \\ \vdots \\ f_d \end{bmatrix} = z . \tag{2}$$

Unfortunately, the aforementioned constructions suffer several limitations when applied in the context of PIOPs. Firstly, succinct verification requires a preprocessing step, meaning that the evaluation point u must be known when public parameters are generated and cannot be chosen adaptively. Further, only [ACL+22; BCFL22] offer extractable polynomial commitments which unfortunately suffer from the following limitations: (i) they rely on a knowledge assumption, making it awkward to set concrete parameters to match a required security parameter, (ii) message space can only consist of short vectors, and (iii) they only support linear functions with short coefficients. This makes proving relations as in (2) cumbersome for large degrees d . Even though one of the issues was circumvented by a promising recent work from Wee and Wu [WW23], which allows committing to vectors of arbitrarily large coefficients, their soundness analysis is left for future work. Therefore, constructing extractable polynomial commitments with succinct verification from lattices still remains an open problem.

1.1 Our Contributions

In this work we propose a lattice-based PIOP-friendly polynomial commitment scheme. Concretely, our construction supports committing to arbitrary polynomials $f \in R[X]$ of bounded degree d over R , and proving evaluations for any point $u \in R$ with no preprocessing necessary. Extractability holds in the random oracle model via the Fiat-Shamir transformation [FS86] under a variant of the BASIS assumption defined recently by Wee and Wu [WW23], which we call **PowerBASIS**.

At the core of our construction lie two split-and-fold interactive protocols for proving polynomial evaluations. The first one, which brings resemblance to lattice Bulletproofs [BLNS20; ACK21; AL21], enjoys proof size and verification complexity polylogarithmic in the degree d . Unfortunately, due to certain restrictions on the challenge space, which are inherited from the aforementioned works, the protocol achieves only $1/\text{poly}(\lambda)$ knowledge soundness error. Even though soundness can be

¹ We excluded generic constructions which simply commit to a polynomial and use a general-purpose SNARK to prove correctness of the evaluation.

scheme	commit time	prover time	verifier time	crs size	commitment size	asymptotic proof size	commitment size	concrete proof size
Construction 1 (Section 5.2)	$O(d^2)$	$O(d)$	$O(\log d)$	$O(d^2)$	$O(1)$	$O(\log d)$	255 KB	349MB
Construction 2 (Section 5.3)	$O(d^2)$	$O(d)$	$d^{O(1/\log \log d)}$	$O(d^2)$	$O(1)$	$d^{O(1/\log \log d)}$	930 KB	6MB

Table 1: Efficiency overview of our polynomial commitment scheme. In this setting, we commit to polynomials of degree at most d over the ring $R := \mathbb{Z}_q[X]/(X^N + 1)$. We count the runtime (resp. sizes) in the number of ring operations (resp. elements), which take time (resp. size) $\text{polylog}(d)$ each. For clarity, we ignore the terms related to the security parameter λ . When computing concrete proof sizes, we set $\lambda = 128$ and $d = 2^{20}$.

amplified via parallel repetition [AF22] for the interactive protocol, this is not necessarily the case in the non-interactive setting when applying the Fiat-Shamir transformation, as discussed in [AFK22]. To this end, we propose the second protocol, which achieves negligible soundness error in one-shot at the cost of *quasi*-polylogarithmic $d^{O(1/\log \log d)}$ proof size and verification runtime. Furthermore, the non-interactive version of the scheme can be proven secure in the random oracle using the framework by Attema et al. [AFK22]. Last but not least, we show how to upgrade the evaluation proof to achieve zero-knowledge using the standard Fiat-Shamir-with-aborts paradigm [Lyu09; Lyu12; BTT22]. We summarise the efficiency of both schemes in Table 1.

As a direct application, we combine our polynomial commitment scheme, which includes batch evaluation proofs, with the Marlin Polynomial IOP [CHM+20] to obtain a trusted-setup (zero-knowledge) succinct non-interactive arguments of knowledge for Rank-1 Constraint System (R1CS). Practically, for $\approx 2^{20}$ constraints our construction achieves proofs of size 26MB, which is around 10X smaller than the only concretely instantiated lattice-based proof system with succinct verification by Albrecht et al. [ACL+22]. Moreover, we obtain a square-root improvement over [ACL+22] in terms of the prover runtime. In comparison with other lattice-based arguments which admit linear verification time, our scheme produces comparable proofs to the recent “square-root” protocol by Nguyen and Seiler [NS22] for bigger R1CS instances, such as 2^{30} constraints, but still more than two orders of magnitude larger than the current state-of-the-art by Beullens and Seiler [BS22]. We refer to Table 2 for full comparison and Section 6 for more details on sizes.

1.2 Technical Overview

We provide a brief overview of our techniques. Let λ be a security parameter, q be an odd prime, and N be a power-of-two. Define the polynomial rings $\mathcal{R} := \mathbb{Z}[X]/(X^N + 1)$ and $\mathcal{R}_q := \mathbb{Z}_q[X]/(X^N + 1)$. Let \mathcal{R}_q^\times be the set of invertible elements in \mathcal{R}_q . For a base $\delta \geq 2$ and $n \geq 1$, we define the gadget matrix as $\mathbf{G}_n := [1 \ \delta \ \dots \ \delta^{\tilde{q}}] \otimes \mathbf{I}_n \in \mathcal{R}_q^{n \times n\tilde{q}}$ where $\tilde{q} := \lceil \log_\delta q \rceil + 1$. For simplicity, we omit the subscript n and write $\mathbf{G} := \mathbf{G}_n$ when it is clear from the context. Further, for a fixed matrix $\mathbf{T} \in \mathcal{R}_q^{n \times k}$ and matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, we denote by $\mathbf{S} \leftarrow \mathbf{A}_\sigma^{-1}(\mathbf{T})$ sampling $\mathbf{S} \in \mathcal{R}_q^{m \times k}$ from the discrete Gaussian distribution with Gaussian parameter $\sigma > 0$ conditioned on $\mathbf{AS} = \mathbf{T}$ over \mathcal{R}_q .

1.2.1 BASIS Commitment Scheme

Until lately, lattice-based commitment schemes were split into two disjoint classes: Hashed-Message Commitments [Ajt96] and Unbounded-Message Commitments [BDL+18]. The former one has the

scheme	assumptions	TP	NI	time		size		concrete proof size
				prover	verifier	crs	proof	
[BBC+18]	(M-)SIS, RO	✓	✓	$O(\ell)$	$O(\ell)$	$O(1)$	$O(\sqrt{\ell})$	-
[BLNS20]	(M-)SIS, RO	✓	✓	$O(\ell)$	$O(\ell)$	$O(1)$	$O(\ell^\varepsilon)$	-
Lattice Bulletproofs [BLNS20; AL21; ACK21]	M-SIS	✓	✗	$O(\ell)$	$O(\ell)$	$O(1)$	$O(\log \ell)$	-
[BF22]	(M-)SIS, RO	✓	✓	$O(\ell)$	$O(\ell)$	$O(1)$	$O(\log \ell)$	-
[NS22]	M-SIS, RO	✓	✓	$O(\ell)$	$O(\ell)$	$O(1)$	$O(\sqrt{\ell})$	6MB
Labrador [BS22]	M-SIS, RO	✓	✓	$O(\ell)$	$O(\ell)$	$O(1)$	$O(\log \ell)$	49KB
[ACL+22]	Knowledge k -M-SIS	✗	✓	$O(\ell^4 \log \ell)$	$O(\log \ell)$	$O(\ell^2)$	$O(\log \ell)$	261MB
This Work	PowerBASIS, RO	✗	✓	$O(\ell^2)$	$\ell^{O(1/\log \log \ell)}$	$O(\ell^2)$	$\ell^{O(1/\log \log \ell)}$	26MB

Table 2: Comparison of lattice-based publicly verifiable proof systems for NP relations of size ℓ with sublinear communication complexity. We count the runtime (resp. sizes) in the number of ring operations (resp. elements), which take time (resp. size) $\text{polylog}(d)$ each, and we ignore the terms related polynomially in the security parameter λ . We exclude the preprocessing step from the verifier runtime. Here $0 < \varepsilon < 1$ is a constant. The “TP” column specifies whether the scheme has transparent setup, and “NI” means whether the protocol can be made non-interactive with negligible soundness error. The concrete proof sizes correspond to proving R1CS with $\ell = 2^{20}$ as reported in the respective works.

property that the sizes of commitments are almost independent of the sizes of the committed values, and thus the commitments are *compressing*. This comes at the cost of the restricted message space being only vectors of small norm. On the other hand, the main characteristic of the latter class is the unbounded message space, but the commitment size is linear in the size of the message.

Recently, Wee and Wu [WW23] proposed the first lattice-based commitment scheme which is compressing, and simultaneously supports arbitrarily large messages over \mathcal{R}_q . The downside of the construction is a requirement on having a trusted setup, which was not necessary in prior works, as well as the quadratic committing time in the message length. In the following, we describe the main intuition behind the construction by Wee and Wu. To this end, we recall the BASIS assumption², which lies at the core of the binding property of the commitment.

BASIS assumption. As in the (Module-)SIS problem [LS15], the adversary’s final goal is to find a non-zero vector \mathbf{s} of small norm such that $\mathbf{A}\mathbf{s} = \mathbf{0}$ for a uniformly random matrix $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$. However, in the BASIS setting the adversary is given more information. Namely, let $(\mathbf{B}, \text{aux}) \leftarrow \text{Samp}(\mathbf{A})$ be an efficient algorithm, which given matrix \mathbf{A} as input, outputs another matrix $\mathbf{B} \in \mathcal{R}_q^{n' \times m'}$ along with some auxiliary information aux . Then, in addition to the challenge matrix \mathbf{A} , the adversary is given a tuple $(\mathbf{B}, \text{aux}, \mathbf{T})$, where \mathbf{T} is a trapdoor³ for \mathbf{B} . In particular, \mathbf{T} can be used to efficiently emulate sampling from $\mathbf{B}_\sigma^{-1}(\mathbf{t})$ for any image $\mathbf{t} \in \mathcal{R}_q^{n'}$ under certain conditions on the parameter $\sigma > 0$.

Note that hardness of the BASIS assumption heavily depends on the Samp algorithm. For instance, if $\text{Samp}(\mathbf{A})$ is an identity function and simply outputs $\mathbf{B} := \mathbf{A}$, then using the trapdoor \mathbf{T}

² BASIS stands for Basis-Augmented Shortest Integer Solution.

³ In [WW23], the trapdoor \mathbf{T} is generated by sampling $\mathbf{T} \leftarrow \mathbf{B}_\sigma^{-1}(\mathbf{G})$. Since the matrix $\mathbf{T} \in \mathcal{R}_q^{m' \times n' \tilde{q}}$ is short and $\mathbf{B}\mathbf{T} = \mathbf{G}$, it can be used in Micciancio-Peikert trapdoor sampling [MP12] to efficiently generate preimages under \mathbf{B} .

we can find a short non-zero solution to \mathbf{A} by sampling $\mathbf{s} \leftarrow \mathbf{B}_\sigma^{-1}(\mathbf{0})$. In this paper, we consider the following three instantiations of the `Samp` algorithm:

- **StructBASIS**: The sampling algorithm `Samp`(\mathbf{A}) first generates a row $\mathbf{a}^\top \leftarrow \mathcal{R}_q^\ell$ and sets

$$\mathbf{A}^* := \begin{bmatrix} \mathbf{a}^\top \\ \mathbf{A} \end{bmatrix} \in \mathcal{R}_q^{(n+1) \times \ell} . \quad (3)$$

Next, it samples square matrices $\mathbf{W}_1, \dots, \mathbf{W}_\ell \in \mathcal{R}_q^{(n+1) \times (n+1)}$ and outputs

$$\mathbf{B}_\ell := \begin{bmatrix} \mathbf{W}_1 \mathbf{A}^* & & & -\mathbf{G}_{n+1} \\ & \ddots & & \vdots \\ & & \mathbf{W}_\ell \mathbf{A}^* & -\mathbf{G}_{n+1} \end{bmatrix} \quad \text{and} \quad \text{aux} := (\mathbf{W}_1, \dots, \mathbf{W}_\ell) .$$

- **PowerBASIS**: `Samp`(\mathbf{A}) generates a row $\mathbf{a}^\top \leftarrow \mathcal{R}_q^\ell$ and sets \mathbf{A}^* as in (3). Then, it samples a single square matrix $\mathbf{W} \leftarrow \mathcal{R}_q^{(n+1) \times (n+1)}$ and outputs

$$\mathbf{B}_\ell := \begin{bmatrix} \mathbf{W}^0 \mathbf{A}^* & & & -\mathbf{G}_{n+1} \\ & \ddots & & \vdots \\ & & \mathbf{W}^{\ell-1} \mathbf{A}^* & -\mathbf{G}_{n+1} \end{bmatrix} \quad \text{and} \quad \text{aux} := \mathbf{W} . \quad (4)$$

- **PRISIS**⁴: `Samp`(\mathbf{A}) samples a row $\mathbf{a}^\top \leftarrow \mathcal{R}_q^\ell$ and sets \mathbf{A}^* as in (3). Then, it samples a uniformly random polynomial $w \leftarrow \mathcal{R}_q$ and outputs

$$\mathbf{B}_\ell := \begin{bmatrix} w^0 \mathbf{A}^* & & & -\mathbf{G}_{n+1} \\ & \ddots & & \vdots \\ & & w^{\ell-1} \mathbf{A}^* & -\mathbf{G}_{n+1} \end{bmatrix} \quad \text{and} \quad \text{aux} := w .$$

Observe that the only difference between these variants is how the square matrices $\mathbf{W}_1, \dots, \mathbf{W}_\ell$ are generated. For **StructBASIS** they are picked independently and uniformly at random, while for **PowerBASIS** (resp. **PRISIS**) each matrix \mathbf{W}_i is defined as $\mathbf{W}_i := \mathbf{W}^{i-1}$ for $i \in [\ell]$, where $\mathbf{W} \leftarrow \mathcal{R}_q^{(n+1) \times (n+1)}$ (resp. $\mathbf{W} := w \cdot \mathbf{I}_{n+1}$ for $w \leftarrow \mathcal{R}_q$). Not to mention the fact that the functional commitment from [WW23] can be built on top of all three **BASIS** instantiations⁵.

In this work, we analyse hardness of the three newly introduced assumptions for $\ell = 2$. Concretely, we prove that

$$\text{StructBASIS} \xleftrightarrow{\text{Lemma 3.5}} \text{PowerBASIS} \quad \text{and} \quad \text{PRISIS} \xleftrightarrow{\text{Lemma 3.6}} \text{MSIS} + \text{NTRU} .$$

Unfortunately, the techniques do not translate well for larger values of ℓ , as we argue in Section 3.2. Therefore, hardness of the **BASIS** assumption for $\ell > 2$ is left as an open problem.

⁴ The name stands for Power-Ring-BASIS.

⁵ A reader familiar with the work of [WW23] can notice a difference between **StructBASIS** and the original **BASIS**_{struct} from [WW23, Assumption 3.3]. Namely, the latter one directly sets the matrix $\mathbf{A}^* := \mathbf{A}$ without appending an additional row \mathbf{a}^\top at the top (as in **BASIS**_{rand} [WW23, Assumption 3.3]). Note that it is possible to build a commitment scheme based on such a variant, as described in [WW23, Section 4], but this would increase the commitment, as well the opening sizes, by a factor of $n\tilde{q}$. Hence, for efficiency we consider the modified version of **BASIS**_{struct} as presented here.

Commitment construction. We describe a commitment scheme based on the PowerBASIS assumption. Trivial modifications can be made in order to make the scheme secure under the StructBASIS or PRISIS assumptions.

Consider a message space of arbitrary vectors in \mathcal{R}_q^{d+1} of length $d + 1$. The setup algorithm generates a (pseudo-)random matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, along with a uniformly random invertible matrix $\mathbf{W} \in \mathcal{R}_q^{n \times n}$. Further, it computes a trapdoor \mathbf{T} for the matrix

$$\mathbf{B} := \left[\begin{array}{ccc|c} \mathbf{W}^0 \mathbf{A} & & & -\mathbf{G} \\ & \ddots & & \vdots \\ & & \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right]. \quad (5)$$

Then, the common reference string is $\text{crs} := (\mathbf{A}, \mathbf{W}, \mathbf{T})$.

In order to commit to a vector $\mathbf{f} = (f_0, f_1, \dots, f_d) \in \mathcal{R}_q^{d+1}$, one uses the trapdoor \mathbf{T} to sample short $\mathbf{s}_0, \dots, \mathbf{s}_d \in \mathcal{R}_q^m$ and $\hat{\mathbf{t}} \in \mathcal{R}_q^{n\hat{q}}$ as follows:

$$\begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \leftarrow \mathbf{B}_\sigma^{-1} \left(\begin{bmatrix} -f_0 \mathbf{W}^0 \mathbf{e}_1 \\ -f_1 \mathbf{W}^1 \mathbf{e}_1 \\ \vdots \\ -f_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix} \right)$$

where $\mathbf{e}_1 := (1, 0, \dots, 0)^\top \in \mathcal{R}_q^n$. The commitment becomes $\mathbf{t} := \mathbf{G}\hat{\mathbf{t}}$, and the opening consists of $(\mathbf{s}_i)_{i \in [0, d]}$. The opening algorithm, given the common reference string crs , commitment $\mathbf{t} \in \mathcal{R}_q^n$ and openings $(\mathbf{s}_i)_{i \in [0, d]}$ as input, checks whether for all $i = 0, 1, \dots, d$:

$$\mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \quad \text{and} \quad \|\mathbf{s}_i\| \leq \beta$$

for some norm parameter $\beta > 0$.

Security properties. In this paper, we consider the notion of *relaxed binding* [ALS20]. Namely, we say that a relaxed opening for a commitment \mathbf{t} consists of (i) a vector of openings $\mathbf{s} = (\mathbf{s}_0, \dots, \mathbf{s}_d)$, (ii) a message $\mathbf{f} = (f_0, \dots, f_d) \in \mathcal{R}_q^{d+1}$, and (iii) a vector of relaxation factors $\mathbf{c} := (c_0, \dots, c_d) \in \mathcal{R}_q^{d+1}$, which together satisfy:

$$\mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t}, \quad \|c_i \cdot \mathbf{s}_i\| \leq \beta, \quad \|c_i\|_1 \leq \kappa \quad \text{and} \quad c_i \in \mathcal{R}_q^\times$$

for $i = 0, 1, \dots, d$ and some $\kappa \geq 1$. In particular, vectors \mathbf{s}_i do not need to be short.

Now, we show that the commitment scheme is binding w.r.t. relaxed openings under the PowerBASIS assumption. Indeed, let \mathcal{B} be the following adversary for the PowerBASIS security game, which is given as input a tuple $(\mathbf{A}, \mathbf{B}, \mathbf{W}, \mathbf{T})$ from the challenger, where \mathbf{B} is defined as in (4) for $\ell = d + 1$, and \mathbf{A}^* is constructed as in (3). First, \mathcal{B} aborts if \mathbf{W} is not invertible⁶. Otherwise, \mathcal{B} passes $\text{crs} := (\mathbf{A}^*, \mathbf{W}, \mathbf{T})$ to the adversary \mathcal{A} against the relaxed binding game. Suppose \mathcal{A} comes up with two relaxed openings $(\mathbf{s}, \mathbf{f}, \mathbf{c})$ and $(\mathbf{s}', \mathbf{f}', \mathbf{c}')$ for the same commitment \mathbf{t} and $\mathbf{f} \neq \mathbf{f}'$. Thus, for some index i we have $f_i \neq f'_i$. Then, by definition of relaxed openings we have

$$\mathbf{A}^*(\mathbf{s}_i - \mathbf{s}'_i) + (f_i - f'_i) \mathbf{e}_1 = \mathbf{0}.$$

⁶ Unlike in PowerBASIS, the commitment construction requires that matrix \mathbf{W} is invertible. However, by carefully choosing parameters q and N , one can argue that the probability of $\mathbf{W} \leftarrow \mathcal{R}_q^{n \times n}$ not being invertible is negligible (c.f. [BTT22, Appendix C.3] and [EZS+19, Appendix C]).

Since $f_i - f'_i \neq 0$, we must have $\bar{\mathbf{s}}_i := \mathbf{s}_i - \mathbf{s}'_i \neq 0$. Hence by definition of \mathbf{A}^* , $\bar{\mathbf{s}}_i$ is a non-zero solution for the matrix \mathbf{A} , but not necessarily a short one. To conclude the proof, note that $c_i c'_i \bar{\mathbf{s}}_i$ is still a non-zero vector, due to the invertibility property of c_i, c'_i , and at the same time:

$$\|c_i c'_i \bar{\mathbf{s}}_i\| \leq \|c'_i (c_i \mathbf{s}_i)\| + \|c_i (c'_i \mathbf{s}'_i)\| \leq 2\kappa\beta . \quad (6)$$

Thus, $c_i c'_i \bar{\mathbf{s}}_i$ is a valid solution for the PowerBASIS problem.

Finally, the statistical hiding property is directly inherited from the original construction of the BASIS commitment by Wee and Wu [WW23].

1.2.2 Framework for Proving Polynomial Evaluations

We use the construction above to build our polynomial commitment scheme. Namely, given a polynomial $f \in \mathcal{R}_q[X]$ of degree at most d over \mathcal{R}_q , we commit to f by committing to its coefficient vector $\mathbf{f} = (f_0, f_1, \dots, f_d) \in \mathcal{R}_q^{d+1}$, as described in Section 1.2.1, to obtain a commitment $\mathbf{t} \in \mathcal{R}_q^n$ along with a short opening $(\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_d)$, where each $\mathbf{s}_i \in \mathcal{R}_q^m$.

An essential property of polynomial commitments is being able to prove that the committed polynomial was evaluated correctly, i.e. $f(u) = z$ for public u and z in \mathcal{R}_q . In the setting of our commitment scheme, we are interested in the following ternary relation⁷:

$$\mathbf{R}_{d,\beta} := \left\{ ((\mathbf{A}, \mathbf{W}, \mathbf{T}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_{0 \leq i \leq d})) \mid \begin{array}{l} \forall 0 \leq i \leq d, \mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \wedge \|\mathbf{s}_i\| \leq \beta \\ \wedge f(u) = z \end{array} \right\} . \quad (7)$$

The key ingredient for proving such relations efficiently will be the compressed Σ -protocol in Figure 1, which we will use recursively.

We take inspiration from a common split-and-fold technique used by prior works, e.g. FRI [BBHR19] and DARK [BFS20]. Concretely, take $k \in \mathbb{N}$ and suppose $d+1 = k^h$ for some $h \in \mathbb{N}$. Let us write the polynomial $f(X) = \sum_{i=0}^d f_i X^i$ as

$$f(X) = \sum_{t=1}^k f_t(X^k) X^{t-1}, \quad \text{where } f_t(X) := \sum_{i=0}^{\frac{d+1}{k}-1} f_{ki+t-1} X^i \quad \text{for } t = 1, 2, \dots, k .$$

Then, we want to prove that $f(u) = \sum_{t=1}^k f_t(u^k) u^{t-1} = z$. To this end, we let the prover send these partial evaluations $z_t := f_t(u^k)$ for $t \in [k]$, and the verifier manually checks whether

$$\sum_{t=1}^k z_t u^{t-1} = z . \quad (8)$$

Further, the verifier returns a challenge $\alpha := (\alpha_1, \dots, \alpha_k)$ from a challenge space $\mathcal{C} \subseteq \mathcal{R}_q^k$. We denote $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$. Later we will discuss concrete instantiations for \mathcal{C} .

Now, consider the folded polynomial $g(X) = \sum_{t=1}^k \alpha_t f_t(X)$ which is of degree at most $d' := (d+1)/k - 1 = k^{h-1} - 1$. The crucial observation here is that using the structure of the PowerBASIS commitment⁸ from Section 1.2.1 we get for every $i = 0, 1, \dots, d'$:

$$(\mathbf{W}^k)^{-i} \left(\sum_{t=1}^k \alpha_t \mathbf{W}^{-(t-1)} \right) \mathbf{t} = \sum_{t=1}^k \alpha_t \mathbf{W}^{-(ki+t-1)} \mathbf{t}$$

⁷ We use the standard notation that the first entry corresponds to the common reference string, the second one is the statement, and the last one is the witness. Also, \mathbf{T} is not going to be used by the prover, nor by the verifier.

⁸ We note that a similar result could be obtained using PRISIS.

Σ -Protocol for $\mathcal{R}_{d,\beta}$




Fig. 1: Compressed Σ -protocol for the relation $\mathcal{R}_{d,\beta}$ from (7). Here, $\text{crs} = (\mathbf{A}, \mathbf{W}, \mathbf{T})$ is the common reference string for our polynomial commitment scheme and $d+1 = k^h$. We denote $d' := (d+1)/k - 1$ to be degree of the polynomial g , and $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$.

$$\begin{aligned}
 &= \mathbf{A} \left(\sum_{t=1}^k \alpha_t \mathbf{s}_{ki+t-1} \right) + \left(\sum_{t=1}^k \alpha_t f_{ki+t-1} \right) \mathbf{e}_1 \\
 &= \mathbf{A}\mathbf{z}_i + g_i \mathbf{e}_1
 \end{aligned}$$

where $\mathbf{z}_i := \sum_{t=1}^k \alpha_t \mathbf{s}_{ki+t-1}$ satisfies $\|\mathbf{z}_i\| \leq \beta' := w \beta$. In other words, $(\sum_{t=1}^k \alpha_t \mathbf{W}^{-(t-1)}) \mathbf{t}$, which can be computed by the verifier in time $O(k)$, is a commitment to the polynomial g with the opening $(\mathbf{z}_j)_{j \in [0, d']}$ w.r.t. the new common reference string $\text{crs}' := (\mathbf{A}, \mathbf{W}^k, \mathbf{T})$. Further, by definition of g :

$$g(u^k) = \sum_{t=1}^k \alpha_t f_t(u^k) = \sum_{t=1}^k \alpha_t z_t .$$

Thus, we can conclude that:

$$\left((\mathbf{A}, \mathbf{W}^k, \mathbf{T}), \left(\sum_{t=1}^k \alpha_t \mathbf{W}^{-(t-1)} \mathbf{t}, u^k, \sum_{t=1}^k \alpha_t z_t \right), (g, (\mathbf{z}_i)_{i \in [0, d']}) \right) \in \mathcal{R}_{d', w \beta} . \quad (9)$$

In our Σ -protocol, the prover directly outputs $(g, (\mathbf{z}_i)_{j \in [0, d']})$ to the verifier, who checks Equations (8) and (9). To achieve succinct proofs and verification, we let the prover recursively run the Σ -protocol on the new instance tuple (9) until the degree of the folded polynomial is zero⁹. Overall, the protocol has $2h + 1$ rounds and the last prover message is a pair of the form $(g, \mathbf{z}) \in \mathcal{R}_q \times \mathcal{R}_q^m$, where $\|\mathbf{z}\| \leq \beta' := w^h \beta$. Performance-wise (excluding the $\text{poly}(\lambda)$ factors), the prover sends $O(hk)$ elements in \mathcal{R}_q , while the verifier makes in total $O(hk)$ operations in \mathcal{R}_q .

We now focus on knowledge soundness. As common in the lattice setting, we aim to extract a witness with respect to the relaxed relation:

$$\tilde{\mathbf{R}}_{d, \beta, \kappa} := \left\{ \left((\mathbf{A}, \mathbf{W}, \mathbf{T}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_{0 \leq i \leq d}, (c_i)_{0 \leq i \leq d}) \right) \left| \begin{array}{l} \forall 0 \leq i \leq d, \mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \\ \wedge \|c_i \cdot \mathbf{s}_i\| \leq \beta \wedge \|c_i\|_1 \leq \kappa \\ \wedge c_i \in \mathcal{R}_q^\times \wedge f(u) = z \end{array} \right. \right\} .$$

In other words, the witness is now a *relaxed opening* for the commitment \mathbf{t} . Note that the relation is still meaningful as long as the commitment scheme is binding w.r.t. relaxed openings.

The knowledge extraction strategy for $\tilde{\mathbf{R}}_{\beta, \kappa}$ will strongly depend on the instantiation of the challenge space \mathcal{C} . In this work, we consider two variants described below.

Construction 1: Monomial protocol. As the name suggests, we will make use of certain invertibility properties of the set of signed monomials in \mathcal{R}_q , following the approach from lattice Bulletproofs [BLNS20; ACK21; AL21]. Namely, we set $(k, h) = (2, \log(d + 1))$ and define the challenge space

$$\mathcal{C} := \{(1, X^i) : i \in \mathbb{Z}\} \subseteq \mathcal{R}_q^k .$$

By construction, $w = 2$ and $|\mathcal{C}| = 2N$. Now, we show that for the challenge space \mathcal{C} above, the Σ -protocol in Figure 1 is special-sound w.r.t. the relaxed relation $\tilde{\mathbf{R}}$. The methodology can then be extended to show that our recursive protocol is $(2, \dots, 2)$ -special sound. Thus, the general parallel repetition results [AF22], as well as security of the Fiat-Shamir transformation in the random oracle model [AFK22] would directly apply here.

To this end, suppose we are given two transcripts

$$\text{tr}_j := ((z_1, z_2), (1, \alpha_j), (g_j, (\mathbf{z}_{j,i})_{i \in [0, d']})) \quad \text{for } j = 0, 1$$

with the same first message (z_1, z_2) and two distinct challenges $(1, \alpha_0) \neq (1, \alpha_1)$ in \mathcal{C} such that

$$\begin{cases} ((\mathbf{A}, \mathbf{W}^2, \mathbf{T}), ((\mathbf{I}_n + \alpha_j \mathbf{W}^{-1}) \mathbf{t}, u^2, z_1 + \alpha_j z_2), (g_j, (\mathbf{z}_{j,i})_{i \in [0, d']})) \in \mathbf{R}_{d', \beta'} \\ z_1 + u z_2 = z \end{cases}$$

where $\beta' := w \beta = 2\beta$. Observing that $\alpha_0 - \alpha_1 \in \mathcal{R}_q^\times$, we define for $i = 0, 1, \dots, d' := (d - 1)/2$

$$\bar{f}_{2i+1} := \frac{g_{0,i} - g_{1,i}}{\alpha_0 - \alpha_1}, \quad \bar{f}_{2i} := \frac{\alpha_1 g_{0,i} - \alpha_0 g_{1,i}}{\alpha_1 - \alpha_0} \quad (10)$$

and similarly

$$\bar{\mathbf{s}}_{2i+1} := \frac{\mathbf{z}_{0,i} - \mathbf{z}_{1,i}}{\alpha_0 - \alpha_1}, \quad \bar{\mathbf{s}}_{2i} := \frac{\alpha_1 \mathbf{z}_{0,i} - \alpha_0 \mathbf{z}_{1,i}}{\alpha_1 - \alpha_0} .$$

⁹ For concrete efficiency, it might be more beneficial to apply the protocol recursively until the degree of the folded polynomial is *sufficiently* small, instead of going down to zero.

Denote $\mathbf{2} := (2, \dots, 2) \in \mathcal{R}_q^{d+1}$. We claim that

$$\left((\mathbf{A}, \mathbf{W}, \mathbf{T}), (\mathbf{t}, u, z), (\bar{f}, (\bar{s}_i)_{i \in [0, d]}, \mathbf{2}) \right) \in \tilde{\mathcal{R}}_{d, 2N\beta', 2}.$$

Let us start with proving correctness of the relaxed opening. By careful inspection:

$$\begin{aligned} \mathbf{A}\bar{s}_{2i+1} + \bar{f}_{2i+1}\mathbf{e}_1 &= \frac{1}{\alpha_0 - \alpha_1} \left((\mathbf{A}\mathbf{z}_{0,i} + g_{0,i}\mathbf{e}_1) - (\mathbf{A}\mathbf{z}_{1,i} + g_{1,i}\mathbf{e}_1) \right) \\ &= \frac{\mathbf{W}^{-2i}}{\alpha_0 - \alpha_1} \left((\mathbf{I}_n + \alpha_0\mathbf{W}^{-1})\mathbf{t} - (\mathbf{I}_n + \alpha_1\mathbf{W}^{-1})\mathbf{t} \right) \\ &= \mathbf{W}^{-(2i+1)}\mathbf{t} \end{aligned}$$

and similarly $\mathbf{A}\bar{s}_{2i} + \bar{f}_{2i}\mathbf{e}_1 = \mathbf{W}^{-2i}\mathbf{t}$. As for shortness, we use the result from [BCK+14] which says that $\left\| \frac{2}{\alpha_0 - \alpha_1} \right\|_\infty = 1$ for any distinct $\alpha_0, \alpha_1 \in \{X^i : i \in \mathbb{Z}\}$. Thus, for any $i \in [0, d]$ we have

$$\|2 \cdot \bar{s}_{2i+1}\| \leq \left\| \frac{2}{\alpha_0 - \alpha_1} \cdot (\mathbf{z}_{0,i} - \mathbf{z}_{1,i}) \right\| \leq \left\| \frac{2}{\alpha_0 - \alpha_1} \right\|_1 \cdot \|\mathbf{z}_{0,i} - \mathbf{z}_{1,i}\| \leq 2N\beta'$$

and similarly

$$\|2 \cdot \bar{s}_{2i}\| \leq \left\| \frac{2}{\alpha_1 - \alpha_0} \cdot (\alpha_1\mathbf{z}_{0,i} - \alpha_0\mathbf{z}_{1,i}) \right\| \leq \left\| \frac{2}{\alpha_1 - \alpha_0} \right\|_1 \cdot \|\alpha_1\mathbf{z}_{0,i} - \alpha_0\mathbf{z}_{1,i}\| \leq 2N\beta'.$$

Finally, we need to prove that the extracted polynomial \bar{f} satisfies $\bar{f}(u) = z$. From the verification equations we know that $g_0(u^2) = z_1 + \alpha_0 z_2$ and $g_1(u^2) = z_1 + \alpha_1 z_2$. Hence,

$$\begin{aligned} \bar{f}(u) &= \sum_{i=0}^{d'} \bar{f}_{2i} u^{2i} + \sum_{i=0}^{d'} \bar{f}_{2i+1} u^{2i+1} \\ &= \sum_{i=0}^{d'} \frac{\alpha_1 g_{0,i} - \alpha_0 g_{1,i}}{\alpha_1 - \alpha_0} \cdot u^{2i} + \sum_{i=0}^{d'} \frac{g_{0,i} - g_{1,i}}{\alpha_0 - \alpha_1} \cdot u^{2i+1} \\ &= \frac{\alpha_1 g_0(u^2) - \alpha_0 g_1(u^2)}{\alpha_1 - \alpha_0} + \frac{g_0(u^2) - g_1(u^2)}{\alpha_0 - \alpha_1} \cdot u \\ &= z_1 + u z_2 \\ &= z \end{aligned}$$

which concludes the proof of the claim.

An almost identical strategy can be applied to our recursive protocol when given a general $(2, \dots, 2)$ -tree of transcripts [ACK21]. In this case, we can extract a relaxed opening $(\bar{f}, (\bar{s}_i)_{i \in [0, d]}, \mathbf{2}^h)$ to the commitment \mathbf{t} which satisfies

$$\left((\mathbf{A}, \mathbf{W}, \mathbf{T}), (\mathbf{t}, u, z), \left(\bar{f}, (\bar{s}_i)_{i \in [0, d]}, \mathbf{2}^h \right) \right) \in \tilde{\mathcal{R}}_{d, (2N)^h \beta', 2^h}$$

where $\beta' := 2^h \beta$ and $\mathbf{2}^h := (2^h, \dots, 2^h)$. In terms of performance, the communication complexity and the verifier runtime (in terms of operations in \mathcal{R}_q) are $O(\log d)$.

Using the knowledge soundness result from [ACK21], we deduce that the soundness error for our protocol is $h/|\mathcal{C}| = h/(2N)$. Since $N = \text{poly}(\lambda)$, we only manage to obtain an inverse-polynomial soundness error. Even though this can be further reduced via parallel repetition in the interactive case [AF22], such amplification does not combine with the Fiat-Shamir transformation [AFK22]. Our second construction circumvents this issue by achieving negligible soundness error in one-shot.




Fig. 2: Visualisation of the notion of coordinate-wise special soundness (CWSS) for $k = 4$ coordinates. Here, $\alpha_i^* \neq \alpha_i$ for all $i \in [4]$.

Construction 2: Large sampling set protocol. In this scenario, we define the challenge space as

$$\mathcal{C} := \{(\alpha_1, \dots, \alpha_k) : \forall i \in [k], \|\alpha_i\|_\infty \leq \beta_{\mathcal{C}}\}$$

for some suitable parameter $\beta_{\mathcal{C}} \geq 1$. Hence, by construction $w \leq k\beta_{\mathcal{C}}N$.

One could naively adapt the strategy from Construction 1 to prove knowledge soundness of the Σ -protocol as follows. To begin with, we aim to extract k accepting transcripts with k pairwise distinct challenges $\boldsymbol{\alpha}_j \in \mathcal{C}$ for $j = 1, \dots, k$. Further, we compute the extracted polynomial f by inverting the $k \times k$ matrix \mathbf{C} , where the j -th row corresponds to the challenge $\boldsymbol{\alpha}_j$ in the j -th transcript. Unfortunately, this approach contains a few critical issues. Firstly, it is unclear whether the matrix \mathbf{C} is invertible. But even if it is, the resulting polynomial f may contain large coefficients, or in the context of relaxed openings, there might be no sufficiently short element $v \in \mathcal{R}_q$ such that $v \cdot f_i$ is short for all coefficients f_i .

We propose an alternative approach which relies on a notion, called *coordinate-wise special soundness*¹⁰ (CWSS). As in special soundness, it says that given $k + 1$ valid transcripts $\text{tr}_j = (\mathbf{a}_j, \boldsymbol{\alpha}_j, \mathbf{z}_j)$ for $j = 0, 1, \dots, d$, such that $\boldsymbol{\alpha}_0, \dots, \boldsymbol{\alpha}_k \in \mathcal{C}$ satisfy a certain relation, then one can extract the witness. The relation is defined as follows: for every $j \in [k]$, vectors $\boldsymbol{\alpha}_0 = (\alpha_{0,1}, \dots, \alpha_{0,k})$ and $\boldsymbol{\alpha}_j = (\alpha_{j,1}, \dots, \alpha_{j,k})$ differ *exactly* in the j -th coordinate, i.e. $\forall i \in [k] \setminus \{j\}, \alpha_{j,i} = \alpha_{0,i}$ and $\alpha_{j,j} \neq \alpha_{0,j}$ (see Figure 2 for visualisation). We prove that for Σ -protocols, CWSS implies knowledge soundness. Furthermore, the argument can be easily generalised using techniques from [ACK21] to the multi-round setting, and the methodology from [AFK22] to argue knowledge soundness of the Fiat-Shamir transformation.

In the following, we show that our Σ -protocol satisfies CWSS. Suppose we are given $k + 1$ valid transcripts

$$\text{tr}_j := ((z_1, \dots, z_k), \boldsymbol{\alpha}_j = (\alpha_{j,1}, \dots, \alpha_{j,k}), (g_j, (\mathbf{z}_{j,i})_{i \in [0,d]})) \quad \text{for } j = 0, 1, \dots, k .$$

Let us fix $j \in [k]$ and consider the transcripts tr_0 and tr_j . From the verification equations we have for $i = 0, \dots, d'$:

$$\begin{aligned}
\mathbf{A}\mathbf{z}_{0,i} + g_{0,i}\mathbf{e}_1 &= \mathbf{W}^{-ki} \left(\sum_{t=1}^k \alpha_{0,t} \mathbf{W}^{-(t-1)} \right) \mathbf{t} \\
\mathbf{A}\mathbf{z}_{j,i} + g_{j,i}\mathbf{e}_1 &= \mathbf{W}^{-ki} \left(\sum_{t=1}^k \alpha_{j,t} \mathbf{W}^{-(t-1)} \right) \mathbf{t}.
\end{aligned}$$

¹⁰ As far as we are aware, this strategy was first introduced by Baum et al. [BBC+18] in the context of amortised lattice-based zero-knowledge proofs.

Since α_0 and α_j are the same in all coordinates apart from the j -th one, by subtracting the two equations we obtain

$$\mathbf{A}(\mathbf{z}_{0,i} - \mathbf{z}_{j,i}) + (g_{0,i} - g_{j,i})\mathbf{e}_1 = (\alpha_{0,j} - \alpha_{j,j})\mathbf{W}^{-(ki+j-1)}\mathbf{t} .$$

Now, by choosing parameters $q, N, \beta_{\mathcal{C}}$ appropriately, and using the result by Lyubashevsky and Seiler that short elements in \mathcal{R}_q are invertible [LS18], we deduce that $\alpha_{0,j} - \alpha_{j,j} \in \mathcal{R}_q^\times$ and thus can define the extracted openings

$$\bar{s}_{ki+j-1} := \frac{\mathbf{z}_{0,i} - \mathbf{z}_{j,i}}{\alpha_{0,j} - \alpha_{j,j}} \quad \text{and} \quad \bar{f}_{ki+j-1} := \frac{g_{0,i} - g_{j,i}}{\alpha_{0,j} - \alpha_{j,j}}$$

and the partial vector of relaxation factors $\mathbf{c}_j := (\alpha_{0,j} - \alpha_{j,j}, \dots, \alpha_{0,j} - \alpha_{j,j}) \in \mathcal{R}_q^{d'+1}$. Then, by construction we have $\mathbf{A}\bar{s}_{ki+j-1} + \bar{f}_{ki+j-1}\mathbf{e}_1 = \mathbf{W}^{-(ki+j-1)}\mathbf{t}$, and further

$$\|(\alpha_{0,j} - \alpha_{j,j}) \cdot \bar{s}_{ki+j-1}\| \leq 2w\beta \quad \text{and} \quad \|\alpha_{0,j} - \alpha_{j,j}\| \leq 2\beta_{\mathcal{C}}N .$$

From the other verification checks we similarly conclude that $\sum_{i=0}^{d'} \bar{f}_{ki+j-1}u^{ki} = z_j$.

Eventually, by running the argument above for $j = 1, 2, \dots, k$, we reconstruct a polynomial $f \in \mathcal{R}_q^{\leq d}[\mathbf{X}]$, along with $(\bar{s}_i)_{i \in [0,d]}$, and the vector $\mathbf{c} := (\mathbf{c}_1, \dots, \mathbf{c}_k)$ of relaxation factors so that

$$((\mathbf{A}, \mathbf{W}, \mathbf{T}), (\mathbf{t}, u, z), (\bar{f}, (\bar{s}_i)_{i \in [0,d]}, \mathbf{c})) \in \tilde{\mathbf{R}}_{d, 2w\beta, 2\beta_{\mathcal{C}}N} .$$

In terms of security, we show that the knowledge soundness error of our Σ -protocol is bounded by $k/(2\beta_{\mathcal{C}} + 1)^N$, where $(2\beta_{\mathcal{C}} + 1)^N$ is the number of all possible choices for a single coordinate in \mathcal{C} . Consequently, by picking $k, \beta_{\mathcal{C}} \geq 1$ and $N = \text{poly}(\lambda)$ appropriately, we achieve negligible soundness error in one-shot.

This strategy can be further applied in our recursive protocol. That is, analogously as for special-soundness, we first generalise the notion of coordinate-wise special soundness in the multi-round setting, and then prove that our protocol satisfies CWSS as above. By following the methodology from [ACK21], we obtain the knowledge soundness error equal to $hk/(2\beta_{\mathcal{C}} + 1)^N$, while the knowledge extractor runs the prover expected $(k + 1)^h$ times, and outputs a relaxed opening $(\bar{f}, (\bar{s}_i)_{i \in [0,d]}, \mathbf{c})$ such that

$$((\mathbf{A}, \mathbf{W}, \mathbf{T}), (\mathbf{t}, u, z), (\bar{f}, (\bar{s}_i)_{i \in [0,d]}, \mathbf{c})) \in \tilde{\mathbf{R}}_{d, \gamma, \xi}$$

where $\gamma := (2^h(2\beta_{\mathcal{C}}N)^{2^h-h-1}w^h) \cdot \beta$ and $\xi := 2\beta_{\mathcal{C}}(2\beta_{\mathcal{C}}N)^{2^h-2}N$. We highlight that the norm blow-up is much larger here than in the monomial case due to certain technical differences¹¹. As a result, we cannot pick $k = 2$ and $h = O(\log d)$ since then one would require $\log q = O(d)$ for relaxed binding to hold (c.f. Equation (6)); thus making the proof size and verifier time polynomial in d . Instead, we instantiate the protocol by choosing $k = O(d^{\frac{1}{\log \log d}})$ and $h = O(\log \log d)$. In this case, $\log q = O(\log^2 d)$ and the proof size and verifier complexity, in terms of operations over \mathcal{R}_q , become $O(d^{\frac{1}{\log \log d}} \log \log d) = d^{O(1/\log \log d)}$.

¹¹ Roughly speaking, in Construction 1 we managed to keep the norm growth smaller due to the fact that the relaxation factors $\mathbf{2}^{\mathbf{h}}$ are independent of the extracted transcripts, which is not the case for the relaxation factors \mathbf{c} in Construction 2. We refer to Section 5.3 for more details.

1.2.3 Polynomial Commitments over Finite Fields

Until now, we were focusing on polynomial commitments over the ring $\mathcal{R}_q := \mathbb{Z}_q[X]/(X^N + 1)$. Here, we sketch how to obtain a polynomial commitment over a *finite field*, which is required by Polynomial IOPs [BFS20; CHM+20] to compile into succinct arguments. The key ingredient, which allows us to do that is the ability to commit to *arbitrarily* large elements in \mathcal{R}_q .

Let $l \geq 1$ be a divisor of N . It is a well-known fact [LS18] that if $q \equiv 2N/l + 1 \pmod{4N/l + 1}$, then there exists a ring isomorphism φ from $\mathbb{F}^{N/l}$ to \mathcal{R}_q , where \mathbb{F} is a finite field of size q^l . Thus, we define a map $\varphi_{\mathbb{F}} : \mathbb{F} \rightarrow \mathcal{R}_q$ as $x \mapsto \varphi(x, 0, \dots, 0)$, and denote the image of $\varphi_{\mathbb{F}}$ as \mathcal{S}_q . We will make use of the fact that \mathcal{S}_q is an ideal of \mathcal{R}_q .

Suppose we want to commit to a polynomial $F \in \mathbb{F}^{\leq d}[X]$ and prove that $F(x) = y$ for $x, y \in \mathbb{F}$. Using the homomorphic property of $\varphi_{\mathbb{F}}$, it is easy to see that this is equivalent to proving $f(u) = z$ over \mathcal{R}_q , where $f[X] := \sum_{i=0}^d \varphi_{\mathbb{F}}(F_i)X^i \in \mathcal{S}_q[X]$, $u = \varphi_{\mathbb{F}}(x) \in \mathcal{S}_q$ and $z = \varphi_{\mathbb{F}}(y) \in \mathcal{S}_q$. Therefore, we commit to the polynomial $f \in \mathcal{R}_q[X]$ and prove evaluation of u at the point z as before.

What we need to take care of is proving that all coefficients of f indeed lie in \mathcal{S}_q . This allows us to extract the polynomial $\bar{F} \in \mathbb{F}[X]$ by taking the inverse of $\varphi_{\mathbb{F}}$ coefficient-wise. Looking at our underlying Σ protocol in Figure 1, the additional proof comes without any change on the prover’s side, while the verifier also checks whether $g \in \mathcal{S}_q[X]$, which is the case since \mathcal{S}_q is an ideal. To see why this modification is sufficient, consider the extraction strategy in Equation (10). Since now $g_{0,i}, g_{1,i} \in \mathcal{S}_q$, we again use the fact that \mathcal{S}_q is an ideal and conclude that $\bar{f}_{2i+1} = (g_{0,i} - g_{1,i})/(\alpha_0 - \alpha_1)$ also lies in \mathcal{S}_q . Identical reasoning follows for both Construction 1 and 2.

1.3 Related Works

The first lattice-based interactive proof with sublinear communication complexity for arithmetic ℓ -gate circuit satisfiability was formally proposed by Baum et al. [BBC+18], where the authors achieve $O(\sqrt{\ell})$ size proofs. The construction was later generalised by Bootle et al. [BLNS20] who define so-called “levelled commitments” and give $O(\ell^{1/k})$ size proofs for proving knowledge of a commitment opening with $k = O(1)$ levels. The main drawback of the scheme is that the modulus for the proof system increases exponentially in k and thus considering more than 2-3 levels seems impractical. Recently, Nguyen and Seiler [NS22] combined the square-root approach from [BBC+18] with the CRT-packing technique from [ENS20] to obtain a practically efficient square-root NIZK, with 6MB proofs for circuits of size $\ell = 2^{20}$.

Bootle et al. [BLNS20] also proposed the first lattice adaptation of the Bulletproofs protocol [BCC+16; BBB+18] over polynomial rings $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^N + 1)$ which offers $\text{polylog}(\ell)$ proof sizes. This approach was later improved independently by Attema et al. [ACK21] and Albrecht and Lai [AL21] in terms of tighter soundness analysis, and also generalised to a more abstract setting by Bootle et al. [BCS21]. While the *split-and-fold* strategy from Bulletproofs is very attractive in the discrete logarithm setting and keeps asymptotic efficiency in the lattice scenario, it does not mix well with the shortness condition required in lattice-based cryptography. Consequently, this leads to a concrete blow-up of the parameters as well as the proof size. Roughly speaking, for the knowledge soundness argument it must be possible to invert the folding in the extraction such that the extracted solution vector is still short. To this end, one needs a challenge space of the underlying compressed Σ -protocol to have a property that (a scaled) inverse of a difference of any two distinct challenges is still short - such sets are called *subtractive*. Hence, Bootle et al. [BLNS20] picked the challenge space to consist of monomial challenges $\mathcal{C} := \{X^i : i \in \mathbb{Z}\} \subseteq \mathcal{R}_q$, which is indeed subtractive as shown in [BCK+14]. Since the Σ -protocol is 3-special-sound, norm of the extracted

solution vector grows by a factor of $O(N^3)$ for *every* level of folding. Then, the parameters must be chosen such that Module-SIS is hard with respect to the norm of the extracted solution vector, resulting in the need for a huge modulus q . Note that a similar issue occurs in our Construction 1 (c.f. Section 5.2). However, since our underlying compressed Σ -protocol is only 2-special-sound, norm of the extracted vector grows by only a factor of $O(N)$ for each folding level (but at the price of having a trusted setup).

In addition to the norm growth of the extracted witness, the restriction on the challenges has a negative impact on the soundness error. Indeed, since the challenge space \mathcal{C} in [BLNS20] has size $2N$, the soundness error becomes only $1/\text{poly}(\lambda)$. Furthermore, it was proven by Albrecht and Lai [AL21] that *all* subtractive set over \mathcal{R}_q have size $O(N)$. This becomes problematic especially in the non-interactive setting due to the result by Attema et al. [AFK22], who showed that the Fiat-Shamir transformation of a parallel repetition of special-sound protocols *does not* necessarily decrease the soundness error. A promising solution to circumvent this limitation was recently proposed by Bünz and Fisch [BF22], who suggested a new knowledge extraction strategy, i.e. the notion of *almost special soundness*, which does not require subtractive sets. Instead, the challenges are picked from the exponential-sized set of integers $[0, 2^\lambda - 1)$. Unfortunately, the former issue with the norm growth for each folding level is still present in [BF22].

Recently, Beullens and Seiler [BS22] showed that by combining a split-and-fold approach with algebraic techniques introduced in linear-sized lattice-based NIZKs [LNP22], it is possible to achieve negligible soundness error whilst controlling the norm growth. This is evidenced with impressive 50KB proofs for circuits of size $\ell = 2^{20}$.

Major downside of all the aforementioned works is a linear verification time, which can be the main efficiency bottleneck when proving satisfiability of large circuits. Until now, the only lattice-based publicly verifiable succinct argument of knowledge with efficient verification (excluding the preprocessing step) was proposed by Albrecht et al. [ACL+22]. The construction is obtained as a direct application of functional commitments [LRY16] and soundness holds under a knowledge assumption. However, similar to our scheme, a trusted setup is required, and more importantly, the prover algorithm runs in time $O(\ell^4 \log \ell)$ which makes it unappealing to implement in practice.

Prior to [ACL+22], all lattice-based zk-SNARKs were in the designated-verifier setting [GMNO18; ISW21; SSEK22]. The constructions use the Linear-PCP compiler [BCI+13] to transform into succinct arguments. Notably, the most recent work by Steinfeld et al. [SSEK22] achieves proofs of size 6KB for $\ell = 2^{20}$ constraints at the cost of very large crs (in the order of tens of gigabytes).

Naturally, there is a line of research focusing on the security of lattice-based zero-knowledge proofs against *quantum adversaries* [DFM20; Kat21; LMS22]. Particularly, Lai et al. [LMS22] show that any multi-round protocol, which satisfies special soundness and *collapsing*, is knowledge sound in the post-quantum setting. As a special case, they demonstrate that the lattice Bulletproofs protocol [BLNS20] is knowledge sound against quantum provers. Since our constructions not only satisfy (coordinate-wise) special soundness but also follow the split-and-fold strategy from [BLNS20], we believe that the general result from [LMS22] can be adapted to our setting.

1.4 Paper Organisation

We start by covering relevant preliminaries in Section 2. This includes the necessary background on lattices, interactive proofs, as well as knowledge extraction strategies and the notion of coordinate-wise special soundness. Section 3 focuses on the general BASIS assumption and its three concrete instantiations: StructBASIS, PowerBASIS and PRISIS. Next, we construct a commitment scheme

based on the PowerBASIS assumption in Section 4. Further, Section 5 shows how to efficiently prove polynomial evaluations, including batching (Section 5.4) and making the protocol zero-knowledge (Section 5.5). In combination with Section 4, this yields a polynomial commitment scheme. Finally, in Section 6 we instantiate our polynomial commitment and propose concrete parameters and sizes.

Acknowledgements. We thank Martin Albrecht and Sasha Lapiha for discussion on the PowerBASIS assumption. Ngoc Khanh Nguyen is supported by the Protocol Labs RFP-013: Cryptonet network grant.

2 Preliminaries

Notation. We denote the security parameter by λ , which is implicitly given to all algorithms unless specified otherwise. Further, we write $\text{negl}(\lambda)$ (resp. $\text{poly}(\lambda)$) to denote an unspecified negligible function (resp. polynomial) in λ . In this work, we implicitly assume that the vast majority of the key parameters, e.g. the ring dimension, and the dimensions of matrices and vectors, are $\text{poly}(\lambda)$. However, the modulus used in this work may be super-polynomial in λ .

For $a, b \in \mathbb{N}$ with $a < b$, write $[a, b] := \{a, a + 1, \dots, b\}$, $[a] := [1, a]$. For $q \in \mathbb{N}$ write \mathbb{Z}_q for the integers modulo q . We denote vectors with lowercase boldface (i.e. \mathbf{u}, \mathbf{v}) and matrices with uppercase boldface (i.e. \mathbf{A}, \mathbf{B}). For a vector \mathbf{x} we write x_i or $\mathbf{x}[i]$ for its i -th entry.

Norms. We define the ℓ_p norm on \mathbb{C}^n as $\|\mathbf{x}\|_p = (\sum_i |x_i|^p)^{1/p}$ for $p < \infty$ and $\|\mathbf{x}\|_\infty := \max_i |x_i|$. Unless otherwise specified, we use $\|\cdot\|$ for the ℓ_2 norm. We let the norm of a matrix be defined as the norm taken over the concatenation of columns of the matrix.

Linear algebra. We let \mathbf{e}_i be the vector with 1 in its i -th entry, 0 everywhere else. For $\mathbf{B} \in \mathbb{R}^{n \times m}$ we let $s_1(\mathbf{B}) = \sup\{\|\mathbf{B}\mathbf{v}\| : \mathbf{v} \in \mathbb{R}^m \wedge \|\mathbf{v}\| = 1\}$ be the **spectral norm** of \mathbf{B} . We also denote by $\tilde{\mathbf{B}}$ the Gram-Schmidt orthonormalization of \mathbf{B} . The Gram-Schmidt norm of \mathbf{B} is defined as

$$\|\tilde{\mathbf{B}}\| := \max_{i \in [m]} \|\tilde{\mathbf{b}}_i\|$$

where $\tilde{\mathbf{b}}_i$ is the i -th column of $\tilde{\mathbf{B}}$.

For a ring R , we define $\text{GL}(n, R)$ to be the group of $n \times n$ invertible matrices over R .

2.1 Lattices

A subset $\Lambda \subseteq \mathbb{R}^m$ is a lattice if the following conditions hold:

- $\mathbf{0} \in \Lambda$, and for $\mathbf{x}, \mathbf{y} \in \Lambda$, $\mathbf{x} + \mathbf{y} \in \Lambda$.
- For every $\mathbf{x} \in \Lambda$, there exists $\epsilon > 0$ such that $\{\mathbf{y} \in \mathbb{R}^m : \|\mathbf{x} - \mathbf{y}\| < \epsilon\} \cup \Lambda = \{\mathbf{x}\}$.

We say $\mathbf{B} \in \mathbb{R}^{m \times k}$ is a basis for Λ if its columns are linearly independent and $\Lambda = \mathcal{L}(\mathbf{B}) := \{\mathbf{B}\mathbf{z} : \mathbf{z} \in \mathbb{Z}^k\}$. If $k = m$ then we say that Λ is full-rank. The span (as a vector space) of the basis of a lattice is the span of a lattice denoted as $\text{Span}(\Lambda)$. We also let Λ^* be the dual lattice defined as $\Lambda^* = \{\mathbf{w} \in \text{Span}(\Lambda) : \langle \Lambda, \mathbf{w} \rangle \subseteq \mathbb{Z}\}$. If $\Lambda \subseteq \mathbb{Z}^m$, we call it an integral lattice. For I an ideal of \mathbb{R}^m , we let $I \cdot \Lambda = \{i \cdot \mathbf{x} : i \in I, \mathbf{x} \in \Lambda\}$, which is also a lattice. For a lattice Λ we denote

$$\lambda_1(\Lambda) := \min_{\mathbf{0} \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\| \quad \text{and} \quad \lambda_1^\infty(\Lambda) := \min_{\mathbf{0} \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\|_\infty .$$

For $\mathbf{t} \in \text{Span}(\Lambda)$, we also define the shifted lattice $\mathbf{t} + \Lambda := \{\mathbf{t} + \mathbf{x} : \mathbf{x} \in \Lambda\}$. We also consider q -ary lattices, namely those with $q\mathbb{Z} \subseteq \Lambda$. For an arbitrary $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ we define the full rank q -ary lattice

$$\begin{aligned}\Lambda^\perp(\mathbf{A}) &= \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{A}\mathbf{z} = \mathbf{0} \pmod{q}\} \\ \Lambda(\mathbf{A}) &= \{\mathbf{z} \in \mathbb{Z}^m : \exists \mathbf{s} \in \mathbb{Z}_q^n, \mathbf{A}\mathbf{z} = \mathbf{s} \pmod{q}\}\end{aligned}$$

For any $\mathbf{u} \in \mathbb{Z}_q^n$ such that there exists \mathbf{x} with $\mathbf{A}\mathbf{x} = \mathbf{u}$, we define $\Lambda_{\mathbf{u}}^\perp(\mathbf{A}) := \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{A}\mathbf{z} = \mathbf{u} \pmod{q}\} = \Lambda^\perp(\mathbf{A}) + \mathbf{x}$.

2.2 Power-of-Two Cyclotomic Rings

Let N be a power-of-two and $\mathcal{K} = \mathbb{Q}[X]/(X^N + 1)$ be the $2N$ -th cyclotomic field. Denote $\mathcal{R} = \mathbb{Z}[X]/(X^N + 1)$ to be the ring of integers of \mathcal{K} . For an odd prime q , we write $\mathcal{R}_q := \mathcal{R}/(q)$. We denote \mathcal{R}_q^\times to be the set of invertible elements in \mathcal{R}_q .

We recall the following inequality, which allows to bound norms on products in the ring \mathcal{R} .

Lemma 2.1. *Let $u, v \in \mathcal{R}$. Then $\|uv\| \leq \|u\|_1 \cdot \|v\|$.*

Proof. Let $u := u_0 + u_1X + \dots + u_{N-1}X^{N-1} \in \mathcal{R}$. Then, by the triangle inequality we get

$$\|uv\| \leq \sum_{i=0}^{N-1} \|u_i v \cdot X^i\| = \sum_{i=0}^{N-1} \|u_i v\| = \sum_{i=0}^{N-1} |u_i| \cdot \|v\| = \|u\|_1 \cdot \|v\| .$$

□

Coefficient embedding. For $x \in \mathcal{K}$, we can consider the additive group isomorphism

$$\begin{aligned}\text{vec} : \mathcal{K} &\rightarrow \mathbb{Q}^N \\ a_0 + a_1X + \dots + a_{N-1}X^{N-1} &\mapsto (a_0, \dots, a_{N-1})^\top\end{aligned}$$

and we refer this as the coefficient embedding of \mathcal{K} . Note that, for $f, g \in \mathcal{K}$, $\langle f, g \rangle = \langle \text{vec}(f), \text{vec}(g) \rangle$ and thus $\|\text{vec}(f)\| = \|f\|$. Furthermore, vec restricts to an isomorphism between $\mathcal{R}_q \cong \mathbb{Z}_q^N$ and $\mathcal{R} \cong \mathbb{Z}^N$. We also extend this to a mapping $\mathcal{K}^m \rightarrow \mathbb{Q}^{mN}$ by applying it component-wise. For $f \in \mathcal{K}$, we let

$$\text{rot}(f) := (\text{vec}(f), \text{vec}(X \cdot f), \dots, \text{vec}(X^{N-1} \cdot f)) \in \mathbb{Q}^{N \times N} ,$$

noting that $\text{rot}(f)\text{vec}(g) := \text{vec}(fg)$ and $\text{rot}(f)\text{rot}(g) = \text{rot}(fg)$. We extend this to matrices $\mathbf{B} \in \mathcal{K}^{m \times n}$ by writing

$$\text{rot}(\mathbf{B}) := \begin{bmatrix} \text{rot}(b_{1,1}) & \dots & \text{rot}(b_{1,n}) \\ \vdots & \ddots & \vdots \\ \text{rot}(b_{m,1}) & \dots & \text{rot}(b_{m,n}) \end{bmatrix} \in \mathbb{Q}^{mN \times nN} .$$

Module lattices. For $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, $\mathbf{x} \in \mathcal{R}_q^m$, $\mathbf{u} = \mathbf{A}\mathbf{x}$, define

$$\begin{aligned}\Lambda^\perp(\mathbf{A}) &:= \{\mathbf{z} \in \mathcal{R}^m : \mathbf{A}\mathbf{z} = \mathbf{0} \pmod{q}\} \\ \Lambda_{\mathbf{u}}^\perp(\mathbf{A}) &:= \{\mathbf{z} \in \mathcal{R}^m : \mathbf{A}\mathbf{z} = \mathbf{u} \pmod{q}\} = \Lambda^\perp(\mathbf{A}) + \mathbf{x} .\end{aligned}$$

Then, $\Lambda^\perp(\mathbf{A}) = \text{vec}^{-1}(\Lambda^\perp(\text{rot}(\mathbf{A})))$ and $\Lambda_{\mathbf{u}}^\perp(\mathbf{A}) = \text{vec}^{-1}(\Lambda_{\text{vec}(\mathbf{u})}^\perp(\text{rot}(\mathbf{A})))$.

Spectral norm. Let $s_1(\mathbf{R}) := \sup\{\|\mathbf{R}\mathbf{v}\| : \mathbf{v} \in \mathcal{K}^w \wedge \|\mathbf{v}\| = 1\}$ be the spectral norm of $\mathbf{R} \in \mathcal{R}^{m \times w}$. Clearly, $s_1(\text{rot}(\mathbf{R})) = s_1(\mathbf{R})$, where the spectral norm of the left-hand side is over \mathbb{R} . Here, we recall a simple bound.

Lemma 2.2. *Let $\mathbf{R} \in \mathcal{R}_q^{m \times t}$. Then $s_1(\mathbf{R}) \leq \sqrt{N} \cdot \|\mathbf{R}\|$.*

Proof. Let $\mathbf{r}_1, \dots, \mathbf{r}_m$ be the rows of \mathbf{R} . Note that by the Cauchy-Schwarz inequality, for any \mathbf{u} with $\|\mathbf{u}\| = 1$ we have that

$$\|\langle \mathbf{r}_i, \mathbf{u} \rangle\|^2 \leq \left(\sum_{j \in [t]} \|r_{i,j} s_j\| \right)^2 \leq N \left(\sum_{j \in [t]} \|r_{i,j}\| \cdot \|s_j\| \right)^2 \leq N \|\mathbf{r}_i\|^2 \cdot \|\mathbf{u}\|^2 \leq N \|\mathbf{r}_i\|^2 .$$

Thus, $\|\mathbf{R}\mathbf{u}\|^2 \leq N \|\mathbf{R}\|^2$ which concludes the proof. \square

Subtractive sets for monomials. We recall the following widely-used result from [BCK+14], which says that the (scaled) inverse of two distinct monomials in \mathcal{R} has coefficients in $\{-1, 0, 1\}$.

Lemma 2.3. *Let $\mathcal{C} := \{X^i : i \in \mathbb{Z}\} \subseteq \mathcal{R}$. Then, for any two distinct $x, y \in \mathcal{C}$, we have $\|\frac{2}{x-y}\|_\infty = 1$.*

Short elements are invertible. For $\kappa > 0$, we define $S_\kappa := \{x \in \mathcal{R}_q : \|x\|_\infty \leq \kappa\}$ to be the set of ring elements in \mathcal{R}_q with infinity norm at most κ . We recall the following invertibility result by Lyubashevsky and Seiler [LS18].

Lemma 2.4. *Let $1 \leq l < N$ be a power-of-two and suppose $q \equiv 2N/l + 1 \pmod{4N/l}$. Then, every non-zero element in S_κ is invertible over \mathcal{R}_q as long as $\kappa < \sqrt{l/N} \cdot q^{l/N}$.*

2.3 Discrete Gaussian Distributions

Let $\sigma > 0$ be a parameter and Λ be a m -dimensional lattice. We then define the discrete Gaussian distribution $\mathcal{D}_{\sigma, \mathbf{c}, \Lambda}$ over a lattice coset $\mathbf{c} + \Lambda$ as follows.

$$\rho_{\sigma, \mathbf{c}}(\mathbf{z}) := \exp\left(-\frac{\pi \|\mathbf{z} - \mathbf{c}\|^2}{\sigma^2}\right) \text{ and } \mathcal{D}_{\sigma, \mathbf{c}, \Lambda}(\mathbf{z}) := \frac{\rho_{\sigma, \mathbf{c}}(\mathbf{z})}{\sum_{\mathbf{x} \in \Lambda} \rho_{\sigma, \mathbf{c}}(\mathbf{x})} .$$

When $\mathbf{c} = \mathbf{0}$ or $\Lambda = \mathbb{Z}^m$, we will omit it from the notation. We naturally extend this notion for lattices over the ring of integers \mathcal{R} , and for matrices by sampling column-wise.

Smoothing parameter. The smoothing parameter $\eta_\epsilon(\Lambda)$ of a lattice is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^*) \leq 1 + \epsilon$. Below we recall the standard upper-bounds on the smoothing parameter [MR07; GPV08].

Lemma 2.5. *Let $\Lambda \subseteq \mathbb{R}^m$ be a lattice, and let $\epsilon > 0$. Then,*

$$\eta_\epsilon(\Lambda) \leq \frac{1}{\lambda_1^\infty(\Lambda^*)} \cdot \sqrt{\frac{\ln(2m(1 + 1/\epsilon))}{\pi}}$$

and in fact, for every basis \mathbf{B} of Λ ,

$$\eta_\epsilon(\Lambda) \leq \|\tilde{\mathbf{B}}\| \cdot \sqrt{\frac{\ln(2m(1 + 1/\epsilon))}{\pi}} .$$

We also recall the bound from [GPV08, Lemma 5.3] and [WW23, Lemma 2.5] for the block-diagonal matrices. Here, we consider the ring setting which can be easily adapted from the aforementioned results.

Lemma 2.6. *Let $\ell, \delta > 1$ and suppose q is prime and $m \geq 2n \log_\delta q$. Then, there exists a negligible function ε such that for all $\mathbf{A}_2, \dots, \mathbf{A}_\ell \in \mathcal{R}_q^{n \times m}$:*

$$\Pr \left[\eta_\varepsilon(\Lambda^\perp(\text{diag}(\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_\ell)) \leq \delta \cdot \log(\ell m N) : \mathbf{A}_1 \leftarrow \mathcal{R}_q^{n \times m} \right] \geq 1 - q^{nN} .$$

Further, we recall the regularity lemma from [LPR13].

Lemma 2.7 (Regularity Lemma). *Let $N = \text{poly}(\lambda)$ and k, n be positive integers such that $\text{poly}(\lambda) \geq m \geq n + \omega(\log \lambda)$. Take $\mathfrak{s} > 2N \cdot q^{n/m+2/(Nm)}$. Then, the following distributions are statistically close:*

$$\left\{ (\mathbf{A}, \mathbf{Ax}) \left| \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m} \\ \mathbf{x} \leftarrow \mathcal{D}_\mathfrak{s}^{mN} \end{array} \right. \right\} \text{ and } \left\{ (\mathbf{A}, \mathbf{u}) \left| \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m} \\ \mathbf{u} \leftarrow \mathcal{R}_q^n \end{array} \right. \right\} .$$

This is slightly modified from the original result in [LPR13, Corollary 7.5] and [BTT22, Lemma 4.2] in a sense that \mathbf{A} might not be full-rank. However, the case $m \geq n + \omega(\log \lambda)$ makes sure the event happens with negligible probability [Ezs+19, Appendix C].

Tail bounds. When sampling over a sufficiently wide discrete Gaussian distribution, a small portion of the probability mass will be in the tail of the distribution, and thus with overwhelming probability the sampled lattice elements will have short norm. The following lemma from [MR07] formalises this intuition.

Lemma 2.8. *For any $0 < \epsilon < 1$, lattice $\Lambda \subseteq \mathbb{R}^m$, center $\mathbf{c} \in \text{Span}(\Lambda)$ and $\sigma > \eta_\epsilon(\Lambda)$,*

$$\Pr \left[\|\mathbf{z}\| \geq \sigma \cdot \sqrt{m} : \mathbf{z} \leftarrow \mathcal{D}_{\sigma, \Lambda, \mathbf{c}} \right] \leq \frac{1 + \epsilon}{1 - \epsilon} 2^{-m} .$$

We also recall the tail bounds for the regular discrete Gaussian distribution over integers [Lyu12].

Lemma 2.9. *Let $\mathbf{z} \leftarrow \mathcal{D}_\mathfrak{s}^m$. Then $\Pr \left[\|\mathbf{z}\| > t \cdot \mathfrak{s} \sqrt{\frac{m}{2\pi}} \right] < \left(t e^{\frac{1-t^2}{2}} \right)^m$.*

By setting $t = \sqrt{2\pi}$, the right-hand side can be upper-bounded by 2^{-2m} .

Preimage sampling for module lattices. Let $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ be a matrix over \mathcal{R}_q and take any $\mathbf{u} \in \mathcal{R}_q^n$. We write $\mathbf{s} \leftarrow \mathbf{A}_\sigma^{-1}(\mathbf{u})$ to denote sampling $\mathbf{s} \leftarrow \mathcal{D}_\sigma^{mN}$ conditioned on $\mathbf{As} = \mathbf{u}$. Assuming there is some $\mathbf{x} \in \mathcal{R}_q^m$ which satisfies $\mathbf{Ax} = \mathbf{u}$, this is the same as sampling $\mathbf{s} \leftarrow \mathcal{D}_{\sigma, \mathbf{x}, \Lambda^\perp(\mathbf{A})}$.

We will need the following lemma from [WW23, Lemma 2.7] for proving hiding property of the commitment scheme.

Lemma 2.10. *Let $n, m, q > 0$. Take any matrices $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, $\mathbf{B} \in \mathcal{R}_q^{n \times \ell}$ where $\ell = \text{poly}(n, \log q)$. Suppose the columns of \mathbf{A} generate \mathcal{R}_q and let $\mathbf{C} := [\mathbf{A} \mid \mathbf{B}]$. Then, for every target vector $\mathbf{t} \in \mathcal{R}_q^n$ and any $\sigma \geq \eta_\epsilon(\Lambda^\perp(\mathbf{A}))$ for some $\epsilon = \text{negl}(\lambda)$, the following distributions are statistically close:*

$$\left\{ \mathbf{v} \mid \mathbf{v} \leftarrow \mathbf{C}_\sigma^{-1}(\mathbf{t}) \right\} \text{ and } \left\{ \begin{array}{l} \mathbf{v}_1 \\ \mathbf{v}_2 \end{array} \left| \mathbf{v}_2 \leftarrow \mathcal{D}_\sigma^{\ell N}, \mathbf{v}_1 \leftarrow \mathbf{A}_\sigma^{-1}(\mathbf{t} - \mathbf{Bv}_2) \right. \right\} .$$

<u>RejSamp:</u>	<u>SimRS:</u>
1: $(\mathbf{u}, \mathbf{v}) \leftarrow h$	1: $(\mathbf{u}, \mathbf{v}) \leftarrow h$
2: $\mathbf{z} \leftarrow \mathcal{D}_{\sigma, \mathbf{v}, \Lambda + \mathbf{u}}^m$	2: $\mathbf{z} \leftarrow \mathcal{D}_{\sigma, \Lambda + \mathbf{u}}^m$
3: return (\mathbf{v}, \mathbf{z}) with prob. $\min\left(\frac{\mathcal{D}_{\sigma}^m(\mathbf{z})}{M \cdot \mathcal{D}_{\sigma, \mathbf{v}}^m(\mathbf{z})}, 1\right)$	3: return (\mathbf{v}, \mathbf{z}) with prob. $\frac{1}{M}$

Fig. 3: Rejection sampling [BTT22].

Rejection sampling. A crucial component in proving the *zero-knowledge* property of lattice-based (non-interactive) arguments is a rejection sampling procedure [Lyu12]. We recall the generalised version introduced recently by Boschini et al. [BTT22] for discrete Gaussian over arbitrary lattices (here we omit the case for ellipsoidal Gaussians).

Lemma 2.11 (Rejection Sampling [BTT22]). *Take any $\alpha, T > 0$ and $\varepsilon \leq 1/2$. Let $\Lambda \subseteq \mathcal{R}^m$ be a lattice over \mathcal{R} and $\sigma \geq \max(\alpha T, \eta_\varepsilon(\Lambda))$ be a parameter. Let $h : \mathcal{R}^m \times \mathcal{R}^m \rightarrow [0, 1]$ be a probability distribution which returns (\mathbf{u}, \mathbf{v}) where the vector \mathbf{v} satisfies $\|\mathbf{v}\| \leq T$. Further, define $M := \exp(\frac{\pi}{\alpha^2} + 1)$ and $\epsilon := 2^{\frac{1+\varepsilon}{1-\varepsilon}} \exp(-\alpha^2 \cdot \frac{\pi-1}{\pi^2})$. Then, the statistical distance between distributions RejSamp and SimRS defined in Figure 3 is at most $\frac{\epsilon}{2M} + \frac{2\varepsilon}{M}$. Moreover, the probability that RejSamp outputs something is at least $\frac{1-\epsilon}{M} \left(1 - \frac{4\varepsilon}{(1+\varepsilon)^2}\right)$.*

2.4 Hardness Assumptions

We recall the standard lattice-based computational assumptions: Module-SIS [LS15] and NTRU [HPS98; LTV12; SS13].

Definition 2.12 (Module-SIS). *Let $q = q(\lambda)$, $n = n(\lambda)$, $m = m(\lambda)$, $\beta = \beta(\lambda)$ and $N = N(\lambda)$. We say that the $\text{MSIS}_{n,m,N,q,\beta}$ assumption holds if for any PPT adversary \mathcal{A} , the following holds:*

$$\Pr \left[\mathbf{A}\mathbf{s} = \mathbf{0} \wedge 0 < \|\mathbf{x}\| \leq \beta \mid \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m} \\ \mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}) \end{array} \right] \leq \text{negl}(\lambda) .$$

Definition 2.13 (NTRU). *Let $q = q(\lambda)$, $N = N(\lambda)$ and $\chi = \chi(\lambda)$ be a probability distribution on \mathcal{R} . We say the (decisional) $\text{NTRU}_{q,N,\chi}$ assumption holds if the distributions*

$$\left\{ h \mid \begin{array}{l} f, g \leftarrow \chi \\ h = g \cdot f^{-1} \pmod{q} \end{array} \right\} \text{ and } \{h \mid h \leftarrow \mathcal{R}_q\}$$

are computationally indistinguishable, conditioned on f being invertible over \mathcal{R}_q .

2.5 NTRU Lattices

As defined before, let N be a power of two, q a positive integer and $f, g \in \mathcal{R}$ such that f is invertible over \mathcal{R}_q . Let $h = g/f \in \mathcal{R}_q$. The NTRU lattice associated to h and q is defined as

$$\Lambda_{h,q} := \{(u, v) \in \mathcal{R}^2 : u + vh = 0 \pmod{q}\} .$$

Then, $\Lambda_{h,q}$ is a $2N$ -dimensional full-rank lattice generated by the rows of

$$A_{h,q} := \begin{bmatrix} -\text{rot}(h) \mathbf{I}_N & \\ q \cdot \mathbf{I}_N & 0 \end{bmatrix} \in \mathbb{Z}^{2N \times 2N} .$$

We recall that there is an efficient algorithm NTRU.TrapGen [DLP14], which given modulus q and the ring dimension N , outputs $h \in \mathcal{R}_q$ and a short basis of $\Lambda_{h,q}$.

Lemma 2.14 (NTRU Trapdoor Generation). *There is a PPT algorithm $\text{NTRU.TrapGen}(q, N)$ which outputs $h \in \mathcal{R}_q$ and a basis \mathbf{T}_{NTRU} of $\Lambda_{h,q}$ such that $\|\tilde{\mathbf{T}}_{\text{NTRU}}\| \leq 1.17\sqrt{q}$. Further, if $\mathfrak{s} = \Omega(\sqrt{q})$ then h is computationally indistinguishable from random by the $\text{NTRU}_{q,N,\mathcal{D}_{\mathfrak{s}}}$ assumption.*

The short basis can now be used for preimage sampling using the well-known GPV framework [GPV08] and its concrete instantiation in [DLP14]. Namely, for any $c \in \mathcal{R}$, one can efficiently sample $(u, v) \in \mathcal{R}^2$ from a discrete Gaussian distribution conditioned on $u + vh = c \pmod{q}$.

Lemma 2.15 (NTRU Preimage Sampling). *Let $\mathfrak{s} = \Omega(\sqrt{q})$. There is a PPT algorithm NTRU.SamplePre , which takes $(h, \mathbf{T}_{\text{NTRU}}) \leftarrow \text{NTRU.TrapGen}(q, N)$, a target vector $c \in \mathcal{R}_q$ and standard deviation $\mathfrak{s} > 0$ as input, and outputs a pair $(u, v) \in \mathcal{R}_q^{m+2}$ such that*

$$\Delta \left([h \ 1]_{\mathfrak{s}}^{-1}(c), \text{NTRU.SamplePre}(h, \mathbf{T}_{\text{NTRU}}, c, \mathfrak{s}) \right) \leq \text{negl}(\lambda) .$$

2.6 Gadget Trapdoors

In this section, we recall the notion of gadget trapdoors as in [MP12], reformulate them for the module setting and state the key results on efficient sampling preimages using trapdoors.

We say that a matrix $\mathbf{G} \in \mathcal{R}_q^{n \times t}$ is primitive if its columns generate \mathcal{R}_q^n , i.e. if $\mathbf{G} \cdot \mathcal{R}^t = \mathcal{R}_q^n$. Note that if \mathbf{G} is primitive, then $\text{rot}(\mathbf{G})$ also is w.r.t. \mathbb{Z}_q^{nN} (i.e. $\text{rot}(\mathbf{G})\mathbb{Z}^{tN} = \mathbb{Z}_q^{nN}$). We also recall the notion of a gadget trapdoor.

Definition 2.16. *Let $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, $\mathbf{H} \in \mathcal{R}_q^{n \times n}$, $\mathbf{G} \in \mathcal{R}_q^{n \times t}$ with $t \geq n$ and \mathbf{H} invertible over \mathcal{R}_q . A \mathbf{G} -trapdoor for \mathbf{A} with tag \mathbf{H} is a matrix $\mathbf{R} \in \mathcal{R}_q^{m \times t}$ with $\mathbf{AR} = \mathbf{HG}$. The quality of a trapdoor is $s_1(\mathbf{R})$.*

When not specified, we set the tag $\mathbf{H} := \mathbf{I}$. In fact, all the theorems in this section can be generalised with a tag.

In this work, we consider one particular primitive matrix that naturally represents δ -base decomposition which we call the gadget matrix.

Definition 2.17 (Gadget Matrix). *Let $\delta \geq 2$. We set $\tilde{q} := \lceil \log_{\delta} q \rceil + 1$, and $\mathbf{g}^{\top} = [1, \delta, \dots, \delta^{\tilde{q}-1}] \in \mathcal{R}_q^{1 \times \tilde{q}}$ and $\mathbf{G}_n := \mathbf{I}_n \otimes \mathbf{g}^{\top} \in \mathcal{R}_q^{n \times n\tilde{q}}$. When the dimension are clear from context we simply write \mathbf{G} . Write $\mathbf{G}_n^{-1} : \mathcal{R}_q^{n \times t} \rightarrow \mathcal{R}_q^{n\tilde{q} \times t}$ for the inverse function that takes a matrix of entries in \mathcal{R}_q , and decomposes each entry w.r.t. We also write \mathbf{g}^{-1} for \mathbf{G}_1^{-1} .*

[MP12, Lemma 5.3] says that having a \mathbf{G} -trapdoor for some matrix \mathbf{A} enables to translate any nice basis of \mathbf{G} 's induced lattice into one for \mathbf{A} 's, whose shortness is proportional to the quality of the trapdoor.

Lemma 2.18. Let $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, $\mathbf{G} \in \mathcal{R}_q^{n \times t}$ be the gadget matrix with decomposition base δ , and suppose there exists a \mathbf{G} -trapdoor \mathbf{R} for \mathbf{A} . Then, there is a basis $\mathbf{S}_\mathbf{A}$ of $\Lambda^\perp(\mathbf{A})$ which satisfies $\|\tilde{\mathbf{S}}_\mathbf{A}\| \leq (s_1(\mathbf{R}) + 1)\sqrt{\delta^2 + 1}$. In particular, if $\|\mathbf{R}\| \leq \beta$ then for $\epsilon = \text{negl}(\lambda)$:

$$\eta_\epsilon(\Lambda^\perp(\mathbf{A})) \leq \beta\delta \cdot \omega(\sqrt{N \log mN}) .$$

We now give crucial properties about the trapdoor generation from [MP12].

Lemma 2.19 (Trapdoor Generation). Let $N, n > 0, t = n\tilde{q}$ and $\mathbf{G}_n \in \mathcal{R}_q^{n \times t}$ be the gadget matrix. Take $m \geq t + n + \omega(\log \lambda)$. Then, there is a PPT algorithm $\text{TrapGen}(n, m)$ that with an overwhelming probability returns two matrices $(\mathbf{A}, \mathbf{R}) \in \mathcal{R}_q^{n \times m} \times \mathcal{R}_q^{m \times t}$ such that $\mathbf{A}\mathbf{R} = \mathbf{G}_n$ and $\|\mathbf{R}\| \leq \mathfrak{s}\sqrt{2t(m-t)N}$ where $\mathfrak{s} > 2N \cdot q^{\frac{n}{m-t} + \frac{2}{N(m-t)}}$. Moreover, \mathbf{A} is statistically close to a uniformly random matrix in $\mathcal{R}_q^{n \times m}$.

Proof. Let $m' = m - t$. Consider the following algorithm [MP12, Alg 1]:

1. Sample $\bar{\mathbf{A}} \leftarrow \mathcal{R}_q^{n \times m'}$.
2. Sample a matrix $\bar{\mathbf{R}} \leftarrow \mathcal{D}_\mathfrak{s}^{m'N \times tN}$ from a discrete Gaussian distribution.
3. Return $\mathbf{A} := [\bar{\mathbf{A}} | \mathbf{G}_n - \bar{\mathbf{A}}\bar{\mathbf{R}}]$ and $\mathbf{R} := \begin{bmatrix} \bar{\mathbf{R}} \\ \mathbf{I}_t \end{bmatrix}$

First, $\mathbf{A}\mathbf{R} = \mathbf{G}$ as desired and $\|\mathbf{R}\| \leq \sqrt{t(\mathfrak{s}^2 m'N + 1)} \leq \mathfrak{s}\sqrt{2t(m-t)N}$ with an overwhelming probability by Lemma 2.9 for $t = \sqrt{2\pi}$. To argue pseudorandomness, we apply Lemma 2.7 and the hybrid argument to get that $\bar{\mathbf{A}}\bar{\mathbf{R}}$ is statistically close to uniform over $\mathcal{R}_q^{n \times t}$, and thus so is \mathbf{A} . \square

The next lemma states that given a short \mathbf{G} -trapdoor matrix \mathbf{R} for \mathbf{A} , one can efficiently sample preimages of \mathbf{A} according to the discrete Gaussian distribution.

Lemma 2.20 (Preimage Sampling). Let $N, n, m > 0$ and $t = n\tilde{q}$. Then, there exists a PPT algorithm $\text{SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{v}, \sigma)$ that takes as input a matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, a \mathbf{G}_n -trapdoor $\mathbf{R} \in \mathcal{R}_q^{m \times t}$ for \mathbf{A} with a tag \mathbf{H} , a target vector $\mathbf{v} \in \mathcal{R}_q^n$ in the column-span of \mathbf{A} , and a Gaussian parameter σ , and outputs a vector $\mathbf{s} \in \mathcal{R}_q^m$ such that $\mathbf{A}\mathbf{s} = \mathbf{v}$. Further, if $\sigma \geq \delta s_1(\mathbf{R}) \cdot \omega(\sqrt{\log nN})$, then the statistical distance between the following distributions is negligible:

$$\{\mathbf{s} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{v}, \sigma)\} \text{ and } \{\mathbf{s} \leftarrow \mathbf{A}_\sigma^{-1}(\mathbf{v})\} .$$

We extend this algorithm for matrices, i.e. for a matrix $\mathbf{V} \in \mathcal{R}_q^{n \times \ell}$ with columns $\mathbf{v}_1, \dots, \mathbf{v}_\ell$, we define $\text{SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{V}, \sigma)$ to be the algorithm which returns a matrix $\mathbf{S} \in \mathcal{R}_q^{m \times \ell}$, where the i -th column is the output of $\text{SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{v}_i, \sigma)$.

2.7 Commitment Scheme

We recall the notion of a commitment scheme, which is a crucial component of various proof systems. As folklore in lattice-based cryptography, we introduce the slack space, which has a role in the binding property.

Definition 2.21. Let $\text{CM} = (\text{Setup}, \text{Commit}, \text{Open})$ be a triple of PPT algorithms. We say that CM is a commitment scheme over \mathcal{M} with slack space \mathcal{S} if it has the following syntax:

- $\text{Setup}(1^\lambda) \rightarrow \text{crs}$ takes a security parameter λ (specified in unary) and outputs a common reference string crs .

- $\text{Commit}(\text{crs}, m) \rightarrow (C, \text{st})$ takes a common reference string crs a message $m \in \mathcal{M}$ and outputs a commitment C and decommitment state st .
- $\text{Open}(\text{crs}, C, m, \text{st}, c)$ takes a common reference string crs , a commitment C , a message $m \in \mathcal{M}$, a decommitment state st and a relaxation factor¹² $c \in \mathcal{S}$ and outputs a bit indicating whether C is a valid commitment to m under crs .

We define the key properties of the commitment scheme: correctness, (relaxed) binding and hiding. In the following, we denote the message space as \mathcal{M} and the slack space as \mathcal{S} .

Definition 2.22 (Completeness). We say that a commitment scheme $\text{CM} = (\text{Setup}, \text{Commit}, \text{Open})$ satisfies completeness if there exists a global relaxation factor $c^* \in \mathcal{S}$ such that for every $m \in \mathcal{M}$:

$$\Pr \left[\text{Open}(\text{crs}, C, m, \text{st}, c^*) = 1 \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ C, \text{st} \leftarrow \text{Commit}(\text{crs}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda) .$$

Definition 2.23 (Relaxed Binding). A commitment scheme $\text{CM} = (\text{Setup}, \text{Commit}, \text{Open})$ satisfies relaxed binding if for every PPT adversary \mathcal{A} :

$$\Pr \left[\begin{array}{l} m \neq m' \wedge m, m' \in \mathcal{M} \wedge \\ \text{Open}(\text{crs}, C, m, \text{st}, c) = \text{Open}(\text{crs}, C, m', \text{st}', c') = 1 \end{array} \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ (C, (m, \text{st}, c), (m', \text{st}', c')) \leftarrow \mathcal{A}(\text{crs}) \end{array} \right] = \text{negl}(\lambda) .$$

Definition 2.24 (Hiding). A commitment scheme $\text{CM} = (\text{Setup}, \text{Commit}, \text{Open})$ satisfies hiding if for every (stateful) PPT adversary \mathcal{A} :

$$\Pr \left[b' = b \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda), (m_0, m_1) \leftarrow \mathcal{A}(\text{crs}) \\ b \leftarrow \{0, 1\} \\ C, \text{st} \leftarrow \text{Commit}(\text{crs}, m_b) \\ b' \leftarrow \mathcal{A}(C) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda) .$$

2.8 Polynomial Commitment Scheme

We also recall the notion of polynomial commitment schemes. Polynomial commitment schemes extend commitment with the ability to prove evaluations of the committed polynomial.

Definition 2.25. Let $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Eval}, \text{Verify})$ be a tuple of algorithms. PC is a polynomial commitment scheme over a ring R with degree bound d and slack space \mathcal{S} if:

- $(\text{Setup}, \text{Commit}, \text{Open})$ is a commitment scheme over

$$\mathcal{M} := \left\{ (f_0, f_1, \dots, f_d) \in R^{d+1} : \sum_{i=0}^d f_i X^i \in R[X] \right\}$$

with slack space \mathcal{S} .

- $\text{Eval}(\text{crs}, C, u, \text{st}) \rightarrow \pi$ takes a common reference string crs , a commitment C , an evaluation point $u \in R$, auxiliary state st and outputs an evaluation proof π .
- $\text{Verify}(\text{crs}, C, u, z, \pi) \rightarrow 0/1$ takes a common reference string crs , a commitment C , an evaluation point $u \in R$, a claimed image $z \in R$, an evaluation proof π , and outputs a bit indicating whether π is a valid evaluation proof that the polynomial committed to in C evaluates to z at the point u .

¹² We implicitly assume that if $c \notin \mathcal{S}$ then Open automatically returns 0.

We also consider a setting in which `Eval` and `Verify` are replaced with an interactive two-party protocol between a prover and a verifier, and refer to that setting as an interactive polynomial commitment scheme.

Additionally, we require that the evaluations procedure satisfy some additional properties that we detail next. For simplicity, we give these definitions for non-interactive polynomial commitments, the interactive variant follows similarly.

Definition 2.26 (Evaluation Completeness). We say that a polynomial commitment scheme $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Eval}, \text{Verify})$ satisfies completeness if for every polynomial $f \in R^{\leq d}[\mathbf{X}]$ and any evaluation point $u \in R$:

$$\Pr \left[\text{Verify}(\text{crs}, C, u, f(u), \pi) = 0 \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ C, \text{st} \leftarrow \text{Commit}(\text{crs}, f) \\ \pi \leftarrow \text{Eval}(\text{crs}, C, u, \text{st}) \end{array} \right] = \text{negl}(\lambda) .$$

Definition 2.27 (Knowledge Soundness). We say that a polynomial commitment scheme $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Eval}, \text{Verify})$ is knowledge sound with knowledge error κ if for all stateful PPT adversaries \mathcal{P}^* , there exists an expected PPT extractor \mathcal{E} such that

$$\Pr \left[b = 1 \wedge (\text{Open}(\text{crs}, C, f, \text{st}, c) \neq 1 \vee f(u) \neq z) \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ (C, u, z, \pi) \leftarrow \mathcal{P}^*(\text{crs}) \\ b = \text{Verify}(\text{crs}, C, u, z, \pi) \\ (f, \text{st}, c) \leftarrow \mathcal{E}^{\mathcal{P}^*}(\text{crs}, C, u, z, \pi) \end{array} \right] \leq \kappa(\lambda) .$$

Here, the extractor \mathcal{E} has a black-box oracle access to the (malicious) prover \mathcal{P}^* and can rewind it to any point in the interaction.

2.9 Interactive Proofs

Let $R \subseteq \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^*$ be a ternary relation. If $(\mathbf{i}, \mathbf{x}, \mathbf{w}) \in R$, we say that \mathbf{i} is an index, \mathbf{x} is a statement and \mathbf{w} is a witness for \mathbf{x} . We denote $R(\mathbf{i}, \mathbf{x}) = \{\mathbf{w} : R(\mathbf{i}, \mathbf{x}, \mathbf{w}) = 1\}$. In this work, we only consider NP relations R for which a witness w can be verified in time $\text{poly}(|\mathbf{i}|, |u|)$ for all $(\mathbf{i}, \mathbf{x}, \mathbf{w}) \in R$.

A proof system $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ for relation R consists of three PPT algorithms: the `Setup` algorithm, prover \mathcal{P} , and the verifier \mathcal{V} . The latter two are interactive and stateful. We write $(tr, b) \leftarrow \langle \mathcal{P}(\mathbf{i}, \mathbf{x}, \mathbf{w}), \mathcal{V}(\mathbf{i}, \mathbf{x}) \rangle$ for running \mathcal{P} and \mathcal{V} on inputs $\mathbf{i}, \mathbf{x}, \mathbf{w}$ and \mathbf{i}, \mathbf{x} respectively and getting communication transcript tr and the verifier's decision bit b . We use the convention that $b = 0$ means reject and $b = 1$ means accept the prover's claim of knowing \mathbf{w} such that $(\mathbf{x}, \mathbf{w}) \in R$. If tr contains a \perp then we say that \mathcal{P} aborts. Unless stated otherwise, we will assume that the first and the last message are sent from a prover. Hence, the protocol between \mathcal{P} and \mathcal{V} has an odd number of rounds. A Σ -protocol is a three-round protocol. Further, we say a protocol is *public coin* if the verifier's challenges are chosen uniformly at random independently of the prover's messages.

We recall a few basic properties of interactive proof systems: completeness and knowledge soundness.

Definition 2.28 (Completeness). A proof system $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ for the relation R has statistical completeness with correctness error ϵ if for all adversaries \mathcal{A} ,

$$\Pr \left[b = 0 \wedge (\mathbf{i}, \mathbf{x}, \mathbf{w}) \in R \mid \begin{array}{l} \mathbf{i} \leftarrow \text{Setup}(1^\lambda) \\ (\mathbf{x}, \mathbf{w}) \leftarrow \mathcal{A}(\mathbf{i}) \\ (tr, b) \leftarrow \langle \mathcal{P}(\mathbf{i}, \mathbf{x}, \mathbf{w}), \mathcal{V}(\mathbf{i}, \mathbf{x}) \rangle \end{array} \right] \leq \epsilon(\lambda) .$$

Definition 2.29 (Knowledge Soundness). A proof system $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ for the relation R is knowledge sound with knowledge error κ if there exists an expected PPT extractor \mathcal{E} such that for any stateful PPT adversary \mathcal{P}^* :

$$\Pr \left[b = 1 \wedge (\mathbf{i}, \mathbf{x}, \mathbf{w}) \notin R \mid \begin{array}{l} \mathbf{i} \leftarrow \text{Setup}(1^\lambda) \\ (\mathbf{x}, \mathbf{st}) \leftarrow \mathcal{P}^*(\mathbf{i}) \\ (tr, b) \leftarrow \langle \mathcal{P}^*(\mathbf{i}, \mathbf{x}, \mathbf{st}), \mathcal{V}(\mathbf{i}, \mathbf{x}) \rangle \\ \mathbf{w} \leftarrow \mathcal{E}^{\mathcal{P}^*}(\mathbf{i}, \mathbf{x}) \end{array} \right] \leq \kappa(\lambda) .$$

Here, the extractor \mathcal{E} has a black-box oracle access to the (malicious) prover \mathcal{P}^* and can rewind it to any point in the interaction.

2.10 Coordinate-Wise Special-Soundness

We generalise the notion of *special-soundness* the following way. Let S be a set and $\ell \in \mathbb{N}$. Namely, take two vectors $\mathbf{x} := (x_1, \dots, x_\ell), \mathbf{y} := (y_1, \dots, y_\ell) \in S^\ell$. Then, we define the following relation “ \equiv_i ” for fixed $i \in [\ell]$ as:

$$\mathbf{x} \equiv_i \mathbf{y} \iff x_i \neq y_i \wedge \forall j \in [\ell] \setminus \{i\}, x_j = y_j .$$

That is, vectors \mathbf{x} and \mathbf{y} have the same values in all coordinates apart from the i -th one. For $\ell = 1$, the relations boils down to checking whether two elements are distinct. Further, we can define the set

$$\text{SS}(S, \ell) := \left\{ (\mathbf{x}_1, \dots, \mathbf{x}_{\ell+1}) \in (S^\ell)^{\ell+1} : \exists k \in [\ell+1], \forall i \in [\ell], \exists j \in [\ell+1] \setminus \{k\}, \mathbf{x}_k \equiv_i \mathbf{x}_j \right\} .$$

As a simple example, $((0, 0), (1, 0), (0, 1)) \in \text{SS}(\mathbb{Z}_2, 2)$ – the vector $(0, 0)$ differs from $(1, 0)$ (resp. $(0, 1)$) exactly in the first (resp. second) coordinate. Note that for $\ell = 1$, this set simply contains pairs of distinct elements in S .

We are ready to define the notion of coordinate-wise special-soundness. We start with the case for Σ -protocols.

Definition 2.30 (Coordinate-Wise Special-Soundness). Let $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ be public-coin three-round interactive proof system for relation R , and suppose the challenge space of \mathcal{V} is $\mathcal{C} = S^\ell$. We say that Π is ℓ -coordinate-wise special-sound if there exists a polynomial time algorithm that on input an index \mathbf{i} , statement \mathbf{x} and $\ell + 1$ accepting transcripts $(a, \mathbf{c}_i, z_i)_{i \in [\ell+1]}$, with $(\mathbf{c}_1, \dots, \mathbf{c}_{\ell+1}) \in \text{SS}(S, \ell)$ and common first message a , outputs a witness $\mathbf{w} \in R(\mathbf{i}, \mathbf{x})$.

Clearly, we obtain the standard special-soundness property if $\ell = 1$. Next, we extend this notion to multi-round protocols via a tree of transcripts. For simplicity, we assume that in each round the verifier picks challenge uniformly at random from the same challenge space S^ℓ , which will be the case for most of our protocols.

Definition 2.31 (Tree of Transcripts). Let $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ be public-coin $(2\mu + 1)$ -round interactive proof system for relation R , where in each round the verifier picks a uniformly random challenge from S^ℓ . A tree of transcripts is a set of $K = (\ell + 1)^\mu$ arranged in the following tree structure. The nodes in the tree correspond to the prover’s messages and the edges correspond to the verifier’s challenges. Each node at depth i has exactly $\ell + 1$ children corresponding to $\ell + 1$ distinct challenges which, as a vector, lie in $\text{SS}(S, \ell)$. Every transcript corresponds to exactly one path from the root to a leaf node.

We say that Π is ℓ -coordinate-wise special-sound if there is a polynomial time algorithm that given an index \mathbf{i} , statement \mathbf{x} and the tree of transcripts, outputs a witness $\mathbf{w} \in R(\mathbf{i}, \mathbf{x})$.

Forking strategies. Next, we show that coordinate-wise special-soundness implies knowledge soundness. Our knowledge extraction approach can be described by the following *collision game* [ACK21]. Consider a binary matrix $H \in \{0, 1\}^{R \times N}$ where $N \in \mathbb{N}$. One interpretation is that the R rows correspond to the prover's randomness and the N columns correspond to the verifier's randomness, or alternatively, the verifier samples a challenge $c \leftarrow \mathcal{C}$ uniformly at random where \mathcal{C} has size N . An entry of H equals 1 if and only if the corresponding protocol transcript is accepting. The knowledge extractor will run the following collision game.

1. First, sample $(r, i) \leftarrow [R] \times [N]$ and check if $H(r, i) = 1$. If not, it aborts.
2. If $H(r, i) = 1$, then it samples $i^* \leftarrow [N]$ without replacement until it obtains $i^* \neq i$ such that $H(r, i^*) = 1$.

The following lemma states the expected runtime and success probability of the algorithm above.

Lemma 2.32 ([ACK21]). *Let $H \in \{0, 1\}^{R \times N}$ and define ϵ to be the fraction of 1-entries in H . Then, the expected number of H -entries queried in the collision game is at most 2 and the probability of the collision-game is at least $\epsilon - \frac{1}{N}$.*

We use this result in the context of proving knowledge soundness of coordinate-wise special-sound protocols. We start from three-round protocols.

Lemma 2.33. *Let $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ be public-coin three-round interactive proof system for relation R and suppose the challenge space of \mathcal{V} is S^ℓ where $\ell = \text{poly}(\lambda)$. If Π is ℓ -coordinate-wise special-sound then it is knowledge sound with knowledge error $\ell/|S|$.*

Proof. Let \mathcal{P}^* be a deterministic, malicious prover which convinces the verifier with probability ϵ . We can define the following ℓ binary matrices $H_1, \dots, H_\ell \in \{0, 1\}^{R \times N}$ where $R = |S|^{\ell-1}$ and $N = |S|$. For a matrix H_i , the rows are indexed by all the possible choices of

$$\bar{\mathbf{c}}_i := (c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_\ell) \in S^{\ell-1} \quad (11)$$

and the columns are indexed by all possible choices of c_i . We define $H_i(\bar{\mathbf{c}}_i, c_i) = 1$ if and only if \mathcal{P}^* convinces the verifier for the challenge (c_1, \dots, c_ℓ) . This implies

$$H_1(\bar{\mathbf{c}}_1, c_1) = \dots = H_\ell(\bar{\mathbf{c}}_\ell, c_\ell) \quad (12)$$

where $\bar{\mathbf{c}}_i$ are defined as in (11). Further, the fraction of 1-entries in each H_i is exactly ϵ and checking one entry of any matrix H_i requires running \mathcal{P}^* once.

We define the knowledge extractor \mathcal{E} as follows:

1. First, sample $\mathbf{c}_0 := (c_1, \dots, c_\ell) \leftarrow S^\ell$ and check if $H_1(\bar{\mathbf{c}}_1, c_1) = 1$. If not, abort.
2. For $i = 1, 2, \dots, \ell$:
 - (a) Sample $c_i^* \leftarrow [N]$ without replacement until obtaining $c_i^* \neq c_i$ such that $H_i(\bar{\mathbf{c}}_i, c_i^*) = 1$.
 - (b) If no such challenge is found, abort.
 - (c) Set $\mathbf{c}_i := (c_1, \dots, c_{i-1}, c_i^*, c_{i+1}, \dots, c_\ell)$
3. Output the corresponding transcripts for challenges $(\mathbf{c}_0, \dots, \mathbf{c}_\ell)$.

By construction, the output of the extractor satisfies $(\mathbf{c}_0, \dots, \mathbf{c}_\ell) \in \text{SS}(S, \ell)$, and thus one can use the property of coordinate-wise special-soundness to conclude the proof. What we have left to do is to analyse the expected run-time and success probability of \mathcal{E} .

Using the property in Equation (12), it is easy to see that \mathcal{E} runs ℓ copies of the collision game for each matrix H_1, \dots, H_ℓ , where the first step is the same for all copies. Hence, by Lemma 2.32 and the linearity of expectation, the expected number of queries to \mathcal{P}^* by \mathcal{E} is at most $\ell + 1$. Further, by the union bound, the probability that \mathcal{E} fails in at least one of the collision games is at most $\frac{\ell}{|S|}$. \square

Following the argument by Attema et al. [ACK21], one can prove an analogous result for multi-round protocols.

Lemma 2.34. *Let $\Pi = (\text{Setup}, \mathcal{P}, \mathcal{V})$ be public-coin $(2\mu + 1)$ -round interactive proof system for relation R and suppose the challenge space of \mathcal{V} in each round is S^ℓ . If Π is ℓ -coordinate-wise special-sound and $\ell^\mu = \text{poly}(\lambda)$, then it is knowledge sound with knowledge error $\mu\ell/|S|$.*

The resulting knowledge extractor runs the malicious prover $(\ell + 1)^\mu$ times in expectation. Hence, in order to keep the knowledge extractor expected PPT, we need $\ell^\mu = \text{poly}(\lambda)$.

The result can be easily extended to the case, where in each i -th round the challenges from the verifier are picked from S^{ℓ_i} for $\ell_i > 0$. Then, the knowledge error becomes $(\ell_1 + \dots + \ell_\mu)/|S|$ and the extractor runs the malicious prover at most $\prod_{i=1}^\mu (\ell_i + 1)$ times.

Finally, using the exact methodology as in [AFK22], one can deduce that coordinate-wise special soundness implies (adaptive) knowledge soundness of the Fiat-Shamir transformed protocol in the random oracle model with knowledge error $(Q + 1) \cdot \mu\ell/|S|$, where Q is the number of random oracle queries made by an adversary. Since all the proofs remain almost identical (with an additional use of union bounds), we omit the concrete analysis.

3 Power-BASIS Assumption

Our construction of the polynomial commitment will rely on a new lattice-based assumption PowerBASIS which is a special case of the BASIS assumption¹³ introduced by Wee and Wu [WW23]. We begin by adapting the latter assumption to the ring setting. Recall that \mathbf{G}_n is a gadget matrix with base δ as in Definition 2.17. We fix the modulus q and set $\tilde{q} := \lfloor \log_\delta q \rfloor + 1$.

Definition 3.1 (BASIS). *Let $q, n, m, n', m', \ell, N, \sigma, \beta$ be lattice parameters. Let Samp be a PPT algorithm, which given a matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, outputs a matrix $\mathbf{B} \in \mathcal{R}_q^{n' \times m'}$ along with auxiliary information aux . We say the $\text{BASIS}_{n, m, n', m', N, q, \ell, \sigma, \beta}$ assumption holds w.r.t. Samp if for any PPT adversary \mathcal{A} :*

$$\Pr \left[\begin{array}{l} \mathbf{A}\mathbf{s} = \mathbf{0} \\ 0 < \|\mathbf{s}\| \leq \beta \end{array} \middle| \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}, (\mathbf{B}, \text{aux}) \leftarrow \text{Samp}(\mathbf{A}) \\ \mathbf{T} \leftarrow \mathbf{B}_\sigma^{-1}(\mathbf{G}_{n'}) \\ \mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{B}, \mathbf{T}, \text{aux}) \end{array} \right] \leq \text{negl}(\lambda) .$$

Intuitively, the BASIS assumption says that it is hard to find a short solution for \mathbf{A} , even when given a trapdoor for a matrix \mathbf{B} related to \mathbf{A} . The trapdoor allows the adversary to sample preimages of \mathbf{B} , and thus it is easy to break the assumption if \mathbf{B} contains too much information about \mathbf{A} , e.g. when $\mathbf{B} = \mathbf{A}$.

Furthermore, we provide three concrete instantiations of the sampling algorithm Samp .

¹³ BASIS stands for Basis-Augmented Shortest Integer Solution.

Definition 3.2 (BASIS Instantiations). We consider three concrete instantiations of the BASIS assumption:

- **StructBASIS** $_{n,m,N,q,\ell,\sigma,\beta}$: The sampling algorithm $\text{Samp}(\mathbf{A})$ first generates a row $\mathbf{a}^\top \leftarrow \mathcal{R}_q^\ell$ and sets

$$\mathbf{A}^* := \begin{bmatrix} \mathbf{a}^\top \\ \mathbf{A} \end{bmatrix} \in \mathcal{R}_q^{(n+1) \times \ell} . \quad (13)$$

Further, it samples $\mathbf{W}_i \leftarrow \text{GL}(n+1, \mathcal{R}_q)$ for all $i \in [\ell]$, and outputs

$$\mathbf{B}_\ell := \left[\begin{array}{ccc|c} \mathbf{W}_1 \mathbf{A}^* & & & -\mathbf{G}_{n+1} \\ & \ddots & & \vdots \\ & & \mathbf{W}_\ell \mathbf{A}^* & -\mathbf{G}_{n+1} \end{array} \right] \quad \text{and} \quad \text{aux} := (\mathbf{W}_1, \dots, \mathbf{W}_\ell) .$$

- **PowerBASIS** $_{n,m,N,q,\ell,\sigma,\beta}$: Here, $\text{Samp}(\mathbf{A})$ generates a row $\mathbf{a}^\top \leftarrow \mathcal{R}_q^\ell$ and sets \mathbf{A}^* as in (13). Then, it samples $\mathbf{W} \leftarrow \text{GL}(n+1, \mathcal{R}_q)$, and outputs

$$\mathbf{B}_\ell := \left[\begin{array}{ccc|c} \mathbf{W}^0 \mathbf{A}^* & & & -\mathbf{G}_{n+1} \\ & \ddots & & \vdots \\ & & \mathbf{W}^{\ell-1} \mathbf{A}^* & -\mathbf{G}_{n+1} \end{array} \right] \quad \text{and} \quad \text{aux} := \mathbf{W} .$$

- **PRISIS** $_{n,m,N,q,\ell,\sigma,\beta}$: $\text{Samp}(\mathbf{A})$ samples a row $\mathbf{a}^\top \leftarrow \mathcal{R}_q^\ell$ and sets \mathbf{A}^* as in (13). Then, it samples $w \leftarrow \text{GL}(1, \mathcal{R}_q)$, and outputs

$$\mathbf{B}_\ell := \left[\begin{array}{ccc|c} w^0 \mathbf{A}^* & & & -\mathbf{G}_{n+1} \\ & \ddots & & \vdots \\ & & w^{\ell-1} \mathbf{A}^* & -\mathbf{G}_{n+1} \end{array} \right] \quad \text{and} \quad \text{aux} := w .$$

Informally, the StructBASIS variant corresponds to the structured version of the BASIS assumption used to build functional commitments [WW23]. PowerBASIS is the special case, where instead of picking ℓ uniformly random invertible matrices \mathbf{W}_i , one takes a single invertible matrix, and sets $\mathbf{W}_i := \mathbf{W}^{i-1}$ for $i \in [\ell]$. Finally, PRISIS is the instance where each $\mathbf{W}_i := w^{i-1} \mathbf{I}_{n+1}$ for $i \in [\ell]$ and $w \in \mathcal{R}_q$ is an invertible element.

Intuitively, StructBASIS seems to be the hardest variant to break out of the three since it carries the least structure. Then, PowerBASIS should be an easier problem due to the very specific relation between matrices \mathbf{W}_i . Finally, PRISIS carries a lot of structure, since it introduces commutativity between the matrices \mathbf{W}_i and \mathbf{A}^* , i.e. $w^{i-1} \mathbf{A} = \mathbf{A} (w^{i-1} \cdot \mathbf{I}_m)$, which can somehow be useful for the adversary to break the assumption.

Remark 3.3. To simplify reductions in the paper, we explicitly require the matrices \mathbf{W}_i to be invertible (unlike in [WW23]). Note that this condition can be dropped by arguing that, depending on the parameters q and N , with overwhelming probability a uniformly random matrix \mathbf{W} is invertible over \mathcal{R}_q (see [EZS+19, Appendix C] and [BTT22, Appendix C.3] for the bounds).

3.1 Hardness of BASIS for Low Dimensions

We analyse the relationship between the three newly introduced instantiations for the dimension $\ell = 2$. To this end, we analyse the following technical lemma which will be used in all our results of this section. Intuitively, it says that if one can find a short solution to a specific linear equation, then one can also build a BASIS trapdoor.

Lemma 3.4. *Let $n, m, N > 0$ and $\alpha \geq 1$. Denote $t = n\tilde{q}$. Then, there exists an efficient deterministic algorithm, that given as input a matrix $\mathbf{A}^* \in \mathcal{R}_q^{n \times m}$, invertible $\mathbf{W}_1, \mathbf{W}_2, \mathbf{H} \in \text{GL}(n, \mathcal{R}_q)$ and two matrices $\mathbf{T}_1, \mathbf{T}_2 \in \mathcal{R}_q^{m \times t}$, which satisfy $\|(\mathbf{T}_1, \mathbf{T}_2)\| \leq \alpha$ for $i = 1, 2$ and*

$$\mathbf{W}_1 \mathbf{A} \mathbf{T}_1 - \mathbf{W}_2 \mathbf{A} \mathbf{T}_2 = \mathbf{H} \mathbf{G}_n ,$$

outputs a tag $\mathbf{H}^ \in \text{GL}(2n, \mathcal{R}_q)$ and a \mathbf{G}_{2n} -trapdoor \mathbf{S} for the matrix \mathbf{B} defined as:*

$$\mathbf{B} := \begin{bmatrix} \mathbf{W}_1 \mathbf{A}^* & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & \mathbf{W}_2 \mathbf{A}^* & -\mathbf{G} \end{bmatrix}$$

with a tag \mathbf{H}^ , where $\|\mathbf{S}\| \leq \sqrt{2(\alpha^2 + t^2 N)}$.*

Proof. Define the following matrices:

$$\begin{aligned} \mathbf{S}_{1,3} &:= \mathbf{G}^{-1}(\mathbf{W}_1 \mathbf{A}^* \mathbf{T}_1 - \mathbf{H} \mathbf{G}_n) = \mathbf{G}^{-1}(\mathbf{W}_2 \mathbf{A} \mathbf{T}_2) \\ \mathbf{S}_{2,3} &:= \mathbf{G}^{-1}(-\mathbf{W}_1 \mathbf{A}^* \mathbf{T}_2 - \mathbf{H} \mathbf{G}_n) = \mathbf{G}^{-1}(-\mathbf{W}_1 \mathbf{A}^* \mathbf{T}_1). \end{aligned}$$

Then, by construction we get:

$$\begin{bmatrix} \mathbf{W}_1 \mathbf{A}^* & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & \mathbf{W}_2 \mathbf{A}^* & -\mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{T}_1 & -\mathbf{T}_1 \\ \mathbf{T}_2 & -\mathbf{T}_2 \\ \mathbf{S}_{1,3} & \mathbf{S}_{2,3} \end{bmatrix} = \begin{bmatrix} \mathbf{H} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{H} \mathbf{G} \end{bmatrix} = \begin{bmatrix} \mathbf{H} & \mathbf{0} \\ \mathbf{0} & \mathbf{H} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{G} \end{bmatrix}.$$

By setting

$$\mathbf{S} := \begin{bmatrix} \mathbf{T}_1 & -\mathbf{T}_1 \\ \mathbf{T}_2 & -\mathbf{T}_2 \\ \mathbf{S}_{1,3} & \mathbf{S}_{2,3} \end{bmatrix} \quad \text{and} \quad \mathbf{H}^* := \begin{bmatrix} \mathbf{H} & \mathbf{0} \\ \mathbf{0} & \mathbf{H} \end{bmatrix} ,$$

we observe that \mathbf{S} is a \mathbf{G}_{2n} -trapdoor for \mathbf{B} with a tag \mathbf{H}^* and $\|\mathbf{S}\|^2 \leq 2\alpha^2 + 2t^2 N$, which concludes the proof. \square

Our first result says that StructBASIS and PowerBASIS are equivalent for the dimension $\ell = 2$.

Lemma 3.5 (StructBASIS \iff PowerBASIS). *Let $n, N, \beta \geq 1$ and $t := (n + 1)\tilde{q}$. Suppose $m \geq t + (n + 1) + \omega(\log \lambda)$ and $\mathfrak{s} > 2N \cdot q^{\frac{n+1}{m-t} + \frac{2}{N(m-t)}}$. If σ_0, σ_1 satisfy the following inequalities:*

$$\sigma_0 \geq \delta \mathfrak{s} N \cdot \omega(\sqrt{t(m-t) \log mN}), \quad \sigma_1 \geq \delta \sqrt{2tN(\sigma_1^2 m' + t)N} \cdot \omega(\sqrt{\log nN}),$$

where $m' = 2m + t$, then the following statements are true:

1. StructBASIS $_{n,m,N,q,2,\sigma_0,\beta}$ assumption holds under the PowerBASIS $_{n,m,N,q,2,\sigma_1,\beta}$ assumption.
2. PowerBASIS $_{n,m,N,q,2,\sigma_0,\beta}$ assumption holds under the StructBASIS $_{n,m,N,q,2,\sigma_1,\beta}$ assumption.

Proof. We only show the first statement since the other direction follows identically. Let \mathcal{A} be a PPT adversary for the StructBASIS $_{n,m,N,q,2,\sigma,\beta}$ problem and suppose it wins with probability ϵ . We provide a PPT algorithm \mathcal{B} for solving PowerBASIS $_{n,m,N,q,2,\sigma,\beta}$ which does the following. First, \mathcal{B} is given a tuple $(\mathbf{A}, \mathbf{B}, \mathbf{T}, \mathbf{W})$ where

$$\mathbf{B} := \begin{bmatrix} \mathbf{A}^* & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & \mathbf{W} \mathbf{A}^* & -\mathbf{G} \end{bmatrix} \quad \text{and} \quad \mathbf{T} := \begin{bmatrix} \mathbf{T}_{1,1} & \mathbf{T}_{1,2} \\ \mathbf{T}_{2,1} & \mathbf{T}_{2,2} \\ \mathbf{T}_{3,1} & \mathbf{T}_{3,2} \end{bmatrix}.$$

First, we claim that the following probability is negligible:

$$\epsilon_{\text{smooth}} := \Pr \left[\sigma_0 < \eta_\epsilon(\Lambda^\perp(\mathbf{B})) \mid \mathbf{A}^* \leftarrow \mathcal{R}_q^{(n+1) \times m} \right] .$$

Indeed, note that by Lemma 2.19 we obtain:

$$\Pr \left[\sigma_0 < \eta_\epsilon(\Lambda^\perp(\mathbf{B})) \mid (\mathbf{A}^*, \mathbf{R}) \leftarrow \text{TrapGen}(n+1, m) \right] \geq \epsilon_{\text{smooth}} - \text{negl}(\lambda) .$$

If $(\mathbf{A}^*, \mathbf{R}) \leftarrow \text{TrapGen}(n+1, m)$ then the following matrix \mathbf{R}^* is a \mathbf{G}_{2n} -trapdoor for \mathbf{B} with a tag \mathbf{H}^* , where:

$$\mathbf{R}^* := \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad \text{and} \quad \mathbf{H}^* := \begin{bmatrix} \mathbf{I}_{n+1} & \mathbf{0} \\ \mathbf{0} & \mathbf{W} \end{bmatrix} .$$

Moreover, $\|\mathbf{R}^*\| \leq 2s\sqrt{t(m-t)N}$ with an overwhelming probability. If this is the case then by assumption $\sigma_0 \geq \delta \cdot \|\mathbf{R}^*\| \cdot \omega(\sqrt{t(m-t)\log mN})$. Then, by combining Lemma 2.18 with Lemma 2.2, we obtain

$$\text{negl}(\lambda) = \Pr \left[\sigma_0 < \eta_\epsilon(\Lambda^\perp(\mathbf{B})) \mid (\mathbf{A}^*, \mathbf{R}) \leftarrow \text{TrapGen}(n+1, m) \right] \geq \epsilon_{\text{smooth}} - \text{negl}(\lambda)$$

and thus $\sigma_0 \geq \eta_\epsilon(\Lambda^\perp(\mathbf{B}))$ with an overwhelming probability, where \mathbf{B} is the matrix received by \mathcal{B} . Thus, we can apply Lemma 2.8 to deduce that with an overwhelming probability

$$\left\| \begin{bmatrix} \mathbf{T}_{1,1} \\ \mathbf{T}_{1,2} \end{bmatrix} \right\| \leq \alpha := \sigma_0 \sqrt{m'tN} .$$

Further, by simple calculation we can deduce that

$$\mathbf{A}^* \mathbf{T}_{1,1} - \mathbf{W} \mathbf{A}^* \mathbf{T}_{1,2} = \mathbf{G} .$$

The reduction \mathcal{B} now samples a uniformly random $\mathbf{W}_1 \leftarrow \text{GL}(n+1, \mathcal{R}_q)$ and defines $\mathbf{W}_2 := \mathbf{W}_1 \mathbf{W}$. Thus

$$\mathbf{W}_1 \mathbf{A}^* \mathbf{T}_{1,1} - \mathbf{W}_2 \mathbf{A}^* \mathbf{T}_{1,2} = \mathbf{W}_1 \mathbf{G} .$$

By applying Lemma 3.4, \mathcal{B} can obtain a $\mathbf{G}_{2(n+1)}$ -trapdoor \mathbf{S} for

$$\mathbf{B}' := \begin{bmatrix} \mathbf{W}_1 \mathbf{A}^* & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & \mathbf{W}_2 \mathbf{A}^* & -\mathbf{G} \end{bmatrix}$$

with the tag $\mathbf{H}^* := \mathbf{I}_2 \otimes \mathbf{W}_1$ where $\|\mathbf{S}\| \leq \sqrt{2(\alpha^2 + t^2N)} \leq \sqrt{2tN(\sigma_1^2 m' + t)}$. Then, the algorithm \mathcal{B} runs $\mathbf{T}' \leftarrow \text{SamplePre}(\mathbf{B}', \mathbf{S}, \mathbf{G}_{2(n+1)}, \sigma_1)$. Finally, \mathcal{B} sends $(\mathbf{A}, \mathbf{B}', \mathbf{T}', \text{aux}' := (\mathbf{W}_1, \mathbf{W}_2))$ to \mathcal{A} and returns what \mathcal{A} outputs.

To argue correctness of the reduction, first note that aux' and \mathbf{B}' are correctly generated. Further, by assumption we have $\sigma_1 \geq \delta \|\mathbf{S}\| \cdot \omega(\sqrt{N \log nN})$ and thus by Lemma 2.20, the distribution of $\text{SamplePre}(\mathbf{B}', \mathbf{S}, \mathbf{G}_{2(n+1)}, \sigma_1)$ is statistically close to $\mathbf{B}'_{\sigma_1^{-1}}(\mathbf{G}_{2(n+1)})$. Consequently, \mathcal{A} outputs a valid answer to \mathcal{B} with probability $\epsilon - \text{negl}(\lambda)$. Finally, a valid solution for StructBASIS implies a valid solution for PowerBASIS, which concludes the proof. \square

The next result focuses on the PRISIS variant. It turns out that the commutative property of the assumption allows to reduce to standard assumptions.

Lemma 3.6 (PRISIS \implies NTRU + MSIS). *Let $n > 0$ and denote $t = (n + 1)\tilde{q}$. Take $m \geq t + (n + 1) + \omega(\log \lambda)$ and $\mathfrak{s} > 2N \cdot q^{\frac{n+1}{m-t} + \frac{2}{N(m-t)}}$. Suppose $q \equiv 2N/l + 1 \pmod{4N/l}$ is prime and q^{-l} is negligible. If*

$$\sigma_0 \geq \omega(\sqrt{q \log N}) \quad \text{and} \quad \sigma_1 \geq \delta \cdot \max \left(\mathfrak{s} \sqrt{2m(m-t)}, 2\sqrt{t(\sigma_0^2 m + t)} \right) \cdot \omega(N \sqrt{\log nN}),$$

then PRISIS $_{n,m,N,q,2,\sigma_1,\beta}$ is hard under the NTRU $_{q,N,\mathcal{D}_{\sigma_0}}$ and MSIS $_{n,m,N,q,\beta}$ assumptions.

Proof. Suppose there is a PPT algorithm \mathcal{A} which wins PRISIS $_{n,m,N,q,2,\sigma,\beta}$ with probability ϵ . We revisit the PRISIS security game and introduce a series of game hops. The purpose of the hybrid argument will be to plug in the NTRU trapdoor inside the auxiliary information w . We define ϵ_i to be the probability that \mathcal{A} wins Game i .

Game 1: This is the standard PRISIS security game. To recall, the challenger samples $\mathbf{a} \leftarrow \mathcal{R}_q^m$, $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$ and sets \mathbf{A}^* as in (13). Then, it generates an invertible element $w \leftarrow \mathcal{R}_q^\times$ and computes the matrix:

$$\mathbf{B} := \begin{bmatrix} \mathbf{A}^* & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & \mathbf{W}\mathbf{A}^* & -\mathbf{G} \end{bmatrix}.$$

Then, it samples $\mathbf{T} \leftarrow \mathbf{B}_\sigma^{-1}(\mathbf{G}_{2(n+1)})$ and outputs $(\mathbf{A}, \mathbf{B}, \mathbf{T}, w)$ to the adversary \mathcal{A} . By definition, $\epsilon_1 = \epsilon$.

Game 2: In this game, we do not require the w to be invertible, and instead the challenger samples $w \leftarrow \mathcal{R}_q$ uniformly at random. Then, we have $|\epsilon_2 - \epsilon_1| \leq \frac{N}{lq^l}$ and thus $\epsilon_2 \geq \epsilon_1 - \text{negl}(\lambda)$.

Game 3: Here, the challenger obtains \mathbf{A}^* by generating $(\mathbf{A}^*, \mathbf{R}) \leftarrow \text{TrapGen}(n + 1, m)^{14}$. By Lemma 2.19, $\epsilon_3 \geq \epsilon_2 - \text{negl}(\lambda)$.

Game 4: The challenger now makes use of the trapdoor \mathbf{R} given by TrapGen. Namely, it can manually construct a $\mathbf{G}_{2(n+1)}$ -trapdoor \mathbf{R}^* for \mathbf{B} with a tag \mathbf{H}^* where

$$\mathbf{R}^* := \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad \text{and} \quad \mathbf{H}^* := \begin{bmatrix} \mathbf{I}_{n+1} & \mathbf{0} \\ \mathbf{0} & w\mathbf{I}_{n+1} \end{bmatrix}.$$

Hence, the challenge now runs $\mathbf{T} \leftarrow \text{SamplePre}(\mathbf{B}, \mathbf{R}^*, \mathbf{G}_{2(n+1)}, \sigma_1)$ and outputs the tuple $(\mathbf{A}, \mathbf{B}, \mathbf{T}, w)$. Note that with an overwhelming probability $\|\mathbf{R}^*\| \leq \mathfrak{s} \sqrt{2m(m-t)N}$. By assumption on σ_0, σ_1 and Lemma 2.20, we have $\epsilon_4 \geq \epsilon_3 - \text{negl}(\lambda)$. Note that here the challenger runs *efficiently*.

Game 5: We substitute $w \leftarrow \mathcal{R}_q$ by running $(w, \mathbf{T}_{\text{NTRU}}) \leftarrow \text{NTRU.TrapGen}(q, N)$. By Lemma 2.14 and the NTRU assumption, we get $\epsilon_5 \geq \epsilon_4 - \text{negl}(\lambda)$.

¹⁴ Thus, the challenge matrix \mathbf{A} is defined as the submatrix of \mathbf{A}^* .

Game 6: We reverse **Game 3** and **4**. That is, the challenger samples $\mathbf{T} \leftarrow \mathbf{B}_\sigma^{-1}(\mathbf{G}_{2(n+1)})$ and \mathbf{A}^* is generated as in the real PRISIS game. As before, we argue that $\varepsilon_6 \geq \varepsilon_5 - \text{negl}(\lambda)$.

Suppose there is an adversary which wins **Game₆**. We now show how to build a PRISIS trapdoor \mathbf{T} given the Module-SIS matrix \mathbf{A} and the NTRU trapdoor \mathbf{T}_{NTRU} . To this end, we will show how to find short matrices $\mathbf{S}_1, \mathbf{S}_2$ such that:

$$\mathbf{A}\mathbf{S}_1 - w\mathbf{A}\mathbf{S}_2 = \mathbf{G} .$$

Let \mathbf{g}_i be the i -th column of \mathbf{G} . Assuming that \mathbf{A}^* is full-rank¹⁵ and using linear algebra, we can find a (possibly large) vector \mathbf{t} such that $\mathbf{A}\mathbf{t} = \mathbf{g}_i$. Now, using the `NTRU.SamplePre` algorithm and the NTRU trapdoor \mathbf{T}_{NTRU} , we can sample vectors $(\mathbf{s}_{1,i}, \mathbf{s}_{2,i})$ such that:

$$\mathbf{s}_{1,i} - w\mathbf{s}_{2,i} = \mathbf{t} \text{ and } \|(\mathbf{s}_{1,i}, \mathbf{s}_{2,i})\| \leq \sigma_0\sqrt{2N}$$

with an overwhelming probability by Lemmas 2.8 and 2.15. Therefore

$$\mathbf{A}\mathbf{s}_{1,i} - w\mathbf{A}\mathbf{s}_{2,i} = \mathbf{A}(\mathbf{s}_{1,i} - w\mathbf{s}_{2,i}) = \mathbf{A}\mathbf{t} = \mathbf{g}_i .$$

Thus, we obtain the matrices $\mathbf{S}_1, \mathbf{S}_2$ by concatenation where

$$\left\| \begin{bmatrix} \mathbf{S}_1 \\ \mathbf{S}_2 \end{bmatrix} \right\| \leq \alpha := \sigma_0\sqrt{2mtN} .$$

Consequently, by Lemma 3.4, we can build a $\mathbf{G}_{2(n+1)}$ -trapdoor \mathbf{S} for \mathbf{B} such that

$$\|\mathbf{S}\| \leq \sqrt{2(\alpha^2 + t^2N)} = \sqrt{2tN(2\sigma_0^2m + t)} \leq 2\sqrt{tN(\sigma_0^2m + t)} .$$

Hence, the reduction \mathcal{B} can construct the trapdoor \mathbf{S} as above and then randomise the trapdoor for \mathbf{B} by running $\mathbf{T} \leftarrow \text{SamplePre}(\mathbf{B}, \mathbf{S}, \mathbf{G}_{2(n+1)}, \sigma_1)$. Finally it sends the tuple to \mathcal{A} and returns what it output. By Lemma 2.20, \mathcal{B} wins the Module-SIS game with probability at least $\varepsilon_6 - \text{negl}(\lambda)$, which concludes the proof. \square

3.2 Higher Dimensions

One could hope that the techniques to analyse hardness of the BASIS assumption can be translated to higher dimensions. This could be promising especially for the PRISIS assumption, which we managed to reduce to standard lattice assumptions for the $\ell = 2$ case. Unfortunately, the reduction falls flat when considering higher dimensions.

We showcase this for $\ell = 3$. Following the approach for the smaller dimension, the goal is to find short matrices $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ such that

$$\begin{aligned} \mathbf{A}^*\mathbf{S}_1 - w\mathbf{A}^*\mathbf{S}_2 &= \mathbf{Z}_1 \\ \mathbf{A}^*\mathbf{S}_2 - w\mathbf{A}^*\mathbf{S}_3 &= \mathbf{Z}_2 \end{aligned} \tag{14}$$

¹⁵ This occurs with an overwhelming probability using the analysis from [EZS+19, Appendix C] and the fact that $m - (n + 1) = \omega(\log \lambda)$.

for any $\mathbf{Z}_1, \mathbf{Z}_2$ given the NTRU trapdoor for w . If this is possible, we could set $\mathbf{Z}_1 = \mathbf{G}$ and $\mathbf{Z}_2 = \mathbf{0}$ which would give us:

$$\begin{aligned}\mathbf{A}^* \mathbf{S}_1 - w \mathbf{A}^* \mathbf{S}_2 &= \mathbf{G} \\ w \mathbf{A}^* \mathbf{S}_2 - w^2 \mathbf{A}^* \mathbf{S}_3 &= \mathbf{0}.\end{aligned}$$

Set $\mathbf{S}_4 := \mathbf{G}^{-1}(\mathbf{A}^* \mathbf{S}_1 - \mathbf{G})$. Then, we have:

$$\begin{bmatrix} \mathbf{A}^* & \mathbf{0} & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & w \mathbf{A}^* & \mathbf{0} & -\mathbf{G} \\ \mathbf{0} & \mathbf{0} & w^2 \mathbf{A}^* & -\mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{S}_1 \\ \mathbf{S}_2 \\ \mathbf{S}_3 \\ \mathbf{S}_4 \end{bmatrix} = \begin{bmatrix} \mathbf{G} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}.$$

We proceed similarly for

$$(\mathbf{Z}_1, \mathbf{Z}_2) = (-\mathbf{G}, w^{-1} \mathbf{G}) \quad \text{and} \quad (\mathbf{Z}_1, \mathbf{Z}_2) = (\mathbf{0}, -w^{-1} \mathbf{G}).$$

Thus, we managed to build a $\mathbf{G}_{3(n+1)}$ -trapdoor for \mathbf{B} . What is left to do is to produce short $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ which satisfy (14). To this end, consider the q -ary lattice

$$\Lambda = \left\{ (s_1, s_2, s_3) : \begin{bmatrix} 1 & -w & 0 \\ 0 & w & -w^2 \end{bmatrix} \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix} = \mathbf{0} \pmod{q} \right\}.$$

Suppose we can build a short basis for Λ given the NTRU trapdoor for w . Let $\mathbf{z}_{1,i}, \mathbf{z}_{2,i}$ be the i -th column of \mathbf{Z}_1 and \mathbf{Z}_2 . Now, assuming that \mathbf{A}^* is full-rank, we can find (possibly large) \mathbf{t}_1 and \mathbf{t}_2 such that $\mathbf{A}^* \mathbf{t}_j = \mathbf{z}_{j,i}$ for $j = 1, 2$. Now, using the short basis for Λ , we can sample short vectors $\mathbf{s}_{1,i}, \mathbf{s}_{2,i}, \mathbf{s}_{3,i}$ such that:

$$\begin{aligned}\mathbf{s}_{1,i} - w \mathbf{s}_{2,i} &= \mathbf{t}_1 \\ \mathbf{s}_{2,i} - w \mathbf{s}_{3,i} &= \mathbf{t}_2.\end{aligned}$$

Hence,

$$\begin{aligned}\mathbf{A}^* \mathbf{s}_{1,i} - w \mathbf{A}^* \mathbf{s}_{2,i} &= \mathbf{A}^*(\mathbf{s}_{1,i} - w \mathbf{s}_{2,i}) = \mathbf{A}^* \mathbf{t}_1 = \mathbf{z}_{1,i} \\ \mathbf{A}^* \mathbf{s}_{2,i} - w \mathbf{A}^* \mathbf{s}_{3,i} &= \mathbf{A}^*(\mathbf{s}_{2,i} - w \mathbf{s}_{3,i}) = \mathbf{A}^* \mathbf{t}_2 = \mathbf{z}_{2,i}.\end{aligned}$$

Therefore, we obtain the matrices $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ by concatenation.

Unfortunately, we are only aware of the following two bases of Λ :

$$\begin{bmatrix} w^2 & w & 1 \\ q & 0 & 0 \\ 0 & q & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} u^2 & uv & v^2 \\ \bar{u}^2 & \bar{u}\bar{v} & \bar{v}^2 \\ \bar{u}u & \bar{u}v & \bar{v}v \end{bmatrix},$$

where $\mathbf{T}_{\text{NTRU}} := ((u, v), (\bar{u}, \bar{v}))$ is the short NTRU basis. Since $\|u\|, \|v\| \approx \sqrt{q}$, the latter basis cannot have short coefficients. We leave further analysis of this approach for future work.

4 Power-BASIS Commitment Scheme

In this section we define a compressing commitment scheme which stems from the vector commitment construction of Wee and Wu [WW23]. We inherit a crucial property from the aforementioned work that we support committing to arbitrarily large ring elements. Let $\ell := d + 1$ be the length of the committed vectors over \mathcal{R}_q . Thus, the message space is $\mathcal{M} := \mathcal{R}_q^{d+1}$. We let γ, β_s be the parameters controlling the norm of various vectors. Further, we define the slack space as the vector of short polynomials:

$$\mathcal{S} := \{(c_0, \dots, c_d) : \forall i \in [0, d], c_i \in \mathcal{R}_q^\times \wedge \|c_i\|_1 \leq \beta_s\} .$$

Informally, we say that a slack is a single element $c \in \mathcal{R}_q$ if $(c, \dots, c) \in \mathcal{S}$. Finally, we define $t = n\tilde{q}$ and $\mathbf{G} := \mathbf{G}_n \in \mathcal{R}_q^{n \times t}$.

We now give intuition on the construction, and provide a formal description in Figure 4. The setup algorithm uses the `TrapGen` and `SamplePre` algorithms defined in Section 2.6. Namely, it first generates the two matrices $(\mathbf{A}, \mathbf{R}) \leftarrow \text{TrapGen}(n, m)$ along with a uniformly random invertible $\mathbf{W} \leftarrow \text{GL}(n, \mathcal{R}_q)$. Then, $\mathbf{AR} = \mathbf{G}$, where $\|\mathbf{R}\| \leq \mathfrak{s} \sqrt{2t(m-t)N}$ and $\mathfrak{s} > 2N \cdot q^{\frac{n}{m-t} + \frac{2}{N(m-t)}}$ (c.f. Lemma 2.7). Further, it computes $\mathbf{R}_i := \mathbf{RG}^{-1}(\mathbf{W}^{-i}\mathbf{G})$ for $i = 0, 1, \dots, d$. Note that

$$\mathbf{W}^i \mathbf{AR}_i = \mathbf{W}^i \mathbf{ARG}^{-1}(\mathbf{W}^{-i}\mathbf{G}) = \mathbf{W}^i \mathbf{GG}^{-1}(\mathbf{W}^{-i}\mathbf{G}) = \mathbf{G}$$

and thus \mathbf{R}_i is a \mathbf{G} -trapdoor for $\mathbf{W}^i \mathbf{A}$ and by Lemma 2.2:

$$\|\mathbf{R}_i\| \leq \|\mathbf{R}\| \cdot N\sqrt{nt} \leq \mathfrak{s}Nt\sqrt{2n(m-t)N}.$$

Then, the algorithm computes the PowerBASIS matrix along with its trapdoor:

$$\mathbf{B} := \left[\begin{array}{ccc|c} \mathbf{A} & & & -\mathbf{G} \\ & \ddots & & \vdots \\ & & \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right], \quad \tilde{\mathbf{R}} := \left[\begin{array}{ccc} \mathbf{R}_0 & & \\ & \ddots & \\ & & \mathbf{R}_d \\ \hline & & & \mathbf{0} \end{array} \right]. \quad (15)$$

Indeed, one can check that $\mathbf{B}\tilde{\mathbf{R}} = \mathbf{G}_{n(d+1)}$ and $\|\tilde{\mathbf{R}}\| \leq \mathfrak{s}Nt\sqrt{2(d+1)n(m-t)N}$. Finally, the setup algorithm re-randomises the trapdoor $\tilde{\mathbf{R}}$ by running

$$\mathbf{T} \leftarrow \text{SamplePre}(\mathbf{B}, \tilde{\mathbf{R}}, \mathbf{G}_{n(d+1)}, \sigma_0) ,$$

and thus $\mathbf{BT} = \mathbf{G}_{n(d+1)}$. Finally, the public parameters $\text{crs} := (\mathbf{A}, \mathbf{W}, \mathbf{T})$ are returned.

Suppose we want to commit to a vector (f_0, f_1, \dots, f_d) of length $d + 1$. To this end, we use crs to compute

$$\begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \leftarrow \text{SamplePre} \left(\left(\left[\begin{array}{ccc|c} \mathbf{A} & & & -\mathbf{G} \\ & \ddots & & \vdots \\ & & \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right], \begin{bmatrix} -f_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -f_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}, \mathbf{T}, \sigma_1 \right) .$$

By definition, this means that $\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_d \in \mathcal{R}_q^m$ and $\hat{\mathbf{t}} := \mathbf{G}\hat{\mathbf{t}}$ satisfy:

$$\mathbf{As}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \hat{\mathbf{t}} \quad \text{for } i = 0, 1, \dots, d . \quad (16)$$

The commitment and the decommitment state are $C := \hat{\mathbf{t}}$ and $\text{st} := (\mathbf{s}_i)_{i \in [0, d]}$.

Finally, the opening function takes the public parameters crs , the commitment $\hat{\mathbf{t}}$, a message vector $\mathbf{f} := (f_0, \dots, f_d)$, the decommitment state $(\mathbf{s}_i)_{i \in [0, d]}$ and a relaxation factor $(c_0, \dots, c_d) \in \mathcal{S}$, and accepts if and only if (16) holds and $\|c_i \mathbf{s}_i\| \leq \gamma$ for all $i = 0, 1, \dots, d$.

PowerBASIS Commitment Scheme

Setup(1^λ)

1. Sample $(\mathbf{A}, \mathbf{R}) \leftarrow \text{TrapGen}(n, m)$.
2. Sample $\mathbf{W} \leftarrow \text{GL}(n, \mathcal{R}_q)$
3. Let $\mathbf{R}_i := \mathbf{R}\mathbf{G}^{-1}(\mathbf{W}^{-i}\mathbf{G})$ for $i \in [0, d]$.
4. Set

$$\mathbf{B} := \left[\begin{array}{ccc|c} \mathbf{A} & & & -\mathbf{G} \\ & \ddots & & \vdots \\ & & \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right], \quad \tilde{\mathbf{R}} := \left[\begin{array}{ccc} \mathbf{R}_0 & & \\ & \ddots & \\ & & \mathbf{R}_d \\ \hline & & & \mathbf{0} \end{array} \right].$$

5. Sample $\mathbf{T} \leftarrow \text{SamplePre}(\mathbf{B}, \tilde{\mathbf{R}}, \mathbf{G}_{n(d+1)}, \sigma_0)$.
6. Return $\text{crs} := (\mathbf{A}, \mathbf{W}, \mathbf{T})$.

Commit($\text{crs}, \mathbf{f} \in \mathcal{R}_q^{d+1}$)

1. Parse $\mathbf{f} := (f_0, f_1, \dots, f_d)$
2. Set $\mathbf{u} := \begin{bmatrix} -f_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -f_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}$
3. Sample $\begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \leftarrow \text{SamplePre}(\mathbf{B}, \mathbf{u}, \mathbf{T}, \sigma_1)$.
4. Set $\mathbf{t} := \mathbf{G}\hat{\mathbf{t}}$.
5. Return $(C := \mathbf{t}, \text{st} := (\mathbf{s}_i)_{i \in [0, d]})$.

Open($\text{crs}, C, \mathbf{f} \in \mathcal{R}_q^{d+1}, \text{st}, \mathbf{c} \in \mathcal{R}_q^{d+1}$)

1. Parse $\mathbf{f} := (f_0, f_1, \dots, f_d)$ and $\mathbf{c} := (c_0, \dots, c_d)$.
2. Parse $C := \mathbf{t} \in \mathcal{R}_q^n$ and $\text{st} := (\mathbf{s}_i)_{i \in [0, d]}$.
3. Return 1 if and only if for all $i \in [0, d]$,
 - $\mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t}$.
 - $\|c_i \mathbf{s}_i\| \leq \gamma$.

Fig. 4: PowerBASIS commitment scheme for arbitrary messages in the message space $\mathcal{M} = \mathcal{R}_q^{d+1}$ with the slack space $\mathcal{S} := \{(c_0, \dots, c_d) : \forall i \in [0, d], c_i \in \mathcal{R}_q^\times \wedge \|c_i\|_\infty \leq \beta_s\}$. Here, $\mathbf{G} \in \mathcal{R}_q^{n \times n\bar{q}}$ is the gadget matrix of height n .

4.1 Security Analysis

In the following, we show that the PowerBASIS commitment scheme satisfies completeness, relaxed binding and hiding.

Lemma 4.1 (Completeness). *Suppose $n, N, \beta_s \geq 1$ and denote $t := n\tilde{q}$. Let $m \geq t + n + \omega(\log \lambda)$, $m' := m(d+1) + n\tilde{q}$, $n' := n\tilde{q}(d+1)$ and $t' := \max(n', m')$. Take*

$$\sigma_0 \geq \delta \mathfrak{s} N t \omega(\sqrt{2(d+1)n(m-t)N \log t'N}) \quad \text{and} \quad \sigma_1 \geq \delta \sigma_0 N \cdot \omega(\sqrt{m'n' \log t'N}) .$$

If $\gamma \geq \sigma_1 \sqrt{m'N}$ then the PowerBASIS commitment scheme satisfies completeness.

Proof. In the discussion above, we already showed that Equation (16) is true. We will show that $\|\mathbf{s}_i\| \leq \gamma$ for all i , and thus we can pick the global relaxation to be $(1, \dots, 1) \in \mathcal{S}$.

First, note that the matrix $\tilde{\mathbf{R}} \in \mathcal{R}_q^{m' \times n'}$ satisfies $\|\tilde{\mathbf{R}}\| \leq \mathfrak{s} N t \sqrt{2(d+1)n(m-t)N}$ with high probability by Lemma 2.9. Hence $\sigma_0 \geq \delta \|\tilde{\mathbf{R}}\| \cdot \omega(\sqrt{N \log t'N})$ for $t' = \max(n', m')$ and thus we can apply both Lemma 2.20 and Lemma 2.8 to deduce that with an overwhelming probability $\|\mathbf{T}\| \leq \sigma_0 \sqrt{m'n'N}$. Similarly, we have $\sigma_1 \geq \delta \|\mathbf{T}\| \cdot \omega(\sqrt{N \log t'N})$ and thus $\|\mathbf{s}_i\| \leq \sigma_1 \sqrt{m'N} \leq \gamma$ with an overwhelming probability for all $i = 0, 1, \dots, d$, which concludes the proof. \square

Based on the parameters above, we would require $\sigma_0 = \tilde{O}(\sqrt{d})$ and $\sigma_1 = \tilde{O}(d^{3/2})$, ignoring the polynomial factors related to the security parameter.

Lemma 4.2 (Relaxed Binding). *Let $t = n\tilde{q}$, $m \geq t + n + \omega(\log \lambda)$ and $n' = n\tilde{q}(d+1)$. Take $\mathfrak{s} > 2N \cdot q^{\frac{n}{m-t} + \frac{2}{N(m-t)}}$. If $\sigma_0 \geq \delta \mathfrak{s} N t \omega(\sqrt{2(d+1)n(m-t)N \log n'N})$ then PowerBASIS commitment scheme satisfies binding under the PowerBASIS $_{n-1, m, N, q, d+1, \sigma_1, 2\beta_s \gamma}$ assumption.*

Proof. Let \mathcal{A} be an adversary for the relaxed binding game which succeeds with probability ϵ . We prove the statement using an hybrid argument. We define ϵ_i to be the probability that \mathcal{A} wins Game i .

Game 0: This is the standard relaxed binding game. By definition $\epsilon_0 = \epsilon$.

Game 1: Here, we swap the SamplePre algorithm with sampling truly from a discrete Gaussian distribution. Since $\sigma_0 \geq \delta \mathfrak{s} N t \omega(\sqrt{2(d+1)n(m-t)N \log n'N})$, we can argue as in Lemma 4.1 that $\epsilon_1 \geq \epsilon_0 - \text{negl}(\lambda)$.

Game 2: In this game we do not run TrapGen any more, but instead the matrix $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$ is selected uniformly at random. By Lemma 2.7, we deduce that $\epsilon_2 \geq \epsilon_1 - \text{negl}(\lambda)$.

We claim that $\epsilon_2 = \text{negl}(\lambda)$ under the PowerBASIS assumption. First, by definition of the PowerBASIS assumption, our goal is to extract a short non-zero solution for the matrix \mathbf{A}^* , where

$$\mathbf{A} := \begin{bmatrix} \mathbf{a}^\top \\ \mathbf{A}^* \end{bmatrix} .$$

Denote the tuple \mathcal{A} outputs as:

$$\mathbf{t}, (\mathbf{f}, (\mathbf{v}_0, \dots, \mathbf{v}_d), (c_0, \dots, c_d)), (\mathbf{f}', (\mathbf{v}'_0, \dots, \mathbf{v}'_d), (c'_0, \dots, c'_d)).$$

By definition, whenever \mathcal{A} wins, it must be that openings are valid and $\mathbf{f} \neq \mathbf{f}'$, which implies there is at least an index j with $f_j \neq f'_j$. Thus, by subtracting the verification equations, we have that

$$\mathbf{A}(\mathbf{v}_j - \mathbf{v}'_j) = \begin{bmatrix} f'_j - f_j \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Since $f'_j - f_j \neq 0$, this implies that $\bar{\mathbf{v}} := (\mathbf{v}_j - \mathbf{v}'_j) \neq \mathbf{0}$. Consequently, $\mathbf{A}^* \bar{\mathbf{v}} = \mathbf{0}$. Now, $\bar{\mathbf{v}}$ might not be short. Hence, we consider $c_j c'_j \bar{\mathbf{v}}$ instead. Clearly, this is still a non-zero solution for \mathbf{A}^* since c_j, c'_j are invertible. Further,

$$\|c_j c'_j \bar{\mathbf{v}}\| \leq \|c'_j (c_j \mathbf{v})\| + \|c_j (c'_j \mathbf{v}')\| \leq 2\beta_s \gamma.$$

Therefore, $c_j c'_j \bar{\mathbf{v}}$ is a valid solution to PowerBASIS. \square

Lemma 4.3 (Hiding). *Suppose $n, N \geq 1$ and denote $t := n\tilde{q}$. Let $m \geq t + n + \omega(\log \lambda)$, $m' := m(d+1) + n\tilde{q}$, $n' := n\tilde{q}(d+1)$ and $t' := \max(n', m')$. Take*

$$\begin{aligned} \sigma_0 &\geq \delta \mathfrak{s} N t \omega(\sqrt{2(d+1)n(m-t)N \log t'N}), \\ \sigma_1 &\geq \delta \cdot \max\left(\log((d+1)mN), \sigma_0 N \cdot \omega(\sqrt{m'n' \log t'N})\right). \end{aligned}$$

Then, the PowerBASIS commitment scheme satisfies hiding.

Proof. Take an unbounded adversary \mathcal{A} which wins the hiding game with probability ϵ . We prove the statement via a sequence of games, where in each game we change the algorithm of Commit. Let ϵ_i be the advantage of the adversary against Game i .

Game 1: This is the original hiding game where Commit is defined in Figure 4. For the purpose of the proof, we assume Commit does not output \mathbf{st} . Then, by definition $\epsilon_1 = \epsilon$.

Game 2: In this game, Commit (inefficiently) samples

$$\begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \leftarrow \mathbf{B}_{\sigma_1}^{-1} \left(\begin{bmatrix} -f_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -f_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix} \right)$$

and outputs $\mathbf{t} := \mathbf{G}\hat{\mathbf{t}}$. By our assumption on σ_0, σ_1 we can argue similarly as in Lemma 4.1 to deduce that $|\epsilon_2 - \epsilon_1| = \text{negl}(\lambda)$.

Game 3: Here we make use of the fact that $\mathbf{B} := [\mathbf{E} \mid \mathbf{F}]$ where

$$\mathbf{E} := \begin{bmatrix} \mathbf{A} & & \\ & \ddots & \\ & & \mathbf{W}^d \mathbf{A} \end{bmatrix} \quad \text{and} \quad \mathbf{F} := \begin{bmatrix} -\mathbf{G} \\ \vdots \\ -\mathbf{G} \end{bmatrix}.$$

Concretely, the `Commit` algorithm first samples $\hat{\mathbf{t}} \leftarrow \mathcal{D}_{\sigma_1}^{tN}$, sets

$$\begin{bmatrix} \mathbf{t} \\ \vdots \\ \mathbf{t} \end{bmatrix} := \mathbf{F}\hat{\mathbf{t}}$$

and then generates

$$\begin{bmatrix} \mathbf{s}_1 \\ \vdots \\ \mathbf{s}_d \end{bmatrix} \leftarrow \mathbf{E}_{\sigma_1}^{-1} \left(\begin{bmatrix} -f_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -f_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix} - \begin{bmatrix} \mathbf{t} \\ \vdots \\ \mathbf{t} \end{bmatrix} \right).$$

Finally, the algorithm outputs \mathbf{t} .

By Lemma 2.6, there is a negligible function ε such that $\sigma_1 \geq \eta_\varepsilon(\Lambda^\perp(\mathbf{E}))$. Further, by [EZS+19, Appendix C] the matrix \mathbf{A} is full-rank (and so is \mathbf{E}) with probability at least

$$(1 - 1/q^{m-n+1})^{nN} \geq 1 - nN/q^{m-n+1},$$

which is overwhelming by assumption on q, n, m . Hence, we can apply Lemma 2.10 to conclude $|\epsilon_3 - \epsilon_2| = \text{negl}(\lambda)$.

Game 4: The `Commit` algorithm simply samples $\hat{\mathbf{t}} \leftarrow \mathcal{D}_{\sigma_1}^{tN}$ and outputs $\mathbf{t} := \mathbf{G}\hat{\mathbf{t}}$. Clearly, there is no difference between the outputs of `Game 3` and `4`, thus $\epsilon_4 = \epsilon_3$.

Finally, the output of `Commit` in `Game 4` does not depend on the challenge messages m_0, m_1 from \mathcal{A} . Hence, we get that $\epsilon_4 = 1/2$. By the hybrid argument we obtain $\epsilon = 1/2 + \text{negl}(\lambda)$, which concludes the proof. \square

Efficiency. The main bottleneck of the `Commit` algorithm is the trapdoor sampling procedure, which asymptotically takes $O(d^2)$ operations over \mathcal{R}_q . On the other hand, the opening algorithm makes $O(d)$ operations in \mathcal{R}_q .

5 Efficient Proofs of Polynomial Evaluation

In this section we illustrate how to prove evaluations of a polynomial that is committed using the `PowerBASIS` commitment scheme from Figure 4. We start by presenting a general framework for proving polynomial evaluations in Section 5.1, and then we describe two distinct instantiations in Sections 5.2 and 5.3. For clarity, we give an overview of frequently used parameters in Table 3. We implicitly assume that lattice dimension parameters, such as n, m, N , are $\text{poly}(\lambda)$.

5.1 Framework for Proving Evaluations

The main intuition can be described as follows. We design a relation that captures statements of the form: “the commitment \mathbf{t} has an opening $f \in \mathcal{R}_q^{d+1}$ (with respect to a given `crs`) such that $f(u) = v$, where $f \in \mathcal{R}_q^{\leq d}[\mathbf{X}]$ is now interpreted as polynomial”. The core observation is that there exists a Σ -protocol that interactively reduces an instance of that relation to a related one, in which the size of the committed polynomial is decreased. This new relation is with respect to a *different*

Parameter	Explanation
q	proof system modulus
N	degree of the cyclotomic ring $\mathcal{R} := \mathbb{Z}[X]/(X^N + 1)$
l	power-of-two such that $q \equiv 2N/l + 1 \pmod{4N/l}$
d	degree of the committed polynomial $f \in \mathcal{R}_q[X]$
n	height of the matrix \mathbf{A}
m	width of the matrix \mathbf{A}
δ	decomposition base of the gadget matrix \mathbf{G}
\tilde{q}	$\lceil \log_\delta q \rceil + 1$
n'	$n\tilde{q}(d + 1)$
m'	$m\tilde{q}(d + 1) + n\tilde{q}$
t'	$\max(n', m')$
k	folding factor of the folding protocol
h	$2h + 1$ is the number of rounds
β	initial norm of the witness openings
w	L_1 norm of elements in the challenge space \mathcal{C}
$\beta_{\mathcal{C}}$	L_∞ of elements in \mathcal{C} (used in Section 5.3)
β_h	norm of the opening vectors sent in the last round
β_s	infinity norm of the extracted relaxation factors
γ	extracted norm

Table 3: Overview of parameters and notation.

common reference string, that can be *efficiently computed* from the previous one. We then exploit this recursion to shrink to a commitment with a constant-size opening.

We formalise this discussion by introducing the opening relation below

$$\mathcal{R}_{d,\beta} := \left\{ ((\mathbf{A}, \mathbf{W}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_i)) \mid \forall i \in [0, d], \mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \wedge \|\mathbf{s}_i\| \leq \beta \right\}. \quad (17)$$

We describe the Σ -protocol, upon which our main evaluation protocol is built, in Figure 5. Roughly speaking, the prover divides the initial polynomial f of degree at most d into k polynomials g_1, \dots, g_k of degree at most $d' := (d + 1)/k - 1$ by writing

$$f(X) := \sum_{t \in [k]} X^{t-1} g_t(X^k). \quad (18)$$

Then, it “commits” to the partial polynomials by providing their evaluations at the point u , say $z_i := g_i(u^k)$. Thus, by construction

$$z = f(u) = \sum_{t \in [k]} u^{t-1} g_t(u^k) = \sum_{t=1}^k z_t u^{t-1}. \quad (19)$$

Next, the verifier outputs a challenge $(\alpha_1, \dots, \alpha_k) \leftarrow \mathcal{C} \subseteq \mathcal{R}_q^k$. Note that by considering the folded polynomial $g = \sum_{t=1}^k \alpha_t g_t$ of degree at most d' , we obtain a new polynomial evaluation statement

about g :

$$g(u^k) = \sum_{t=1}^k \alpha_t z_t . \quad (20)$$

The main strength of the PowerBASIS commitment from Figure 4 is that the commitment (resp. openings) to g can be efficiently computed from the commitment \mathbf{t} (resp. openings \mathbf{s}_i) of f given $\alpha_1, \dots, \alpha_k$ in time $O(k)$. This is the key idea for achieving succinct verification. Hence, the prover outputs the polynomial g in the clear, along with its opening vectors. The verifier eventually checks correctness of the openings with respect to the message g , as well as (19) and (20).

We first prove that this protocol transforms an instance of $\mathbf{R}_{d,\beta}$ into a smaller one of $\mathbf{R}_{d',\beta'}$.

Lemma 5.1 (Completeness). *Let $\Pi := \Sigma[d, k, \mathcal{C}, \beta]$ as in Figure 5. Then, Π is an interactive protocol with perfect completeness for $\mathbf{R}_{d,\beta}$.*

Proof. Let $(\mathbf{i}, \mathbf{x}, \mathbf{w}) = ((\mathbf{A}, \mathbf{W}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_{i \in [0,d]})) \in \mathbf{R}_{d,\beta}$. Since $f(u) = z$, the first verification check always succeeds by Equation (19). We are left to show that the new instance is valid. First, $g(u^k) = \sum_{t \in [k]} \alpha_t g_t(u^k) = \sum_{t \in [k]} \alpha_t z_t$. Further, recall that for $i \in [0, d']$ and $t \in [k]$ we have

$$\mathbf{s}_{t,i} = \mathbf{s}_{ki+t-1} \quad \text{and} \quad g_{t,i} = f_{ki+t-1} ,$$

where $g_{t,i}$ is the i -th coefficient of the polynomial g_t . Hence, the i -th coefficient of g satisfies $g_i = \sum_{t \in [k]} \alpha_t g_{t,i} = \sum_{t \in [k]} \alpha_t f_{ki+t-1}$. Therefore,

$$\begin{aligned} \mathbf{A}z_i + g_i \mathbf{e}_1 &= \mathbf{A} \left(\sum_{t \in [k]} \alpha_t \mathbf{s}_{t,i} \right) + \left(\sum_{t \in [k]} \alpha_t f_{ki+t-1} \right) \cdot \mathbf{e}_1 \\ &= \sum_{t \in [k]} \alpha_t (\mathbf{A} \mathbf{s}_{ki+t-1} + f_{ki+t-1} \mathbf{e}_1) \\ &= \sum_{t \in [k]} \alpha_t \left(\mathbf{W}^{-(ki+t-1)} \mathbf{t} \right) \\ &= \left(\sum_{t \in [k]} \alpha_t \mathbf{W}^{-(ki+t-1)} \right) \cdot \mathbf{t} \\ &= (\mathbf{W}^k)^{-i} \left(\sum_{t \in [k]} \alpha_t \mathbf{W}^{-(t-1)} \right) \cdot \mathbf{t}. \end{aligned}$$

Finally, by Lemma 2.1 for $\alpha \in \mathcal{C}$, $\|\mathbf{z}_i\| \leq \sum_{t \in [k]} \|\alpha_t \mathbf{s}_{t,i}\| \leq \sum_{t \in [k]} \|\alpha_t\|_1 \cdot \beta \leq w \beta$ where $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$. This shows that the new instance is in $\mathbf{R}_{d',\beta'}$, and thus the verifier accepts. \square

We now apply the Σ -protocol recursively h times, reducing the final opening size to $(d+1)/k^h$, while increasing the final norm for verification by a factor w^h .

Construction 5.2. Let k, h be integers, and let $\mathcal{C} \subseteq \mathcal{R}_q^k$. We let $\text{Eval}[d, k, h, \mathcal{C}, \beta] := (\mathcal{P}, \mathcal{V})$ be the protocol that we describe in Figure 6.

Completeness of the protocol is easily shown by applying Lemma 5.1 h times.

Σ -Protocol for $\mathcal{R}_{d,\beta}$

Prover

$$\sum_{t \in [k]} X^{t-1} g_t(X^k) =: f(X)$$

$$z_t := g_t(u^k) \text{ for } t \in [k]$$

$$\xrightarrow{(z_t)_{t \in [k]}}$$

Verifier

$$\alpha \leftarrow \mathcal{C} \subseteq \mathcal{R}_q^k$$

$$\xleftarrow{\alpha}$$

$$g := \sum_{t \in [k]} \alpha_t g_t$$

$$\mathbf{z}_i := \sum_{t \in [k]} \alpha_t \mathbf{s}_{t,i} \text{ for } i \in [0, d']$$

$$\xrightarrow{g, (\mathbf{z}_i)_{i \in [0, d']}}$$

$$\beta' := w \beta$$

$$\mathbf{t}' := \left(\sum_{t \in [k]} \alpha_t \mathbf{W}^{-(t-1)} \right) \cdot \mathbf{t}$$

$$\mathbf{i}' := (\mathbf{A}, \mathbf{W}^k)$$

$$\mathbf{x}' := \left(\mathbf{t}', u^k, \sum_{t \in [k]} \alpha_t z_t \right)$$

$$\mathbf{w}' := (g, (\mathbf{z}_i)_{i \in [0, d']})$$

Check:

$$z = \sum_{t \in [k]} u^{t-1} z_t$$

$$(\mathbf{i}', \mathbf{x}', \mathbf{w}') \in \mathcal{R}_{d', \beta'}$$

Fig. 5: The Σ -protocol $\Sigma[d, k, \mathcal{C}, \beta]$ for relation $\mathcal{R}_{d,\beta}$ in Equation (17). Here, $d' := (d + 1)/k - 1$, $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$ and $\mathbf{s}_{t,i} := \mathbf{s}_{ki+t-1}$ for $i \in [0, d']$ and $t \in [k]$.

Interactive Protocol for $R_{d,\beta}$

$\mathcal{P}((\mathbf{A}, \mathbf{W}), (\mathbf{t}, u_0, z_0), (f_0, (\mathbf{s}_{0,i})_{i \in [0,d]}))$

1. Set $d_0 := d$.
2. For $r \in [h]$:
 - (a) Set $d_r := (d_{r-1} + 1)/k - 1$.
 - (b) Write $f_{r-1}(\mathbf{X}) := \sum_{t \in [k]} \mathbf{X}^{t-1} f_{r-1,t}(\mathbf{X}^k)$ for $f_{r-1,1}, \dots, f_{r-1,k} \in \mathcal{R}_q^{\leq d_r}[\mathbf{X}]$.
 - (c) Set $z_{r-1,t} := f_{r-1,t}(u_{r-1}^k)$ for $t \in [k]$.
 - (d) Send $(z_{r-1,t})_{t \in [k]}$ to the verifier.
 - (e) Receive α_r from the verifier.
 - (f) Compute $f_r := \sum_{t \in [k]} \alpha_{r,t} f_{r-1,t}$.
 - (g) Compute $\mathbf{s}_{r,i} := \sum_{t \in [k]} \alpha_{r,t} \mathbf{s}_{r-1,ki+t-1}$ for $i \in [0, d_r]$.
 - (h) Compute $u_r := u_{r-1}^k$.
3. Send $(f_h, (\mathbf{s}_{h,i})_{i \in [0,d_h]})$ to the verifier.

$\mathcal{V}((\mathbf{A}, \mathbf{W}_0), (\mathbf{t}_0, u_0, z_0))$

1. $\beta_0 := \beta$.
2. For $r \in [h]$:
 - (a) Receive $(z_{r-1,t})_{t \in [k]}$ from the prover.
 - (b) Check $z_{r-1} = \sum_{t \in [k]} u_{r-1}^{t-1} z_{r-1,t}$.
 - (c) Sample $\alpha_r \leftarrow \mathcal{C}$ and send it to the prover.
 - (d) Set $\mathbf{W}_r := \mathbf{W}_{r-1}^k$.
 - (e) Set $\mathbf{t}_r := \left(\sum_{t \in [k]} \alpha_{r,t} \mathbf{W}_{r-1}^{-(t-1)} \right) \cdot \mathbf{t}_{r-1}$.
 - (f) Set $\beta_r := w \cdot \beta_{r-1}$.
 - (g) Set $u_r := u_{r-1}^k$.
 - (h) Set $z_r := \sum_{t \in [k]} \alpha_{r,t} z_{r-1,t}$.
3. Receive $(f_h, (\mathbf{s}_{h,i})_{i \in [0,d_h]})$ from the prover.
4. Check:
 - (a) $f_h(u_h) = z_h$.
 - (b) $\mathbf{A} \mathbf{s}_{h,i} + f_{h,i} \mathbf{e}_1 = \mathbf{W}_h^{-i} \mathbf{t}_h$ for $i \in [0, d_h]$.
 - (c) $\|\mathbf{s}_{h,i}\| \leq \beta_h$ for $i \in [0, d_h]$.

Fig. 6: The protocol $\text{Eval}[d, k, h, \mathcal{C}, \beta]$ for $R_{d,\beta}$. As before, we denote $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$.

Lemma 5.3 (Completeness). *Let $\Pi := \text{Eval}[d, k, h, \mathcal{C}, \beta]$. Then, Π is an interactive protocol with perfect completeness for $\mathcal{R}_{d, \beta}$.*

Proof. Denote by $(\mathbf{i}_r, \mathbf{x}_r, \mathbf{w}_r) := ((\mathbf{A}, \mathbf{W}_r), (\mathbf{t}_r, u_r, z_r), (f_r, (\mathbf{s}_{r,i})_{i \in [d_r]}))$ for $r \in [h]$. By Lemma 5.1, $(\mathbf{i}_r, \mathbf{x}_r, \mathbf{w}_r) \in \mathcal{R}_{d_r, \beta_r}$ implies $(\mathbf{i}_{r+1}, \mathbf{x}_{r+1}, \mathbf{w}_{r+1}) \in \mathcal{R}_{d_{r+1}, \beta_{r+1}}$ with probability 1. Since $(\mathbf{i}_0, \mathbf{x}_0, \mathbf{w}_0) \in \mathcal{R}_{d, \beta_0}$, then $(\mathbf{i}_h, \mathbf{x}_h, \mathbf{w}_h) \in \mathcal{R}_{d_h, \beta_h}$, and thus the verifier final checks accept. \square

Remark 5.4. The protocol that we have described has folding factor k constant across every round of interaction. In fact, we can gain more flexibility by allowing each round to use a different folding factor. This can be beneficial, for example, to obtain a constant polynomial in the last round of the protocol when the original degree is not a h -power.

We analyse the communication complexity of $\text{Eval}[d, k, h, \mathcal{C}, \beta]$ in the next lemma.

Lemma 5.5 (Efficiency). *The total communication complexity of $\text{Eval}[d, k, h, \mathcal{C}, \beta]$ (in bits) can be bounded by*

$$h \cdot (kN \lceil \log q \rceil + \lceil \log |\mathcal{C}| \rceil) + \frac{d+1}{k^h} \left(N \lceil \log q \rceil + mN \lceil \log(2^w \beta) \rceil \right) .$$

Further, the prover makes $O(md)$ operations in \mathcal{R}_q while the verifier makes $O((n+m)^2(hk + d/k^h))$ operations in \mathcal{R}_q .

Proof. In each round the prover sends k elements of \mathcal{R}_q to the verifier, and the verifier sends 1 element of \mathcal{C} . In the final round, the prover sends a polynomial with $d_h = (d+1)/k^h$ coefficients, and $d_h + 1$ opening vectors, each of which has norm at most β_h .

We turn to the prover complexity and first consider Step 2. Every r -th round out of $[h]$, the prover makes $O(mkd_r) = O(md_{r-1})$ operations in \mathcal{R}_q . Since $d_0 = O(d)$ and in general $d_r = O(d/k^r)$, the total runtime of the prover can be bounded by

$$O \left(\sum_{r=0}^{h-1} md_r \right) = O \left(m \sum_{r=0}^{h-1} d/k^r \right) = O \left(md \cdot \frac{1 - 1/k^h}{1 - 1/k} \right) = O(md) .$$

We move to the verifier analysis. In Step 2, for every round $r \in [h]$, the verifier makes at most $O(kn^2)$ operations. Hence, the total cost of Step 2 is $O(hkn^2)$. The rest of the algorithm takes $O(d_h(nm + n^2))$ steps. Thus, the total runtime can be bounded by $O((n+m)^2(hk + d/k^h))$ ring operations. \square

Next, we provide two instantiations of the protocol in Figure 6 which will differ in the selection of the challenge space \mathcal{C} . This has direct impact on the knowledge extraction strategy.

5.2 Monomial Protocol

In the following, we describe a so-called *monomial* variant of the protocol, where the name comes from the description of the challenge space \mathcal{C} . Fix $k := 2$, and $\mathcal{C} := \{1\} \times \{X^i : i \in \mathbb{Z}\}$. Note that by definition $w = 2$, and $\alpha, \alpha' \in \mathcal{C}$ with $\alpha \neq \alpha'$ implies that $\alpha_2 - \alpha'_2 \in \mathcal{R}_q^\times$. In this section, we also assume that $2 \in \mathcal{R}_q^\times$ (which can be enforced if $\gcd(2, q) = 1$).

We aim to show that $\Pi := \text{Eval}[d, 2, h, \mathcal{C}, \beta]$ is 2-special sound. In fact, we will not be able to show this *exactly*, as the extraction will introduce some slack. Rather we show that Π is special sound for the *relaxed* opening relation that we describe next:

$$\tilde{\mathcal{R}}_{d,c,\gamma} := \left\{ ((\mathbf{A}, \mathbf{W}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_{i \in [0,d]})) \mid \begin{array}{l} \forall i \in [0, d], \mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \wedge \\ \wedge c \in \mathcal{R}_q^\times \wedge \|c \cdot \mathbf{s}_i\| \leq \gamma \\ \wedge f(u) = z \end{array} \right\}. \quad (21)$$

We will directly show that Eval is special-sound, which also implies special-soundness of the Σ -protocol by noting that the two protocols are equivalent when $h = 1$. To argue soundness we will first prove that there exists an extractor that is able to extract witnesses of the higher layer of the transcript tree from the children.

Lemma 5.6 (Special Soundness for Σ). *Let $c \in \mathcal{R}_q^\times$, and let $\mathfrak{i} = (\mathbf{A}, \mathbf{W})$, $\mathfrak{x} = (\mathbf{t}, u, z)$. There exists an algorithm that, given two transcripts tr_j of the following form*

$$\text{tr}_j := ((z_1, z_2), \alpha_j := (1, \alpha_j) \in \mathcal{C}, \mathfrak{w}'_j := (g_j, (\mathbf{z}_{j,i})_i)) \quad \text{for } j = 0, 1$$

where $\alpha_0 \neq \alpha_1$, outputs $\mathfrak{w} := (\bar{f}, (\bar{\mathbf{s}}_i)_i)$. Furthermore, let $d', \mathfrak{i}', \mathfrak{x}'_0, \mathfrak{x}'_1$ be obtained as in Figure 5. If, for $i \in \{0, 1\}$, $(\mathfrak{i}', \mathfrak{x}'_i, \mathfrak{w}'_i) \in \tilde{\mathcal{R}}_{d',c,\beta}$, and $z = z_1 + uz_2$, then $(\mathfrak{i}, \mathfrak{x}, \mathfrak{w}) \in \tilde{\mathcal{R}}_{d,2c,\gamma}$ where $\gamma := 2N\beta$.

Proof. Consider the following algorithm:

$\mathcal{E}(\text{tr}_0, \text{tr}_1)$:

1. Set $\bar{\mathbf{s}}_{2i} := \frac{\alpha_1 \mathbf{z}_{0,i} - \alpha_0 \mathbf{z}_{1,i}}{\alpha_1 - \alpha_0}$, $\bar{\mathbf{s}}_{2i+1} := \frac{\mathbf{z}_{0,i} - \mathbf{z}_{1,i}}{\alpha_0 - \alpha_1}$ for $i \in [0, (d-1)/2]$.
2. Set $\bar{f}_1 := \frac{\alpha_1 g_0 - \alpha_0 g_1}{\alpha_1 - \alpha_0}$, $\bar{f}_2 := \frac{g_0 - g_1}{\alpha_0 - \alpha_1}$.
3. Set $\bar{f} := f_1(\mathbf{X}^2) + \mathbf{X} \bar{f}_2(\mathbf{X}^2)$.
4. Return $\bar{f}, (\bar{\mathbf{s}}_i)_{i \in [0,d]}$.

Let now $(\bar{f}, (\bar{\mathbf{s}}_i)_i) \leftarrow \mathcal{E}(\text{tr})$. Note that

$$\begin{aligned} \mathbf{A} \bar{\mathbf{s}}_{2i} + \bar{f}_{2i} \mathbf{e}_1 &= \mathbf{W}^{-2i} \mathbf{t} \\ \mathbf{A} \bar{\mathbf{s}}_{2i+1} + \bar{f}_{2i+1} \mathbf{e}_1 &= \mathbf{W}^{-(2i+1)} \mathbf{t} . \end{aligned}$$

Now, we have that:

$$\begin{aligned} \bar{f}(u) &= \bar{f}_1(u^2) + u \bar{f}_2(u^2) \\ &= \frac{\alpha_1 g_0(u^2) - \alpha_0 g_1(u^2)}{\alpha_1 - \alpha_0} + u \frac{g_0(u^2) - g_1(u^2)}{\alpha_0 - \alpha_1} \\ &= z_1 + uz_2 \\ &= z . \end{aligned}$$

Finally, we set $c^* := 2c$. First, note that $c^* \in \mathcal{R}_q^\times$ since $2 \in \mathcal{R}_q^\times$. Now, for $i \in [0, d']$, we have:

$$\begin{aligned} \|c^* \cdot \bar{\mathbf{s}}_{2i}\| &= \left\| \frac{2}{\alpha_1 - \alpha_0} \cdot c \cdot (\alpha_1 \mathbf{z}_{0,i} - \alpha_0 \mathbf{z}_{1,i}) \right\| \\ &\leq \left\| \frac{2}{\alpha_1 - \alpha_0} \right\|_\infty \|c(\alpha_1 \mathbf{z}_{0,i} - \alpha_0 \mathbf{z}_{1,i})\|_1 \end{aligned}$$

$$\begin{aligned}
&= \|c(\alpha_1 \mathbf{z}_{0,i} - \alpha_0 \mathbf{z}_{1,i})\|_1 \\
&\leq \sqrt{N}(\|\alpha_1 c \mathbf{z}_{0,i}\| + \|\alpha_0 c \mathbf{z}_{1,i}\|) \\
&\leq N(\|\alpha_1\| \cdot \|c \mathbf{z}_{0,i}\| + \|\alpha_0\| \cdot \|c \mathbf{z}_{1,i}\|) \\
&\leq 2N\beta = \gamma
\end{aligned}$$

where the second equality follows by Lemma 2.3 and the last inequality by $\|\alpha\| = 1$ for $(1, \alpha) \in \mathcal{C}$. Similarly, $\|c^* \cdot \bar{\mathbf{s}}_{2i+1}\| \leq \gamma$. \square

Using this extractor, we show that Π is $(2, \dots, 2)$ -special sound. The new extractor will start from the leaves of the tree of transcripts, applying the extractor described in Lemma 5.6 to obtain witnesses¹⁶ for the upper layer.

Lemma 5.7 (Special Soundness for Eval). *Let $\mathcal{C} := \{1\} \times \{X^i : i \in \mathbb{Z}\}$ and let $\Pi := \text{Eval}[d, 2, h, \mathcal{C}, \beta]$ be as in Construction 5.2. Set $\gamma := (2N)^h \cdot \beta_h$. Then Π is a special sound proof system for $\tilde{\mathbf{R}}_{d, 2^h, \gamma}$.*

Proof. Let tr be a tree of transcripts, which we index as follows.

- $\alpha_{(r,j)}$ for $(r, j) \in [h] \times [2^r]$ is the j -th challenge in the r -th layer of the transcript.
- $(z_{(r,j),1}, z_{(r,j),2})$ for $(r, j) \in [0, h-1] \times [2^r]$ is the j -th response in the r -th layer of the transcript.
- $(\bar{f}_{(h,j)}, (\bar{\mathbf{s}}_{(h,j),i})_i)$ for $j \in [2^h]$ is the final message sent by the prover.

We introduce the following notation as in the verifier algorithm:

- $d_0 := d$, $d_r := d_{r-1}/2$ for $r \in [h]$
- $\mathbf{W}_0 := \mathbf{W}$, $\mathbf{W}_r := \mathbf{W}_{r-1}^2$ for $r \in [h]$.
- $\mathbf{t}_{(0,1)} := \mathbf{t}$, $\mathbf{t}_{(r,2j-1)} := (1 + \alpha_{(r,2j-1)} \mathbf{W}_{r-1}^{-1}) \mathbf{t}_{(r-1,j)}$, $\mathbf{t}_{(r,2j)} := (1 + \alpha_{(r,2j)} \mathbf{W}_{r-1}^{-1}) \mathbf{t}_{(r-1,j)}$ for $(r, j) \in [h] \times [2^r]$.
- $\beta_0 := \beta$, $\beta_r := 2N \cdot \beta_{r-1}$ for $r \in [h]$.
- $u_0 := u$, $u_r := u_{r-1}^2$ for $r \in [h]$.
- $z_{(r,2j-1)} := z_{(r-1,j),1} + \alpha_{(r,2j-1)} z_{(r-1,j),2}$, $z_{(r,2j)} := z_{(r-1,j),1} + \alpha_{(r,2j)} z_{(r-1,j),2}$ for $(r, j) \in [h] \times [2^r - 1]$.

Denote with $\mathcal{E}^{(1)}$ the extractor of Lemma 5.6.

$\mathcal{E}(\text{tr})$:

1. Set $d_0 := d$, $d_r := d_{r-1}/2$ for $r \in [h]$.
2. For $r := h, \dots, 1$:
 - (a) Set, for $j \in [2^{r-1}]$,

$$\text{tr}_{(r-1,j)} := \left((z_{(r-1,j),1}, z_{(r-1,j),2}), \begin{matrix} \alpha_{(r,2j-1)}, (\bar{f}_{(r,2j-1)}, (\bar{\mathbf{s}}_{(r,2j-1),i})_i) \\ \alpha_{(r,2j)}, (\bar{f}_{(r,2j)}, (\bar{\mathbf{s}}_{(r,2j),i})_i) \end{matrix} \right).$$

- (b) Compute $\bar{f}_{(r-1,j)}, (\bar{\mathbf{s}}_{(r-1,j),i})_{i \in [0, d_{r-1}]} \leftarrow \mathcal{E}^{(1)}(\text{tr}_{(r-1,j)})$ for $j \in [2^{r-1}]$
3. Return $\bar{f}_{(0,1)}, (\bar{\mathbf{s}}_{(0,1),i})_{i \in [d]}$.

¹⁶ We also implicitly collect the corresponding relaxation factors, which are *the same* across the same layer.

We prove that this extractor yields a valid witness by induction on r . First note that, by the verifier checks, for $(r, j) \in [h] \times [2^r]$

$$z_{(r-1,j)} = z_{(r-1,j),1} + u_{r-1} z_{(r-1,j),2} .$$

Write $\mathbf{i}_{(r,j)} := (\mathbf{A}, \mathbf{W}_r)$, $\mathbf{x}_{(r,j)} := (\mathbf{t}_{(r,j)}, u_{(r,j)}, z_{(r,j)})$, $\mathbf{w}_{(r,j)} := (\bar{f}_{(r,j)}, (\bar{\mathbf{s}}_{(r,j),i})_i)$ for $(r, j) \in [h] \times [2^r]$. For $r = h$, since the transcripts are accepting, $(\mathbf{i}_{(h,j)}, \mathbf{x}_{(h,j)}, \mathbf{w}_{(h,j)}) \in \mathbf{R}_{d_h, \beta_h} = \tilde{\mathbf{R}}_{d_h, 1, \beta_h}$ for $j \in [2^h]$. Thus, by Lemma 5.6, $(\mathbf{i}_{(h-1,j)}, \mathbf{x}_{(h-1,j)}, \mathbf{w}_{(h-1,j)}) \in \tilde{\mathbf{R}}_{d_{h-1}, 2, 2N\beta_h}$.

We can continue with the induction, and this yields that for the extracted witness $\mathbf{w}_{(0,1)} := (\bar{f}_{(0,1)}, (\bar{\mathbf{s}}_{(0,1),i})_{i \in [d]})$ we have that:

$$(\mathbf{i}_{(0,1)}, \mathbf{x}_{(0,1)}, \mathbf{w}_{(0,1)}) \in \tilde{\mathbf{R}}_{d, 2^h, (2N)^h \beta_h} .$$

Setting $\gamma := (2N)^h \beta_h$, and noting that $2^h \in \mathcal{R}_q^\times$, this concludes our proof. \square

We can use **Eval** to construct a polynomial commitment scheme. We detail the construction in Theorem 5.8 and summarise the parameters and efficiency features in Table 4.

Parameters	Instantiation
m	$\geq n(1 + \tilde{q}) + \omega(\log \lambda)$
δ	$q^{1/O(1)}$
\mathfrak{s}	$> 2Nq^{\frac{n}{m-n\tilde{q}} + \frac{2}{N(m-n\tilde{q})}}$
σ_0	$\geq \delta \mathfrak{s} N n \tilde{q} \cdot \omega(\sqrt{2(d+1)n(m-n\tilde{q})N \log t'N})$
σ_1	$\geq \delta \sigma_0 N \cdot \omega(\sqrt{m'n' \log t'N})$
β	$\geq \sigma_1 \sqrt{m'N}$
k	2
\mathcal{C}	$\{1\} \times \{X^i : i \in \mathbb{Z}\}$
\mathbf{w}	2
β_h	$\mathbf{w}^h \cdot \beta$
γ	$(2N)^h \cdot \beta_h$
β_s	2^h
Soundness	$\left(\frac{h}{2N}\right)^\ell$
Commitment size	$nN \log q$
Communication complexity	$\ell \cdot (h(2N \log q + \log N + 1) + \frac{d+1}{2^h} (N \log q + mN \log \beta_h))$
Prover time	$O(\ell \cdot md)$
Verifier time	$O(\ell \cdot (n+m)^2 \cdot (2h + d/2^h))$

Table 4: Parameters for the interactive polynomial commitment scheme obtained from Figure 4 and running the ℓ -parallel repetition of **Eval** $[d, 2, h, \mathcal{C}, \beta]$ for proofs of evaluation. We compute the prover and verifier runtime in terms of operations in \mathcal{R}_q .

Theorem 5.8. *Let $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \mathcal{P}^t, \mathcal{V}^t)$ where **Setup**, **Commit**, **Open** are as in Figure 4 and $\mathcal{P}^t, \mathcal{V}^t$ are the t -parallel repetitions of the prover and verifier of **Eval**. Then **PC** is an interactive polynomial commitment scheme with the efficiency properties and parameters shown in*

Table 4. In particular, when $h = O(\log d)$ and $t > \frac{\lambda}{\log N + 1 - \log h}$ we obtain an interactive polynomial commitment scheme with negligible knowledge soundness error, polylogarithmic communication complexity, and polylogarithmic verifier time.

Proof. Completeness and relaxed binding follow from Lemmas 4.1 and 4.2. Perfect evaluation completeness follows from Lemma 5.1. For evaluation knowledge soundness, we apply [AF22, Theorem 4] to Lemma 5.7. Communication complexity follows from Lemma 5.5. Additionally, claims about the prover and verifier runtime hold by Lemma 5.5 and the fact that both $\log q$ and N are polynomial in λ . \square

5.3 Large Sampling Set

We present a second instantiation which allows us to obtain negligible knowledge soundness error *without parallel repetition*, using coordinate-wise special soundness (c.f. Section 2.10) and a large challenge space. We let $t, k \in \mathbb{N}$. Fix also $\beta_C > 0$. Recall that $S_\kappa := \{\alpha \in \mathcal{R}_q : \|\alpha\|_\infty \leq \kappa\}$. We define the challenge space and the slack space as

$$\mathcal{C} := S_{\beta_C}^k \quad \text{and} \quad \mathcal{S}_t := \left\{ \prod_{i \in [t]} \alpha_i - \alpha'_i : \alpha_i, \alpha'_i \in S_{\beta_C}, \alpha_i \neq \alpha'_i \right\}.$$

Note that $|\mathcal{C}| = (2\beta_C + 1)^{kN}$ and $w \leq \beta_C kN$. We also let $\beta_{s,t} := \max_{c \in \mathcal{S}_t} \|c\|_\infty$. Note that, for $c \in \mathcal{S}_t$,

$$\|c\|_\infty \leq \left\| \prod_i (\alpha_i - \alpha'_i) \right\|_\infty \leq \|\alpha_1 - \alpha'_1\|_\infty \cdot \prod_{i \neq 1} \|\alpha_i - \alpha'_i\|_1 \leq 2\beta_C \cdot (2\beta_C N)^{t-1},$$

and thus $\|c\|_1 \leq (2\beta_C N)^t$.

We show a simple invertibility result that will be useful in the proof of soundness.

Lemma 5.9. *Let $1 \leq l < N$ be a power of two, and suppose that $q \equiv 2N/l + 1 \pmod{4N/l}$. If $2\beta_C < \sqrt{l/N} q^{l/N}$, then for any $t \geq 1$, $\mathcal{S}_t \subseteq \mathcal{R}_q^\times$.*

Proof. Let $\alpha \neq \alpha' \in S_{\beta_C}$. Then, $\alpha - \alpha' \neq 0$, and $\|\alpha - \alpha'\|_\infty \leq 2\beta_C$. Thus, by Lemma 2.4, $\alpha - \alpha' \in \mathcal{R}_q^\times$. Elements of \mathcal{S}_t are products of elements of that form, and since the product of invertible elements is itself invertible, the result follows. \square

We will assume thereafter that we are in the regime in which Lemma 2.4 holds (as in Table 3).

We again aim to show that $\text{Eval}[d, k, h, \mathcal{C}, \beta]$ is knowledge sound. As before, we define an opening relation, which will differ from Equation (21) in that the relaxation factors will not be the same across openings, but rather will be included as part of the witness. This will reflect the fact that the extracted opening will have different slack derived from the challenges.

$$\tilde{\mathcal{R}}_{d,\beta,t} := \left\{ ((\mathbf{A}, \mathbf{W}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_i, (c_i)_i)) \left| \begin{array}{l} \forall i \in [0, d], \mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \wedge \\ \wedge c_i \in \mathcal{S}_t \wedge \|c_i \cdot \mathbf{s}_i\| \leq \beta \\ \wedge f(u) = z \end{array} \right. \right\}. \quad (22)$$

As before, to argue that the protocol is knowledge sound, we will first show an extractor to be used to move between layers of the transcript tree. In this case however, we will argue using coordinate-wise special-soundness instead of special soundness.

Lemma 5.10 (Coordinate-Wise Special Soundness for Σ). Let $c \in \mathcal{R}_q^\times$, and let $\mathbf{i} = (\mathbf{A}, \mathbf{W})$, $\mathbf{x} = (\mathbf{t}, u, z)$. There exists an algorithm that, given $k+1$ transcripts $(\text{tr}_j)_{j \in [0, k]}$ of the following form:

$$\text{tr}_j := \begin{pmatrix} (z_1, \dots, z_k) \\ \boldsymbol{\alpha}_j \\ (g_j, (\mathbf{s}_{j,i})_{i \in [0, d']}) \end{pmatrix} \text{ with } (\boldsymbol{\alpha}_j)_j \in \text{SS}(S_{\beta_C}, k) ,$$

and slack $(c_{j,i})_{j,i}$ outputs $\mathbf{w} := (\bar{f}, (\bar{\mathbf{s}}_i)_i, (\bar{c}_i)_i)$. Furthermore, let \mathbf{i}' , $(\mathbf{x}'_j)_{j \in [k]}$ be obtained as in Figure 5 (where \mathbf{x}'_j is obtained from the j -th leaf of the transcript) and $\mathbf{w}'_j := (g_j, (\mathbf{s}_{j,i})_i, (c_{j,i})_i)$. If, for $i \in [0, k]$, $(\mathbf{i}', \mathbf{x}'_i, \mathbf{w}'_i) \in \tilde{\mathbf{R}}_{d', \beta, t}$, and $z = \sum_{t \in [k]} u^{t-1} z_t$, then $(\mathbf{i}, \mathbf{x}, \mathbf{w}) \in \tilde{\mathbf{R}}_{d, \gamma, 2t+1}$ where $\gamma := 2\beta$ if $t = 0$ and $\gamma := 2N\beta_{s,t}\beta$ otherwise.

Proof. Assume, without loss of generality, that the transcripts are arranged so that, for $j \in [k]$, $\boldsymbol{\alpha}_0 \equiv_j \boldsymbol{\alpha}_j$. We thus can write $\boldsymbol{\alpha}_0 = (\alpha_1, \dots, \alpha_k)$ and $\boldsymbol{\alpha}_j := (\alpha_1, \dots, \alpha'_j, \dots, \alpha_k)$ with $\alpha_j \neq \alpha'_j$. Consider the extractor

$\mathcal{E}(\text{tr} = (\text{tr}_0, \dots, \text{tr}_k), (\tilde{c}_{j,i})_{j,i})$:

1. For $j \in [k]$:
 - (a) Set $\bar{f}_j := \frac{g_0 - g_j}{\alpha_j - \alpha'_j}$.
 - (b) For $i \in [0, d']$:
 - i. Set $\bar{\mathbf{s}}_{ki+j-1} := \frac{\mathbf{z}_{0,i} - \mathbf{z}_{j,i}}{\alpha_j - \alpha'_j}$.
 - ii. Set $\bar{c}_{ki+j-1} := (\alpha_j - \alpha'_j) c_{0,i} c_{j,i}$.
2. Set $\bar{f} := \sum_{j \in [k]} \mathbf{X}^{j-1} \bar{f}_j(\mathbf{X}^k)$.
3. Return $(\bar{f}, (\bar{\mathbf{s}}_i)_{i \in [0, d]}, (\bar{c}_i)_{i \in [0, d]})$.

Since the transcript is accepting, for $j \in [0, k]$, $i \in [0, d']$ we have that

$$\mathbf{A} \mathbf{z}_{j,i} + g_{j,i} \mathbf{e}_1 = (\mathbf{W}^k)^{-i} \left(\sum_{t \in [k]} \alpha_{j,t} \mathbf{W}^{t-1} \right) \mathbf{t} .$$

Subtracting the equation for $j = 0$ from the equation for $j \in [k]$ yields that, for $i \in [0, d']$:

$$\mathbf{A} \left(\frac{\mathbf{z}_{0,i} - \mathbf{z}_{j,i}}{\alpha_j - \alpha'_j} \right) + \left(\frac{g_{0,i} - g_{j,i}}{\alpha_j - \alpha'_j} \right) \mathbf{e}_1 = \mathbf{W}^{-(ki+j-1)} \mathbf{t} .$$

To show that the extracted \bar{f} evaluates to z at u , note that:

$$\begin{aligned} \bar{f}(u) &= \sum_{j \in [k]} u^{j-1} \bar{f}_j(u^k) \\ &= \sum_{j \in [k]} u^{j-1} \frac{g_0(u^k) - g_j(u^k)}{\alpha_j - \alpha'_j} \\ &= \sum_{j \in [k]} u^{j-1} \frac{\sum_{t \in [k]} (\alpha_{0,t} - \alpha_{j,t}) z_t}{\alpha_j - \alpha'_j} \\ &= \sum_{j \in [k]} u^{j-1} z_j = z . \end{aligned}$$

Where in the third equality we have used that the verifier check accepts, and for the fourth $\sum_{t \in [k]} (\alpha_{0,t} - \alpha_{j,t}) z_t = (\alpha_j - \alpha'_j) z_j$. We argue that the extracted $\bar{\mathbf{s}}_i$ are (relaxed) short.

$$\begin{aligned} \|\bar{c}_{ki+j-1} \cdot \bar{\mathbf{s}}_{ki+j-1}\| &= \left\| (\alpha_j - \alpha'_j) c_{0,i} c_{j,i} \frac{\mathbf{z}_{0,i} - \mathbf{z}_{j,i}}{\alpha_j - \alpha'_j} \right\| \\ &= \|c_{0,i} c_{j,i} (\mathbf{z}_{0,i} - \mathbf{z}_{j,i})\| \\ &\leq \|c_{j,i} c_{0,i} \mathbf{z}_{0,i}\| + \|c_{0,i} c_{j,i} \mathbf{z}_{j,i}\| \\ &\leq \sqrt{N} \beta_{s,t} (\|c_{0,i} \mathbf{z}_{0,i}\|_1 + \|c_{j,i} \mathbf{z}_{j,i}\|_1) \\ &\leq 2N \beta_{s,t} \beta = \gamma . \end{aligned}$$

If $t = 0$, then the slacks must have been 1, and thus $\|\bar{c}_{ki+j-1} \bar{\mathbf{s}}_{ki+j-1}\| \leq \|\mathbf{z}_{0,i} - \mathbf{z}_{j,i}\| \leq 2\beta$ as desired. Finally, what is left to show is that the new slack is in the prescribed slack space. This is easy to show as the previous two slacks are a product of t differences of challenges, that we then multiply with a new difference, leading to a product of $2t + 1$ differences of challenges. Lemma 5.9 guarantees that this new slack is invertible as long as β_C is small enough. \square

We then use this extractor recursively to show that Eval is coordinate-wise special sound.

Lemma 5.11 (Coordinate-Wise Special Soundness for Eval). *Let $k, h \in \mathbb{N}$, $\beta_C > 0$. Let $\Pi := \text{Eval}[d, k, h, C, \beta]$ be as in Construction 5.2. Then, Π is a k -coordinate-wise special-sound proof system for the relation $\bar{\mathbf{R}}_{d, \gamma, t}$ where*

$$\begin{aligned} \gamma &:= 2^h \cdot (2\beta_C N)^{2^h - h - 1} \cdot \beta_h \\ t &:= 2^h - 1 . \end{aligned}$$

Proof. We index the transcript as in Lemma 5.7. Denote by $\mathcal{E}^{(1)}$ the extractor of Lemma 5.10. Consider the new extractor

$\mathcal{E}(\text{tr})$:

1. Set $\bar{c}_{(h,j)} = 1$ for $j \in [(k+1)^h]$.
2. For $r := h, \dots, 1$:
 - (a) Set for $j \in [(k+1)^{r-1}]$:

$$\text{tr}_{(r-1,j)} := \begin{pmatrix} (\boldsymbol{\alpha}_{(r,(j-1)(k+1)+1)}, (\bar{f}_{(r,(j-1)(k+1)+1)}, (\bar{\mathbf{s}}_{(r,(j-1)(k+1)+1}, i) i))) \\ (z_{(r-1,j),t})_{t \in [k]} \\ \vdots \\ (\boldsymbol{\alpha}_{(r,j(k+1))}, (\bar{f}_{(r,j(k+1))}, (\bar{\mathbf{s}}_{(r,j(k+1))}, i) i)) \end{pmatrix} .$$

- (b) Compute $(\bar{f}_{(r-1,j)}, (\bar{\mathbf{s}}_{(r-1,j),i} i), (\bar{c}_{(r-1,j),i} i)) \leftarrow \mathcal{E}^{(1)}(\text{tr}_{(r-1,j)}, (\bar{c}_{(r,(j-1)(k+1)+t}, i)_{t,i}))$.

3. Return $(\bar{f}_{(0,1)}, (\bar{\mathbf{s}}_{(0,1),t}), (\bar{c}_{(0,1),t})_t$.

We argue that the extractor yields a valid witness inductively. We again note that for $(r, j) \in [h] \times [(k+1)^r]$, since the transcripts are accepting,

$$z_{(r-1,j)} = \sum_{t \in [k]} u_{r-1}^{k-1} z_{(r-1,j),t} .$$

Write $\mathbf{i}_{(r,j)} := (\mathbf{A}, \mathbf{W}_r)$, $\mathbf{x}_{(r,j)} := (\mathbf{t}_{(r,j)}, u_r, (z_{(r,j),i} i))$ and $\mathbf{w}_{(r,j)} := (\bar{f}_{(r,j)}, (\bar{\mathbf{s}}_{(r,j),i} i), (\bar{c}_{(r,j),i} i))$. Since the leaves are accepting (and the relaxed relation is equivalent to the exact one when the relaxation

factors are one), $(\mathfrak{i}_{(h,j)}, \mathfrak{x}_{(h,j)}, \mathfrak{w}_{(h,j)}) \in \tilde{\mathcal{R}}_{d_h, \beta_h, 0}$. Thus, Lemma 5.10 (in the case $t = 0$) implies that $(\mathfrak{i}_{(h-1,j)}, \mathfrak{x}_{(h-1,j)}, \mathfrak{w}_{(h-1,j)}) \in \tilde{\mathcal{R}}_{d_{h-1}, 2\beta_h, 1}$. Now, we define the recurrence relations:

$$t_r := \begin{cases} 1 & \text{if } r = 1 \\ 2t_{r-1} + 1 & \text{otherwise} \end{cases} \quad \text{and} \quad \gamma_r := \begin{cases} 2\beta & \text{if } r = 1 \\ 2N\beta_{s, t_{r-1}}\gamma_{r-1} & \text{otherwise} \end{cases} .$$

Lemma 5.10 implies exactly that, if $(\mathfrak{i}_{(r,j)}, \mathfrak{x}_{(r,j)}, \mathfrak{w}_{(r,j)}) \in \tilde{\mathcal{R}}_{d_{r-i}, \gamma_r, t_r}$, then the extracted witness $(\mathfrak{i}_{(r+1,j)}, \mathfrak{x}_{(r+1,j)}, \mathfrak{w}_{(r+1,j)}) \in \tilde{\mathcal{R}}_{d_{k-r-1}, \gamma_{r+1}, t_{r+1}}$. Unfolding the recurrence relations, we note that $t_r = 2^r - 1$ and

$$\begin{aligned} \gamma_r &= 2^r N^{r-1} \left(\prod_{i=1}^{r-1} \beta_{s, t_i} \right) \beta_h \\ &\leq 2^r N^{r-1} \left(\prod_{i=1}^{r-1} 2\beta_{\mathcal{C}} (2\beta_{\mathcal{C}} N)^{2^i - 2} \right) \beta_h \\ &= 2^r N^{r-1} (2\beta_{\mathcal{C}})^{r-1} (2\beta_{\mathcal{C}} N)^{\sum_{i=1}^{r-1} 2^i - 2} \cdot \beta_h \\ &= 2^j N^{r-1} (2\beta_{\mathcal{C}})^{r-1} (2\beta_{\mathcal{C}} N)^{2^r - 2r} \cdot \beta_h \\ &= 2^r (2\beta_{\mathcal{C}} N)^{2^r - r - 1} \cdot \beta_h \end{aligned}$$

Taking this to its natural conclusion:

$$(\mathfrak{i}_{(0,1)}, \mathfrak{x}_{(0,1)}, \mathfrak{w}_{(0,1)}) \in \tilde{\mathcal{R}}_{d, \gamma_h, t_h} ,$$

and setting $\gamma := \gamma_h$, $t := t_h$ implies the result. \square

Again, we can use Eval to construct a polynomial commitment scheme.

Theorem 5.12. *Let $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Eval}, \text{Verify})$ where Setup, Commit, Open are as in Figure 4 and Eval, Verify are obtained by applying the Fiat-Shamir transform to Eval $[d, k, h, \mathcal{C}, \beta]$ when $k^h = \text{poly}(d)$. Then, PC is an polynomial commitment scheme with the efficiency properties and parameters shown in Table 4.*

Proof. Completeness and relaxed binding follow from Lemmas 4.1 and 4.2. Perfect evaluation completeness follows from Lemma 5.1. Communication complexity and runtimes follow from Lemma 5.5. Knowledge soundness follows from Lemma 2.34 and Lemma 5.11, noting that when $k^h = \text{poly}(d)$ and thus the extractor runs in expected polynomial time. \square

At this point, one might be tempted to instantiate the scheme in Theorem 5.12 with $h = O(\log d)$ and $k = O(1)$ to obtain a protocol with logarithmic communication complexity as in Theorem 5.8 and small soundness error. This unfortunately does not succeed, as the extracted norm in this case grows $\exp(d)$ and thus $\log q \geq \text{poly}(d)$. The resulting protocol will communicate logarithmically many elements of \mathcal{R}_q , but the overall communication complexity will thus be polynomial in d . Thus, h must be at most $O(\log \log d)$. In fact, let $0 < \epsilon < 1$ be a constant and set $h = 1/\epsilon = O(1)$, $k = d^\epsilon$. It is easy to see from Table 5 that then the communication complexity will be $O(d^{1/\epsilon})$ elements of \mathcal{R}_q and we can set $\log q = \text{polylog}(d)$ to obtain overall sublinear communication complexity.

Parameters	Instantiation
m	$\geq n(1 + \tilde{q}) + \omega(\log \lambda)$
δ	$q^{1/O(1)}$
\mathfrak{s}	$> 2Nq^{\frac{n}{m-n\tilde{q}} + \frac{2}{N(m-n\tilde{q})}}$
σ_0	$\geq \delta \mathfrak{s} N n \tilde{q} \cdot \omega(\sqrt{2(d+1)n(m-n\tilde{q})N \log t'N})$
σ_1	$\geq \delta \sigma_0 N \cdot \omega(\sqrt{m'n' \log t'N})$
β	$\geq \sigma_1 \sqrt{m'N}$
\mathcal{C}	$S_{\beta_{\mathcal{C}}}^k$
$\beta_{\mathcal{C}}$	$< \frac{1}{2} \sqrt{l/N} q^{l/N}$
w	$kN\beta_{\mathcal{C}}$
β_h	$w^h \cdot \beta$
γ	$2^h \cdot (2\beta_{\mathcal{C}}N)^{2^h-h-1} \cdot \beta_h$
β_s	$(2\beta_{\mathcal{C}}N)^{2^h-1}$
Soundness	$\frac{(Q+1) \cdot hk}{(2\beta_{\mathcal{C}}+1)^N}$
Commitment size	$nN \log q$
Proof size	$h(kN \log q) + \frac{d+1}{k^h}(N \log q + mN \log \beta_h)$
Prover time	$O(md)$
Verifier time	$O((n+m)^2 \cdot (hk + d/k^h))$

Table 5: Parameters for the polynomial commitment scheme obtained from Figure 4 and the Fiat-Shamir transform of $\text{Eval}[d, k, h, \mathcal{C}, \beta]$ for proofs of evaluation. We let Q be an upper bound on the number of queries an adversary can make to the random oracle.

Accordingly, the verifier time will also be sublinear. In fact, we can further improve on this. Set now $h \approx \log \log d$, and $k \approx d^{1/\log \log d}$. It can be easily verified that in this case we obtain

$$\log q = O\left(\frac{\log^2 d}{\log \log d}\right),$$

and in terms of communication complexity: $O((\log \log d) \cdot d^{1/\log \log d})$ elements of \mathcal{R}_q or $\text{polylog}(d) \cdot d^{1/\log \log d}$ bits (similarly for the verifier complexity). As such, we can conclude that Theorem 5.12 gives rise to a *quasi-polylogarithmic* non-interactive polynomial commitment scheme from lattice assumptions.

5.4 Batching Evaluations

5.4.1 Multiple Evaluations at a Single Point

We show a simple approach to amortise the cost of proving evaluations of multiple evaluations at a single point. More concretely, we have a list of (committed) polynomials f_1, \dots, f_r and want to show that $f_i(u) = z_i$. First we define the corresponding relation, namely:

$$\mathbb{R}_{d,\beta}^r := \left\{ (\mathbf{A}, \mathbf{W}), ((\mathbf{t}_j)_j, u, (z_j)_j), ((f_j)_j, (\mathbf{s}_{j,i})_{j,i}) \mid \left((\mathbf{A}, \mathbf{W}), (\mathbf{t}_j, u, z_j), (f_j, (\mathbf{s}_{j,i})_{j,i}) \right) \in \mathbb{R}_{d,\beta} \right\} \quad \forall j \in [r].$$

The intuition of the protocol that we design is to take a random linear combinations of the polynomials f_1, \dots, f_r , and prove that its evaluation at u is equal to the linear combination of the claimed evaluations. The protocol that we describe in Figure 7 takes this idea and combines it with one round of Figure 5, which is useful for better concrete efficiency.

Lemma 5.13 (Completeness). *Let $\Pi := \text{multiEval}[d, r, k, \mathcal{C}, \beta]$ be the protocol in Figure 7. Then, Π is a Σ -protocol with perfect completeness for $\mathbb{R}_{d,\beta}^r$.*

Proof. It is easy to see that $g(u^k) = \sum_{\iota,t} \alpha_{\iota,t} g_{\iota,t}(u^k) = \sum \alpha_{\iota,t} z_{\iota,t}$. Also, for $i \in [0, d']$,

$$\begin{aligned} \mathbf{A}\mathbf{z}_i + g_i \mathbf{e}_1 &= \sum_{\iota,t \in [r] \times [k]} \alpha_{\iota,t} (\mathbf{A}\mathbf{s}_{\iota,t,i} + g_{\iota,t,i} \mathbf{e}_1) \\ &= \sum_{\iota,t \in [r] \times [k]} \alpha_{\iota,t} (\mathbf{A}\mathbf{s}_{\iota,ki+t-1} + g_{\iota,ki+t-1} \mathbf{e}_1) \\ &= \sum_{\iota,t \in [r] \times [k]} \alpha_{\iota,t} \mathbf{W}^{-(ki+t-1)} \mathbf{t}_\iota. \end{aligned}$$

Finally, $\|\mathbf{z}_i\| = \left\| \sum_{\iota,t} \alpha_{\iota,t} \mathbf{s}_{\iota,t,i} \right\| \leq w \beta = \beta'$ as desired. \square

As before, we define a relaxed opening relation (we use the definition of $\tilde{\mathbb{R}}$ from Equation (22)):

$$\tilde{\mathbb{R}}_{d,\beta,t}^r := \left\{ \left((\mathbf{A}, \mathbf{W}), ((\mathbf{t}_\iota)_\iota, u, (z_\iota)_\iota), ((f_\iota)_\iota, (\mathbf{s}_{\iota,i})_{\iota,i}, (c_{\iota,i})_{\iota,i}) \right) \mid \left((\mathbf{A}, \mathbf{W}), (\mathbf{t}_\iota, u, z_\iota), (f_\iota, (\mathbf{s}_{\iota,i})_{\iota,i}, (c_{\iota,i})_{\iota,i}) \right) \in \tilde{\mathbb{R}}_{d,\beta,t} \right\} \quad \forall \iota \in [r].$$

We now prove coordinate-wise special soundness for the set $\mathcal{C} := S_{\beta\mathcal{C}}^{rk} \subseteq \mathcal{R}_q^{rk}$, where each element has rk coordinates. Then, it is easy to show (e.g. using the composition results as in [BS22, Section 3]) that composing multiEval with Eval yields a knowledge sound protocol for this relaxed relation.

Proving Multiple Evaluations at a Single Point

Prover

$$\sum_{t \in [k]} X^{t-1} g_{\iota,t}(X^k) =: f_{\iota}(X) \text{ for } \iota \in [r]$$

$$z_{\iota,t} := g_{\iota,t}(u^k) \text{ for } (\iota, t) \in [r] \times [k]$$

$$\xrightarrow{(z_{\iota,t})_{(\iota,t) \in [r] \times [k]}}$$

Verifier

$$\alpha = (\alpha_1, \dots, \alpha_r) \leftarrow \mathcal{C} := S_{\beta \mathcal{C}}^{rk}$$

$$\xleftarrow{\alpha}$$

$$g := \sum_{(\iota,t) \in [r] \times [k]} \alpha_{\iota,t} g_{\iota,t}$$

$$\mathbf{z}_i := \sum_{(\iota,t) \in [r] \times [k]} \alpha_{\iota,t} \mathbf{s}_{\iota,t,i} \text{ for } i \in [0, d']$$

$$\xrightarrow{g, (\mathbf{z}_i)_{i \in [0, d']}}$$

$$\beta' := w \beta$$

$$\mathbf{t}' := \left(\sum_{(\iota,t) \in [r] \times [k]} \alpha_{\iota,t} \mathbf{W}^{-(t-1)} \cdot \mathbf{t}_{\iota} \right)$$

$$\mathbf{i}' := (\mathbf{A}, \mathbf{W}^k)$$

$$\mathbf{x}' := \left(\mathbf{t}', u^k, \sum_{(\iota,t) \in [r] \times [k]} \alpha_{\iota,t} z_{\iota,t} \right)$$

$$\mathbf{w}' := (g, (\mathbf{z}_i)_{i \in [0, d']})$$

Check:

$$z_{\iota} = \sum_{t \in [k]} u^{t-1} z_{\iota,t} \text{ for } \iota \in [r]$$

$$(\mathbf{i}', \mathbf{x}', \mathbf{w}') \in \mathbf{R}_{d', \beta'}$$

Fig. 7: The protocol $\text{multiEval}[d, r, k, \mathcal{C}, \beta]$ for proving evaluations of r polynomials at a single point. In the above $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$. As before, we define $d' := (d+1)/k - 1$ and $\mathbf{s}_{\iota,t,i} := \mathbf{s}_{\iota, ki+t-1}$ for $\iota \in [r]$.

Lemma 5.14 (Coordinate-Wise Special Soundness). *Let $\Pi := \text{multiEval}[d, r, k, \mathcal{C}, \beta]$ be the protocol in Figure 7. Let $\mathbf{i} := (\mathbf{A}, \mathbf{W})$, $\mathbf{x} := ((\mathbf{t}_\iota)_\iota, u, (z_\iota)_\iota)$. There exists an algorithm that, given $rk + 1$ transcripts $(\text{tr}_j)_{j \in [0, rk]}$ of the following form:*

$$\text{tr}_j := \begin{pmatrix} (z_{\iota,t})_{\iota,t} \\ \boldsymbol{\alpha}_j \\ (g_j, (\mathbf{z}_{j,i})_{i \in [0, d]}) \end{pmatrix} \text{ with } (\boldsymbol{\alpha}_j)_j \in \text{SS}(S_{\beta_{\mathcal{C}}}, rk) ,$$

and relaxation factors $(c_{j,i})_{j,i}$, outputs $\mathbf{w} := ((\bar{f}_\iota)_\iota, (\bar{\mathbf{s}}_{\iota,i})_{\iota,i}, (\bar{c}_{\iota,i})_{\iota,i})$. Now, set $\mathbf{i}' := (\mathbf{A}, \mathbf{W}^k)$, $\mathbf{x}_j := (\sum_{\iota,t} \alpha_{j,\iota,t} \mathbf{t}_\iota, u^k, \sum_{\iota,t} \alpha_{j,\iota,t} z_{\iota,t})$, $\mathbf{w}_j := (g_j, (\mathbf{z}_{j,i})_i, (c_{j,i})_i)$. If for $j \in [0, r]$, $(\mathbf{i}', \mathbf{x}_j, \mathbf{w}_j) \in \tilde{\mathbf{R}}_{d,\beta,t}$, and $z_\iota = \sum_{t \in [k]} u^{t-1} z_{\iota,t}$ for $\iota \in [r]$, then $(\mathbf{i}, \mathbf{x}, \mathbf{w}) \in \tilde{\mathbf{R}}_{d,\gamma,t'}^r$ where $\gamma := 2N\beta_{s,t}\beta$, $t' := 2t + 1$.

Proof. Again, assume without loss of generality that $\boldsymbol{\alpha}_0 \equiv_j \boldsymbol{\alpha}_j$ for $j \in [rk]$. Now, reindex $\boldsymbol{\alpha}_1 \dots, \boldsymbol{\alpha}_{rk}$ into a $r \times k$ matrix $\boldsymbol{\alpha}_{1,1}, \dots, \boldsymbol{\alpha}_{r,k}$. We write $\boldsymbol{\alpha}_0 = (\alpha_{1,1}^*, \dots, \alpha_{r,k}^*)$ and thus assume that $\boldsymbol{\alpha}_{v,w} = (\alpha_{1,1}^*, \dots, \alpha'_{v,w}, \dots, \alpha_{r,k}^*)$ with $\alpha'_{v,w} \neq \alpha_{v,w}^*$. We also reindex $(g_j)_j, (\mathbf{z}_{j,i})$ accordingly so that $g_{v,w}$ corresponds the $\boldsymbol{\alpha}_{v,w}$ challenge (note that we skip the 0-th challenge $\boldsymbol{\alpha}_0$).

With these conventions, we let the extractor be the following.

$\mathcal{E}(\text{tr})$:

1. For $\iota \in [r], t \in [k]$:

(a) Let $\bar{f}_{\iota,t} := \frac{g_0 - g_{\iota,t}}{\alpha_{\iota,t}^* - \alpha'_{\iota,t}}$.

(b) Let $\bar{\mathbf{s}}_{\iota,ki+t-1} := \frac{\mathbf{z}_{0,i} - \mathbf{z}_{\iota,t,i}}{\alpha_{\iota,t}^* - \alpha'_{\iota,t}}$ for $i \in [0, d']$.

(c) Let $\bar{c}_{\iota,ki+t-1} := (\alpha_{\iota,t}^* - \alpha'_{\iota,t}) c_{0,i} c_{\iota,t,i}$ for $i \in [0, d']$.

2. Set $\bar{f}_\iota := \sum_{t \in [k]} \mathbf{X}^{t-1} \bar{f}_{\iota,t}$ for $\iota \in [r]$.

3. Return $(\bar{f}_\iota)_\iota, ((\bar{\mathbf{s}}_{\iota,i})_i)_\iota, ((\bar{c}_{\iota,i})_i)_\iota$.

First note that by assumption, $g_0(u^k) = \sum_{\iota,t} \alpha_{\iota,t}^* z_{\iota,t}$ and $g_{v,w}(u^k) = \alpha'_{v,w} z_{v,w} + \sum_{(\iota,t) \neq (v,w)} \alpha_{\iota,t}^* z_{\iota,t}$.

Thus, $\bar{f}_{v,w}(u^k) = \frac{g_0 - g_{v,w}}{\alpha_{v,w}^* - \alpha'_{v,w}}(u^k) = z_{v,w}$. Thus, for $\iota \in [r]$:

$$\bar{f}_\iota(u) = \sum_{t \in [k]} u^{t-1} \bar{f}_{\iota,t}(u^k) = \sum_{t \in [k]} u^{t-1} \frac{g_0 - g_{\iota,t}}{\alpha_{\iota,t}^* - \alpha'_{\iota,t}}(u^k) = \sum_{t \in [k]} u^{t-1} z_{\iota,t} = z_\iota .$$

Now, also by assumption:

$$\begin{aligned} \mathbf{A} \mathbf{z}_{0,i} + g_{0,i} \mathbf{e}_1 &= \mathbf{W}^{-i} \left(\sum_{(\iota,t)} \alpha_{\iota,t}^* \mathbf{t}_\iota \right) \\ \mathbf{A} \mathbf{z}_{v,w,i} + g_{v,w,i} \mathbf{e}_1 &= \mathbf{W}^{-i} \left(\alpha'_{v,w} \mathbf{t}_v + \sum_{(\iota,t) \neq (v,w)} \alpha_{\iota,t}^* \mathbf{t}_\iota \right) \\ &\Downarrow \\ \mathbf{A} \left(\frac{\mathbf{z}_{0,i} - \mathbf{z}_{v,w,i}}{\alpha_{v,w}^* - \alpha'_{v,w}} \right) + \left(\frac{g_{0,i} - g_{v,w,i}}{\alpha_{v,w}^* - \alpha'_{v,w}} \right) \cdot \mathbf{e}_1 &= \mathbf{W}^{-(ki+w-1)} \mathbf{t}_v \\ &\Downarrow \\ \mathbf{A} \bar{\mathbf{s}}_{v,ki+w-1} + \bar{f}_{v,ki+w-1} \mathbf{e}_1 &= \mathbf{W}^{-(ki+w-1)} \mathbf{t}_v . \end{aligned}$$

Finally, note that $\|\bar{c}_{\iota,i} \bar{\mathbf{s}}_{\iota,i}\| \leq 2N\beta_{s,t}\beta$ by exactly the same reasoning as in Lemma 5.10. \square

5.4.2 Multiple Evaluations at Distinct Points

Next, we consider the dual problem, namely amortising proving many statements of the form $f_\iota(u_\iota) = z_\iota$ for $\iota \in [r]$ where u_1, \dots, u_r can be potentially distinct. Looking at Lemma 5.5, a large part of the communication complexity is represented by the last round, where the prover has to send openings $\mathbf{s}_0, \dots, \mathbf{s}_{d_h}$. We amortise this by taking a random linear combination of these openings. As before, for concrete efficiency reasons, we integrate this within a round of compression.

The relation that we consider is the following:

$$\mathbb{R}_{d,\beta}^r := \left\{ \left(\begin{array}{l} (\mathbf{A}, \mathbf{W}), \\ (\mathbf{t}_\iota, u_\iota, z_\iota)_\iota \\ (f_\iota, \mathbf{s}_{\iota,i})_{\iota,i} \end{array} \right) \mid \left((\mathbf{A}, \mathbf{W}), (\mathbf{t}_\iota, u_\iota, z_\iota), (f_\iota, \mathbf{s}_{\iota,i})_{\iota,i} \right) \in \mathbb{R}_{d,\beta} \right\} \quad \forall \iota \in [r]$$

The protocol is then described in Figure 8. Now, we show evalMulti has perfect completeness..

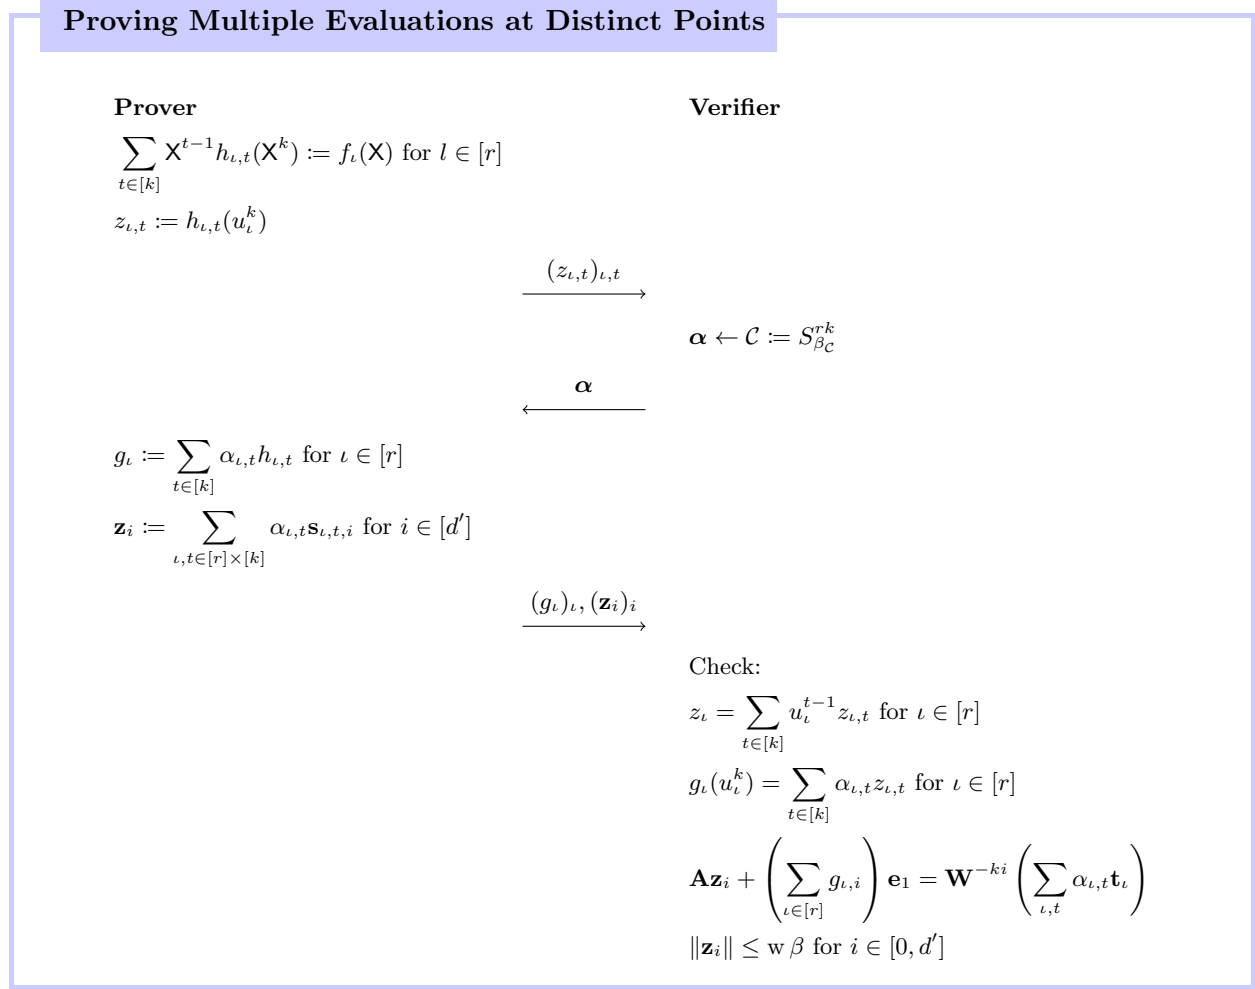


Fig. 8: The protocol $\text{evalMulti}[d, r, k, \mathcal{C}, \beta]$ for proving evaluations of multiple polynomials at multiple points. In the above $w := \max_{\alpha \in \mathcal{C}} \|\alpha\|_1$ and $d' := (d + 1)/k - 1$.

Lemma 5.15 (Completeness). *Let $\Pi := \text{evalMulti}[d, r, k, \mathcal{C}, \beta]$. Then Π is a Σ -protocol with perfect completeness for $\mathbb{R}_{d,\beta}^r$.*

Proof. For the first verifier check,

$$z_\iota = f_\iota(u_\iota) = \sum_{t \in [k]} u_\iota^{t-1} h_{\iota,t}(u_\iota^k) = \sum_{t \in [k]} u_\iota^{t-1} z_{\iota,t} .$$

Next, we check that g_ι evaluates to the correct value.

$$g_\iota(u_\iota^k) = \sum_{t \in [k]} \alpha_{\iota,t} h_{\iota,t}(u_\iota^k) = \sum_{t \in [k]} \alpha_{\iota,t} z_{\iota,t} .$$

Checking validity of the openings is similarly straightforward:

$$\begin{aligned} \mathbf{A} \mathbf{z}_i + \left(\sum_{\iota} g_{\iota,i} \right) \mathbf{e}_1 &= \mathbf{A} \left(\sum_{\iota,t} \alpha_{\iota,t} \mathbf{s}_{\iota,t,i} \right) + \left(\sum_{\iota,t} \alpha_{\iota,t} h_{\iota,t,i} \right) \mathbf{e}_1 \\ &= \sum_{\iota,t} \alpha_{\iota,t} (\mathbf{A} \mathbf{s}_{\iota,t,i} + h_{\iota,t,i} \mathbf{e}_1) \\ &= \sum_{\iota,t} \alpha_{\iota,t} (\mathbf{A} \mathbf{s}_{\iota,ki+t-1} + f_{\iota,ki+t-1} \mathbf{e}_1) \\ &= \sum_{\iota,t} \alpha_{\iota,t} (\mathbf{W}^{-(ki+t-1)} \mathbf{t}_\iota) \\ &= (\mathbf{W}^k)^{-i} \cdot \left(\sum_{\iota,t} \alpha_{\iota,t} \mathbf{W}^{-(t-1)} \mathbf{t}_\iota \right) . \end{aligned}$$

Finally, $\|\mathbf{z}_i\| = \left\| \sum_{\iota,t} \alpha_{\iota,t} \mathbf{s}_{\iota,t,i} \right\| \leq w \beta$. □

For knowledge soundness, we again define a relaxed opening relation, namely:

$$\tilde{\mathbb{R}}_{d,\beta}^r := \left\{ \left(\begin{array}{c} (\mathbf{A}, \mathbf{W}), \\ (\mathbf{t}_\iota, u_\iota, z_\iota)_\iota \\ (f_\iota, \mathbf{s}_{\iota,i}, c_{\iota,i})_{\iota,i} \end{array} \right) \middle| \left((\mathbf{A}, \mathbf{W}), (\mathbf{t}_\iota, u_\iota, z_\iota), (f_\iota, \mathbf{s}_{\iota,i}, c_{\iota,i})_{\iota,i} \right) \in \tilde{\mathbb{R}}_{d,\beta,1} \right\} .$$

Lemma 5.16 (Coordinate-Wise Special Soundness). *Let $\Pi := \text{multiEval}[d, r, k, \mathcal{C}, \beta]$ be the protocol in Figure 7. Then, Π is a rk -coordinate-wise knowledge sound proof system for $\tilde{\mathbb{R}}_{d,2\beta}^r$.*

Proof. For $j \in [0, rk]$, consider transcripts of the following form:

$$\text{tr}_j := \left(\begin{array}{c} (z_{\iota,t})_{\iota,t} \\ \alpha_j \\ ((g_{j,t})_\iota, (\mathbf{z}_{j,i})_i) \end{array} \right) \text{ with } (\alpha_j)_j \in \text{SS}(S_{\beta_C}, rk) ,$$

and again assume, without loss of generality, that the transcripts are arranged so that, for $j \in [r]$, $\alpha_0 \equiv_j \alpha_j$. Reindex and arrange the challenges as in Section 5.4.1.

Consider the following extractor:

$\mathcal{E}(\text{tr}_0, \dots, \text{tr}_{rk})$:

1. For $\iota \in [r], t \in [k]$:
 - (a) Set $\bar{f}_{\iota,t} := \frac{g_0 - g_{\iota,t}}{\alpha_{\iota,t}^* - \alpha'_{\iota,t}}$.
 - (b) Set $\bar{\mathbf{s}}_{\iota,ki+t-1} := \frac{\mathbf{z}_{0,i} - \mathbf{z}_{\iota,t,i}}{\alpha_{\iota,t}^* - \alpha'_{\iota,t}}$ for $i \in [0, d']$.
 - (c) Set $\bar{c}_{\iota,ki+t-1} := \alpha_{\iota,t}^* - \alpha'_{\iota,t}$ for $i \in [0, d']$.
2. Set $\bar{f}_\iota := \sum_{t \in [k]} \mathbf{X}^{t-1} \bar{f}_{\iota,t}$ for $\iota \in [r]$.
3. Return $(\bar{f}_\iota)_\iota, (\bar{\mathbf{s}}_{\iota,i})_{\iota,i}, (\bar{c}_{\iota,i})_{\iota,i}$.

Since the transcripts are accepting, we have that $z_\iota = \sum_{t \in [k]} u_\iota^{t-1} z_{\iota,t}$ for $\iota \in [r]$. Also, $g_{0,\iota}(u_\iota^k) = \sum_{t \in [k]} \alpha_{\iota,t}^* z_{\iota,t}$ and $g_{v,w,\iota}(u_\iota^k) = \alpha'_{v,w} z_{v,w} + \sum_{t \neq w} \alpha_{\iota,t}^* z_{\iota,t}$. Thus, $\frac{g_{0,\iota} - g_{v,w}}{\alpha_{\iota,t}^* - \alpha'_{v,w}}(u_\iota^k) = z_{v,w}$. Now,

$$\bar{f}_\iota(u_\iota) = \sum_{t \in [k]} u_\iota^{t-1} \bar{f}_{\iota,t}(u_\iota^k) = \sum_{t \in [k]} u_\iota^{t-1} \frac{g_0 - g_{\iota,t}}{\alpha_{\iota,t}^* - \alpha'_{\iota,t}}(u_\iota^k) = \sum_{t \in [k]} u_\iota^{t-1} z_{\iota,t} = z_\iota .$$

We also have that

$$\begin{aligned} \mathbf{A} \mathbf{z}_{0,i} + \left(\sum_{\iota} g_{0,\iota,i} \right) \mathbf{e}_1 &= \mathbf{W}^{-ki} \left(\sum_{\iota,t} \alpha_{\iota,t}^* \mathbf{W}^{-(t-1)} \mathbf{t}_\iota \right) \\ \mathbf{A} \mathbf{z}_{v,w,i} + \left(\sum_{\iota} g_{v,w,\iota,i} \right) \mathbf{e}_1 &= \mathbf{W}^{-ki} \left(\alpha'_{v,w} \mathbf{W}^{-(w-1)} \mathbf{t}_v + \sum_{\iota, t \neq (v,w)} \alpha_{\iota,t}^* \mathbf{W}^{-(t-1)} \mathbf{t}_\iota \right) \\ &\Downarrow \\ \mathbf{A} \left(\frac{\mathbf{z}_{0,i} - \mathbf{z}_{v,w,i}}{\alpha_{v,w}^* - \alpha'_{v,w}} \right) + \bar{f}_{v,w,i} \mathbf{e}_1 &= \mathbf{W}^{-ki} \left(\mathbf{W}^{-(w-1)} \mathbf{t}_v \right) \\ &\Downarrow \\ \mathbf{A} \bar{\mathbf{s}}_{v,ki+w-1} + \bar{f}_{v,ki+w-1} \mathbf{e}_1 &= \mathbf{W}^{-(ki+w-1)} \mathbf{t}_v . \end{aligned}$$

Finally, $\|\bar{c}_{\iota,ki+t-1} \bar{\mathbf{s}}_{\iota,ki+t-1}\| \leq \|\mathbf{z}_{0,i}\| + \|\mathbf{z}_{\iota,t,i}\| \leq 2\beta$ as desired. \square

We can combine these two newly presented protocols with `Eval` to obtain a protocol for multiple evaluations. Let $u_1, \dots, u_r \in \mathcal{R}_q$, and suppose we want to show that $f_{\iota,m}(u_\iota) = z_{\iota,m}$ for $\iota \in [r], m \in [r_\iota]$ for committed polynomials $(f_{\iota,m})_{\iota,m}$. Write $w_s := \max_{\alpha \leftarrow S_{\beta_C}^s} \|\alpha\|_1$. The combined protocol runs (in parallel) `multiEval` $[d, r_\iota, k, S_{\beta_C}^{r_\iota \cdot k}, \beta]$ with input $(f_{\iota,m})_{m \in [r_\iota]}$ for $\iota \in [r]$. This outputs r claims, which we handle by running `Eval` $[d/k, k, S_{\beta_C}^k, w_{r_\iota, k} \cdot \beta]$ r -times into parallel. Finally, we run a single instance of `multiEval` $[d/k^{h+1}, r, k, S_{\beta_C}^{r \cdot k}, (\max_\iota w_{r_\iota, k}) \cdot w_k^h \beta]$. The final complexity of this protocol is summarised in Table 6.

5.5 Honest-Verifier Zero-Knowledge

We provide a linear-sized Σ -protocol for the relation $\mathbf{R}_{d,\beta}$ (c.f. Equation (17)) which satisfies honest-verifier zero-knowledge. Combined with the recursive methodology described above, we can achieve zero-knowledge succinct proofs of polynomial evaluation. The strategy can identically be applied when proving knowledge of multiple polynomials at the same query point, which brings resemblance to [BBC+18].

Parameters	Instantiation
m	$\geq n(1 + \tilde{q}) + \omega(\log \lambda)$
δ	$q^{1/O(1)}$
\mathfrak{s}	$> 2Nq^{\frac{n}{m-n\tilde{q}} + \frac{2}{N(m-n\tilde{q})}}$
σ_0	$\geq \delta \mathfrak{s} N n \tilde{q} \cdot \omega(\sqrt{2(d+1)n(m-n\tilde{q})N \log t'N})$
σ_1	$\geq \delta \sigma_0 N \cdot \omega(\sqrt{m'n' \log t'N})$
β	$\geq \sigma_1 \sqrt{m'N}$
β_C	$< \frac{1}{2} \sqrt{l/N} q^{l/N}$
w_s	$sN\beta_C$
β_h	$(\max_l w_{r_l k}) w_k^h w_{rk} \cdot \beta$
γ	$2^{h+2} \cdot (2\beta_C N)^{2^{h+2}-h-3} \cdot \beta_h$
β_s	$(2\beta_C N)^{2^{h+2}-1}$
Soundness	$(Q+1) \cdot \left(\frac{(\max_l r_l + h + r)k}{(2\beta_C + 1)^N} \right)$
Commitment size	$nN \log q \cdot \sum_l r_l$
Proof size	$(\sum_l r_l k N \log q) + r(h+1) \cdot (kN \log q) + \frac{d+1}{k^{h+2}} (rN \log q + mN \log \beta_h)$

Table 6: Parameters and complexity of the multi-evaluation protocol.

Recall that we want to prove knowledge of the polynomial $f \in \mathcal{R}_q[X]$ of degree at most d , and the openings $(\mathbf{s}_i)_{i \in [0, d]}$ such that $f(u) = z$ and $\mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t}$ and $\|\mathbf{s}_i\| \leq \beta$ for $i = 0, 1, \dots, d$. In addition to the public matrices $(\mathbf{A} \in \mathcal{R}_q^{n \times m}, \mathbf{W} \in \mathcal{R}_q^{n \times n})$, this time the index \mathbf{i} contains a short basis \mathbf{T} such that $\mathbf{B}\mathbf{T} = \mathbf{G}_{n(d+1)}$ where ¹⁷

$$\mathbf{B} := \left[\begin{array}{c|c} \mathbf{A} & -\mathbf{G} \\ \cdot & \vdots \\ \cdot & -\mathbf{G} \\ \mathbf{W}^d \mathbf{A} & \end{array} \right] \quad \text{and} \quad \|\mathbf{T}\| \leq \beta_T . \quad (23)$$

This is the case when generating the PowerBASIS commitment in Section 4 since the public parameters are indeed of the form $\text{crs} := (\mathbf{A}, \mathbf{W}, \mathbf{T})$.

We present the protocol in Figure 10. The strategy follows the Fiat-Shamir with Aborts paradigm [Lyu09] using the generalised rejection sampling from [BTT22]. That is, the prover starts by sampling uniformly random $\mathbf{g} := (g_0, \dots, g_d) \leftarrow \mathcal{R}_q^{d+1}$, which corresponds to coefficients of a uniformly random polynomial $g \in \mathcal{R}_q[X]$ of degree at most d . Then, the prover runs the PowerBASIS commitment algorithm for \mathbf{g} (c.f. Figure 4). Namely, it samples

$$\begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} \leftarrow \text{SamplePre}(\mathbf{B}, \mathbf{u}, \mathbf{T}, \sigma), \quad \text{where } \mathbf{u} := \begin{bmatrix} -g_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -g_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix},$$

and sets $\mathbf{t}_y := \mathbf{G} \hat{\mathbf{t}}_y$. The first message sent by the prover is (\mathbf{t}_y, v) where $v := \sum_{i=0}^d g_i u^i$ is the evaluation of g at the point u . Then, the verifier picks a challenge α from the challenge space $\mathcal{C} := S_{\beta_C}$ of short polynomials of infinity norm at most β_C .

¹⁷ See Lemma 4.1 on how to obtain the bound on $\|\mathbf{T}\|$. For presentation, we assume the bound β_T is known.

Next, given a challenge $\alpha \leftarrow \mathcal{C}$ from the verifier, the prover computes

$$\mathbf{z}_i := \mathbf{y}_i + \alpha \mathbf{s}_i \quad \text{and} \quad h_i := g_i + \alpha f_i \quad \text{for } i = 0, 1, \dots, d,$$

and outputs $(\mathbf{z}_i, \mathbf{h}_i)$ after performing the rejection sampling procedure. Note that the distribution of \mathbf{z}_i can be written alternatively as:

$$\begin{bmatrix} \mathbf{z}_0 \\ \vdots \\ \mathbf{z}_d \\ \hat{\mathbf{t}}_z \end{bmatrix} = \begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} + \alpha \begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \quad \text{where} \quad \begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} \leftarrow \text{SamplePre} \left(\left[\begin{array}{c|c} \mathbf{A} & -\mathbf{G} \\ \cdot & \vdots \\ \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right], \begin{bmatrix} -g_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -g_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}, \mathbf{T}, \sigma \right) \quad (24)$$

and $\hat{\mathbf{t}} = \mathbf{G}^{-1}(\mathbf{t})$. Hence, this vector comes from a shifted discrete Gaussian distribution (over a coset of $\Lambda^\perp(\mathbf{B})$), where the norm of the shifted vector can be bounded by:

$$\left\| \begin{bmatrix} \alpha \mathbf{s}_0 \\ \vdots \\ \alpha \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \right\| \leq \beta_{\mathcal{C}} N \cdot \sqrt{(d+1)\beta^2 + n\tilde{q}N}. \quad (25)$$

This interpretation will be useful when analysing the rejection sampling algorithm.

Finally, the verifier checks whether

$$\begin{aligned} \mathbf{A}\mathbf{z}_i + \mathbf{h}_i \mathbf{e}_1 &= \mathbf{W}^{-i}(\mathbf{t}_y + \alpha \mathbf{t}) \quad \text{for } i = 0, 1, \dots, d \\ \|z_i\| &\leq \beta_z \quad \text{for } i = 0, 1, \dots, d \\ \sum_{i=0}^d h_i u^i &= v + \alpha z. \end{aligned}$$

In the following, we give a brief reasoning about completeness, special-soundness and honest-verifier zero-knowledge.

Completeness. By careful inspection, we can deduce from the third verification check:

$$\sum_{i=0}^d h_i u^i = \sum_{i=0}^d g_i u^i + \alpha \sum_{i=0}^d f_i u^i = v + \alpha z,$$

and from the second verification check:

$$\mathbf{A}\mathbf{z}_i + \mathbf{h}_i \mathbf{e}_1 = \mathbf{A}\mathbf{y}_i + \mathbf{g}_i \mathbf{e}_1 + \alpha(\mathbf{A}\mathbf{s}_i + \mathbf{f}_i \mathbf{e}_1) = \mathbf{W}^{-i} \mathbf{t}_y + \alpha \mathbf{W}^{-i} \mathbf{t} = \mathbf{W}^{-i}(\mathbf{t}_y + \alpha \mathbf{t}).$$

What we have left to show is shortness of \mathbf{z}_i . Take the standard deviation

$$\sigma \geq \max \left(O(\sqrt{\lambda}) \cdot \beta_{\mathcal{C}} N \cdot \sqrt{(d+1)\beta^2 + n\tilde{q}N}, \beta_T \cdot \omega(\sqrt{N \log tN}) \right) \quad (26)$$

where $t = \max(n, m)$. By Lemma 2.20, we can swap the `SamplePre` algorithm with truly sampling from a discrete Gaussian. Further, since σ is larger than the shifted vector in (25) by a factor of

HVZK Σ -Protocol for $R_{d,\beta}$

Prover

$\mathbf{s} := (s_0, \dots, s_d, \hat{\mathbf{t}})$ where $\hat{\mathbf{t}} := \mathbf{G}^{-1}(\mathbf{t})$

$\mathbf{g} := (g_0, \dots, g_d) \leftarrow \mathcal{R}_q^{d+1}$

$v := g_0 + g_1 u + \dots + g_d u^d$

$\begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} \leftarrow \text{SamplePre} \left(\left[\begin{array}{c|c} \mathbf{A} & -\mathbf{G} \\ \hline \cdot & \vdots \\ \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right], \begin{bmatrix} -g_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -g_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}, \mathbf{T}, \sigma \right)$

$\mathbf{t}_y := \mathbf{G} \hat{\mathbf{t}}_y$

$\hat{\mathbf{t}}_z := \hat{\mathbf{t}}_y + \alpha \hat{\mathbf{t}}$

for $i = 0, 1, \dots, d$:

$\mathbf{z}_i := \mathbf{y}_i + \alpha \mathbf{s}_i$

$h_i := g_i + \alpha f_i$

$\mathbf{z} := (\mathbf{z}_1, \dots, \mathbf{z}_d, \hat{\mathbf{t}}_z)$

$\rho \leftarrow [0, 1)$

if $\rho > \min \left(\frac{\mathcal{D}_\sigma^{m'N}(\mathbf{z})}{M \cdot \mathcal{D}_{\sigma, \alpha \mathbf{s}}^{m'N}(\mathbf{z})}, 1 \right)$:

$\mathbf{z} := \perp$

Verifier

$\xrightarrow{\mathbf{t}_y, v}$

$\alpha \leftarrow \mathcal{C} := S_{\beta_C}$

$\xleftarrow{\alpha}$

$\xrightarrow{(z_i, h_i)_{i \in [0, d]}}$

Check:

$$\sum_{i \in [0, d]} h_i u^{i-1} = v + \alpha z$$

$$\forall i, \mathbf{A} \mathbf{z}_i + \mathbf{h}_i \mathbf{e}_1 = \mathbf{W}^{-i}(\mathbf{t}_y + \alpha \mathbf{t})$$

$$\forall i, \|\mathbf{z}_i\| \leq \beta_z$$

Fig. 9: The honest-verifier zero-knowledge Σ -protocol for $R_{d,\beta}$. Here, $m' := (d+1)m + n\tilde{q}$ is the width of the matrix \mathbf{B} in (23).

$O(\sqrt{\lambda})$, using rejection sampling (c.f. Lemma 2.11) we enforce the distribution of $(\mathbf{z}_0, \dots, \mathbf{z}_d, \hat{\mathbf{t}}_z)$ from (24) to be from a discrete Gaussian on $\Lambda_{\mathbf{u}}^\perp(\mathbf{B})$ where

$$\mathbf{u} := \begin{bmatrix} -(g_0 + \alpha f_0) \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -(g_d + \alpha f_d) \mathbf{W}^d \mathbf{e}_1 \end{bmatrix} .$$

Thus, by Lemma 2.20, we can set $\beta_z := \sigma \sqrt{(d+1)N}$. The correctness error becomes $\approx 1/M$.

Special-soundness. Given two valid transcripts $(\mathbf{t}_y, v, \alpha, (z_i, h_i)), (\mathbf{t}_y, v, \alpha', (z'_i, h'_i))$ with distinct challenges $\alpha, \alpha' \in \mathcal{C}$, we can define

$$\bar{\mathbf{s}}_i := \frac{\mathbf{z}_i - \mathbf{z}'_i}{\alpha - \alpha'} \quad \text{and} \quad \bar{f}_i := \frac{h_i - h'_i}{\alpha - \alpha'} \quad \text{for } i = 0, 1, \dots, d .$$

Note that $\|\alpha - \alpha'\|_\infty \leq 2\beta_{\mathcal{C}}$. If $\beta_{\mathcal{C}}$ is chosen according to Lemma 2.4 then we deduce that the difference is invertible over \mathcal{R}_q . Further, by construction

$$\bar{f}(u) = \sum_{i=0}^d \bar{f}_i u^i = \frac{1}{\alpha - \alpha'} \sum_{i=0}^d (h_i - h'_i) u^i = \frac{\alpha z - \alpha' z}{\alpha - \alpha'} = z .$$

Furthermore, for $i = 0, 1, \dots, d$ we have $\|(\alpha - \alpha') \mathbf{s}_i\| \leq 2\beta_z$ and

$$\mathbf{A} \bar{\mathbf{s}}_i + \bar{f}_i \mathbf{e}_1 = \frac{1}{\alpha - \alpha'} (\mathbf{A} \mathbf{z}_i + h_i \mathbf{e}_1 - (\mathbf{A} \mathbf{z}'_i + h'_i \mathbf{e}_1)) = \frac{1}{\alpha - \alpha'} (\alpha \mathbf{W}^{-i} \mathbf{t} - \alpha' \mathbf{W}^{-i} \mathbf{t}) = \mathbf{W}^{-i} \mathbf{t} .$$

Thus, $(\bar{\mathbf{s}}_0, \dots, \bar{\mathbf{s}}_d)$ along with the message $(\bar{f}_0, \dots, \bar{f}_d)$ is a relaxed opening for the PowerBASIS commitment \mathbf{t} with the relaxation factor $\alpha - \alpha'$. Hence, we can extract the witness for the relaxed relation $\bar{\mathbf{R}}_{d, 2\beta_z, 1}$ in (22).

Honest-verifier zero-knowledge. We show how to simulate the transcripts when the verifier behaves honestly. To this end, we prove the following lemma which is almost analogous to [BTT22, Lemma B.8].

Lemma 5.17 (Honest-Verifier Zero-Knowledge). *Let σ be chosen as in (26) where $t = \max(n, m)$. Then, the output distributions of \mathcal{T} and \mathcal{S} in Figure 10 are statistically indistinguishable.*

Proof. We prove the statement via a standard hybrid argument.

- Hyb_0 is identical to \mathcal{T} as in Figure 10.
- Hyb_1 is identical to Hyb_0 , but now we define $\hat{\mathbf{t}}_z := \hat{\mathbf{t}}_y + \alpha \hat{\mathbf{t}}$, where $\hat{\mathbf{t}} := \mathbf{G}^{-1}(\mathbf{t})$, and compute $\mathbf{t}_y := \mathbf{G} \hat{\mathbf{t}}_z - \alpha \hat{\mathbf{t}}$. By construction, the output distribution of Hyb_1 is identical to Hyb_0 and

$$\begin{bmatrix} \mathbf{z}_0 \\ \vdots \\ \mathbf{z}_d \\ \hat{\mathbf{t}}_z \end{bmatrix} = \begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} + \alpha \begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \quad \text{where} \quad \begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} \leftarrow \text{SamplePre} \left(\left(\begin{bmatrix} \mathbf{A} & & & -\mathbf{G} \\ & \ddots & & \vdots \\ & & \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{bmatrix}, \begin{bmatrix} -g_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -g_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}, \mathbf{T}, \sigma \right) .$$




Fig. 10: Simulating the transcripts from the Σ -protocol described in Figure 10.

– Hyb_2 is identical to Hyb_1 , but now we compute

$$\begin{bmatrix} \mathbf{z}_0 \\ \vdots \\ \mathbf{z}_d \\ \hat{\mathbf{t}}_z \end{bmatrix} = \begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} + \alpha \begin{bmatrix} \mathbf{s}_0 \\ \vdots \\ \mathbf{s}_d \\ \hat{\mathbf{t}} \end{bmatrix} \quad \text{where} \quad \begin{bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_d \\ \hat{\mathbf{t}}_y \end{bmatrix} \leftarrow \left[\begin{array}{c|c} \mathbf{A} & -\mathbf{G} \\ \hline \cdot & \vdots \\ \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right]_{\sigma}^{-1} \left(\begin{bmatrix} -g_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -g_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix} \right).$$

By Lemma 2.20, Hyb_1 and Hyb_2 are statistically close.

– Hyb_3 is identical to Hyb_2 , but here we directly sample

$$\begin{bmatrix} \mathbf{z}_0 \\ \vdots \\ \mathbf{z}_d \\ \hat{\mathbf{t}}_z \end{bmatrix} \leftarrow \left[\begin{array}{c|c} \mathbf{A} & -\mathbf{G} \\ \hline \cdot & \vdots \\ \mathbf{W}^d \mathbf{A} & -\mathbf{G} \end{array} \right]_{\sigma}^{-1} \left(\begin{bmatrix} -(g_0 + \alpha f_0) \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -(g_d + \alpha f_d) \mathbf{W}^d \mathbf{e}_1 \end{bmatrix} \right)$$

and with probability $1 - 1/M$ we output $\mathbf{z} := \perp$. By the generalised rejection sampling (c.f. Lemma 2.11), Hyb_3 and Hyb_2 are statistically close.

– Hyb_4 is identical to Hyb_3 , except now we efficiently sample:

$$\begin{bmatrix} \mathbf{z}_0 \\ \vdots \\ \mathbf{z}_d \\ \hat{\mathbf{t}}_z \end{bmatrix} \leftarrow \text{SamplePre} \left(\left(\begin{bmatrix} \mathbf{A} & & -\mathbf{G} \\ & \ddots & \vdots \\ & & \mathbf{W}^d \mathbf{A} - \mathbf{G} \end{bmatrix}, \begin{bmatrix} -(g_0 + \alpha f_0) \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -(g_d + \alpha f_d) \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}, \mathbf{T}, \sigma \right) .$$

As before, by Lemma 2.20 we deduce that Hyb_4 and Hyb_3 are statistically close.

– Hyb_5 is identical to Hyb_4 , except now we define $h_i := g_i - \alpha f_i$ for $i = 0, 1, \dots, d$. Thus,

$$\begin{bmatrix} \mathbf{z}_0 \\ \vdots \\ \mathbf{z}_d \\ \hat{\mathbf{t}}_z \end{bmatrix} \leftarrow \text{SamplePre} \left(\left(\begin{bmatrix} \mathbf{A} & & -\mathbf{G} \\ & \ddots & \vdots \\ & & \mathbf{W}^d \mathbf{A} - \mathbf{G} \end{bmatrix}, \begin{bmatrix} -h_0 \mathbf{W}^0 \mathbf{e}_1 \\ \vdots \\ -h_d \mathbf{W}^d \mathbf{e}_1 \end{bmatrix}, \mathbf{T}, \sigma \right) .$$

Furthermore, we set $v := h(v) - \alpha z$. Clearly, the output distributions of Hyb_5 and Hyb_4 are identical.

– Hyb_6 is identical to Hyb_5 , but now we sample each $h_i \leftarrow \mathcal{R}_q$ uniformly at random. Since in Hyb_5 each g_i was sampled uniformly at random from \mathcal{R}_q , we conclude that the output distributions of Hyb_6 and Hyb_5 are identical.

Finally, the output distribution of Hyb_6 is identical to the one by \mathcal{S} which ends the proof. \square

Remark 5.18. Similarly as in Section 5.4, we can combine the HVZK protocol with one round of folding to minimise the total round complexity, and thus the extracted norm growth. This yields an almost identical protocol as in [BBC+18].

5.6 Polynomial Commitments over Finite Fields

So far we showed how to commit and prove evaluations of polynomials over the cyclotomic ring \mathcal{R}_q . We now present how to build polynomial commitments over finite fields of specific form. This will be useful when combining with Polynomial IOPs to obtain succinct arguments of knowledge.

Suppose q is a prime which satisfies $q \equiv 2N/l + 1 \pmod{4N/l}$ for some positive divisor l of N . Then by [LS18, Corollary 1.2], the polynomial $X^N + 1$ factors as:

$$X^N + 1 \equiv \prod_{i=1}^{N/l} (X^l - r_i) \pmod{q}$$

for distinct $r_i \in \mathbb{Z}_q^*$ where all $X^l - r_i$ are irreducible in the ring $\mathbb{Z}_q[X]$. Further, by the Chinese Remainder Theorem, there exists a ring isomorphism $\varphi : \mathbb{F}^{N/l} \rightarrow \mathcal{R}_q$ where \mathbb{F} is a finite field of size q^l . Consider the restricted function:

$$\begin{aligned} \varphi_{\mathbb{F}} : \mathbb{F} &\rightarrow \mathcal{R}_q \\ x &\mapsto \phi(x, 0, \dots, 0). \end{aligned}$$

By construction, the image of $\varphi_{\mathbb{F}}$ can be described as

$$\mathcal{S}_q := \text{Im}(\varphi_{\mathbb{F}}) = \{\phi(x, 0, \dots, 0) : x \in \mathbb{F}\} .$$

The following simple lemma states that \mathcal{S}_q is an ideal of \mathcal{R}_q .

Lemma 5.19. *The set $\mathcal{S}_q \subseteq \mathcal{R}_q$ defined above is an ideal.*

Proof. The fact that \mathcal{S}_q is an additive subgroup of \mathcal{R}_q follows directly from the additively homomorphic properties of φ . Now let $a \in \mathcal{S}_q$, i.e. $\varphi(x, 0, \dots, 0) = a$ for some $x \in \mathbb{F}$. Further, take arbitrary $\gamma \in \mathcal{R}_q$ and let $(\gamma_1, \dots, \gamma_{N/l}) := \varphi^{-1}(\gamma)$. Then, by the multiplicative homomorphism of φ we get

$$\gamma \cdot a = \varphi(\gamma_1, \dots, \gamma_{N/l}) \cdot \varphi(x, 0, \dots, 0) = \varphi(\gamma_1 x, 0, \dots, 0) = \varphi_{\mathbb{F}}(\gamma_1 x) \in \mathcal{S}_q ,$$

which concludes the proof. \square

Suppose we want to commit to a polynomial $F := \sum_{i=0}^d F_i X^i \in \mathbb{F}[X]$ of degree at most d , and prove evaluation $F(x) = y$ for $x, y \in \mathbb{F}$. By the homomorphic property of $\varphi_{\mathbb{F}}$, this is equivalent to proving $f(u) = z$ over \mathcal{R}_q where

$$\begin{cases} f[X] = \sum_{i=0}^d \varphi_{\mathbb{F}}(F_i) X^i \in \mathcal{S}_q[X] \\ u = \varphi_{\mathbb{F}}(x) \in \mathcal{S}_q \\ z = \varphi_{\mathbb{F}}(y) \in \mathcal{S}_q \end{cases} .$$

Hence, we can commit to the polynomial $f \in \mathcal{R}_q[X]$ and prove evaluation of u at the point z as before. What is new is that we additionally need to prove that coefficients of f indeed lie in \mathcal{S}_q . Therefore, we are interested in a stronger relation:

$$\left\{ \left((\mathbf{A}, \mathbf{W}), (\mathbf{t}, u, z), (f, (\mathbf{s}_i)_i) \right) \left| \begin{array}{l} f(u) = z \wedge f \in \mathcal{S}_q[X] \\ \forall i \in [0, d], \mathbf{A}\mathbf{s}_i + f_i \mathbf{e}_1 = \mathbf{W}^{-i} \mathbf{t} \\ \wedge \|\mathbf{s}_i\| \leq \beta \end{array} \right. \right\} . \quad (27)$$

We show how to modify the protocol in Figure 6 to accommodate for this change. Actually, the interaction between the prover and the verifier stays the same but the verifier additionally performs a check whether the final polynomial $f_h \in \mathcal{R}_q[X]$ sent by the prover has coefficients in \mathcal{S}_q .

Completeness follows by induction. We start with the initial polynomial $f_0 := f \in \mathcal{S}_q[X]$. Then for each $r \in [h]$, the prover computes the polynomial $f_r \in \mathcal{R}_q[X]$ as a linear combination of “partial terms” of f_{r-1} :

$$f_r := \sum_{t \in [k]} \alpha_{r,t} f_{r-1,t} .$$

If $f_{r-1} \in \mathcal{S}_q[X]$, then by Lemma 5.19 we deduce that $f_r \in \mathcal{S}_q[X]$.

To argue (coordinate-wise) special soundness, consider the extractor in the proof of Lemma 5.6. The coefficients of the extracted polynomial f are computed as

$$f_{2i} := \frac{\alpha_1 g_{0,i} - \alpha_0 g_{1,i}}{\alpha_1 - \alpha_0}, \quad f_{2i+1} := \frac{g_{0,i} - g_{1,i}}{\alpha_0 - \alpha_1} \quad \text{for } i \in [0, d/2] .$$

If polynomials g_0 and g_1 have coefficients in \mathcal{S}_q , then again by Lemma 5.19 we can deduce that $f \in \mathcal{S}_q[X]$. Identical argument holds when analysing Lemma 5.10.

Finally, to support honest-verifier zero-knowledge in Figure 9, we let the prover pick uniformly random elements g_i from \mathcal{S}_q instead of \mathcal{R}_q in order to fully mask the coefficients f_i . Thus, by construction and Lemma 5.19, $h_i = g_i + \alpha f_i \in \mathcal{S}_q$ for all $i = 0, \dots, d$. Hence, the verifier additionally performs the check whether coefficients h_i lie in \mathcal{S}_q .

6 Concrete Instantiation and Applications to Marlin

Hardness of PowerBASIS. In parameter selection, we make a heuristic assumption that PowerBASIS is exactly as hard as MSIS. Hence, one should treat our computed sizes only as intuition on how practical the polynomial commitment is.

In the literature, hardness of the MSIS problems is often analysed identically as the plain SIS since, so far, the best known attacks do not make use of the algebraic structure of the polynomial ring [ADPS16]. We follow the methodology from Dilithium [DKL+18, Appendix C]. That is, $\text{MSIS}_{n,m,N,q,\beta}$ for matrix \mathbf{A} is equivalent to finding a non-trivial vector of norm smaller than β in the lattice $\Lambda := \Lambda^\perp(\mathbf{A})$. In order to find short non-trivial vectors in Λ , we apply the Block-Korkine-Zolotarev algorithm (BKZ) [SE94; CN11]. As a subroutine, BKZ uses an algorithm for the shortest vector problem (SVP) in lattices of dimension b , where b is called the block size. If we apply the best known algorithm for solving SVP with no memory constraints by Becker et al. [BDGL16], the time required by BKZ to run on the mN -dimensional lattice Λ with block size b is given by $8mN \cdot 2^{0.292b+16.4}$ (one also considers a more *conservative* variant with runtime $2^{0.292b}$). The algorithm outputs a vector of norm $\delta_{\text{rhf}}^{mN} \det(\Lambda)^{\frac{1}{mN}}$ where δ_{rhf} is the root Hermite factor and it is given by

$$\delta_{\text{rhf}} = \left(\frac{b(\pi b)^{1/b}}{2\pi e} \right)^{\frac{1}{2(b-1)}}. \quad (28)$$

For our usual parameter selection, the probability that a random matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ is of full rank is overwhelming (see [Ezs+19, Appendix C]) and thus $\det(\Lambda) = q^{nN}$. Next, Micciancio and Regev [MR09] show that

$$\delta_{\text{rhf}}^{mN} \det(\Lambda)^{\frac{1}{mN}} = \delta_{\text{rhf}}^{mN} q^{\frac{nN}{mN}} \geq 2^{2\sqrt{nN \log q \log \delta}}$$

and the equality holds when $mN = \sqrt{nN \log q / \log \delta}$. Hence, given a bound $\beta < q$ we compute δ_{rhf} from the equation $\beta = 2^{2\sqrt{nN \log q \log \delta}}$. Next, we calculate the minimum block size b from Equation (28), and thus we get the total time for BKZ to solve $\text{MSIS}_{n,m,N,q,\beta}$. Hereafter, we will refer to the “aggressive strategy” to set PowerBASIS as the one using the estimate from Becker et al. [BDGL16], and to the “conservative strategy” as the one using $2^{0.292b}$.

Parameters. Using a combination of randomised and exhaustive search, we found parameters for the schemes in Theorem 5.8 and Theorem 5.12. In Table 7 we detail the parameters obtained for the scheme presented in Theorem 5.12 and in Table 8 for that in Theorem 5.8. Since we aim to support prime order fields, by the reasoning in Section 5.6 we require that $q \equiv 2N + 1 \pmod{4N}$. We also require, for soundness, that the preconditions of Lemma 5.9 holds, when $l = 1$. Since our moduli are quite large, this can be easily verified. We stress that these parameters are presented to give the reader an indication of the concrete efficiency of the scheme. The commitments have sizes on the order of hundreds of kilobytes, while evaluation proofs are on the order of a few megabytes, and so are larger than desirable in most applications. We also emphasise that the assumption that the hardness of PowerBASIS is as hard as MSIS is an heuristic, and thus, until this heuristic is backed or disproved by sufficient cryptanalysis, the sizes should be considered as an optimistic lower bound.

Applications to Polynomial IOPs. Marlin [CHM+20] is a widely deployed preprocessing zkSNARK. As many modern constructions, Marlin is constructed by combining two ingredients:

- a polynomial interactive oracle proof (PIOP) (therein a algebraic holographic proof);

k	h	d	λ	Q	n	m	N	δ	$\log q$	$2\gamma\beta_s$	β	\mathfrak{s}	β_C	β_h	$ \mathbf{t} $	$ \pi $
1024	2	2^{20}	80	64	188	5271	32	13	234	225	155	29	4	183	172 KB	5.5 MB
1024	3	2^{30}	80	64	325	7807	32	24	376	376	221	51	4	263	477 KB	12.2 MB
1024	2	2^{20}	128	64	255	5745	32	18	256	252	173	37	5	204	255 KB	6.5 MB
1024	3	2^{30}	128	64	376	9031	32	26	404	404	229	54	5	275	593 KB	14.2 MB

Table 7: Parameters and concrete sizes for the polynomial commitment described in Theorem 5.12. δ , norms and standard deviation given in log form.

h	d	λ	n	m	N	δ	$\log q$	$2\gamma\beta_s$	β	\mathfrak{s}	β_h	t	$ \mathbf{t} $	$ \mathbf{cc} $
20	2^{20}	80	48	1159	256	33	512	477	236	68	276	18	768 KB	187.0 MB
30	2^{30}	80	36	871	512	48	758	711	320	100	380	16	1.7 MB	368.0 MB
20	2^{20}	128	58	1399	256	33	513	478	237	69	277	28	930 KB	349.3 MB
30	2^{30}	128	41	991	512	46	719	701	310	94	370	26	1.8 MB	651.7 MB

Table 8: Parameters and concrete sizes for the interactive polynomial commitment in Theorem 5.8. δ , norms and standard deviation given in log form.

– and a polynomial commitment scheme.

An interactive oracle proof (IOP) is a generalisation of both probabilistically checkable proofs and interactive proofs. Informally, they are interactive protocols between a prover and a verifier, in which the prover sends *oracle messages*, which the verifier is allowed to not read in their entirety. A PIOP is simply an IOP where the prover messages are guaranteed to be (low degree) polynomials. IOPs and PIOPs are *information theoretic object*, and as such inherit a number of efficiency limitations (for example, IOP proof length are required to be at least linear in the size of the instance), but can be compiled using cryptography (see [BCS16]) to obtain arguments that are both asymptotically and concretely efficient. Informally, to compile a PIOP into an interactive argument, the prover can commit to each polynomial oracle using a polynomial commitment scheme, and then prove to the verifier that the evaluations (at points chosen by the verifier) are as claimed. Then, to obtain a NARK, we can apply the Fiat-Shamir transformation to this interactive protocol. We can thus aim to use our polynomial commitment scheme in Theorem 5.12 as an ingredient of Marlin to obtain a zkSNARK for R1CS. Let d denote the size of the R1CS instance that we aim to prove. As detailed in [CHM+20, Section 9], Marlin after compilation has commitments to 19 total polynomials of degree at most $6d$. The prover has then to produce 19 evaluations proofs for these polynomials, at three distinct points. We can thus apply the techniques in Section 5.4 to batch evaluations together and amortise the cost of the last round. In Table 9 we compute parameters for Marlin instantiated using our polynomial commitment scheme and the PIOP therein described. Again, these sizes are meant to give a rough estimate of the concrete efficiency of the scheme, and the same caveats apply as

k	h	d	λ	Q	n	m	N	δ	$\log q$	$2\gamma\beta_s$	β	\mathfrak{s}	β_C	β_h	$ \mathbf{t} $	$ \pi $
[64, 256, 384]	1	2^{20}	80	64	322	7735	32	24	372	372	205	50	4	259	8.2 MB	12.0 MB
[16, 512, 512, 1536]	2	2^{30}	80	64	518	14770	32	32	600	599	262	67	4	336	21.3 MB	37.2 MB
[64, 256, 384]	1	2^{20}	128	64	399	9583	32	26	402	401	213	53	5	272	11.0 MB	14.7 MB
[16, 512, 512, 1536]	2	2^{30}	128	64	598	16153	32	38	667	666	284	77	5	363	27.4 MB	42.7 MB

Table 9: Parameters and concrete sizes for Marlin when instantiated with the commitment described in Theorem 5.12 with amortisation as in Table 6. δ , norms and standard deviation given in log form. Folding factor varies across rounds as mentioned in Remark 5.4

with the polynomial commitment scheme. We also note that Marlin operates over fields with a large multiplicative (or additive) subgroup with smooth order, which imposes an additional requirement on the size of q . Since our moduli are again quite large, this additional requirement is immaterial.

Falsifiable version of PowerBASIS. Note that the challenger in the PowerBASIS game from Section 3 is not efficient since it needs to sample a random trapdoor \mathbf{T} according to a discrete Gaussian distribution. In order to make the assumption falsifiable, one could let the challenger sample efficiently using the SamplePre algorithm, e.g. as in the Setup algorithm of Figure 4, and only ensure that the sampled matrix \mathbf{A} from $(\mathbf{A}, \mathbf{R}) \leftarrow \text{TrapGen}(n, m)$ is *computationally* indistinguishable from random. Thus, we would rely on the Module-LWE [LS15] as opposed to the regularity lemma (c.f. Lemma 2.7), which results in picking moderately smaller values for m . However, we do not apply this heuristic in our parameter selection.

References

- [ACK21] T. Attema, R. Cramer, and L. Kohl. “A Compressed \mathbb{Z} -Protocol Theory for Lattices”. In: *CRYPTO (2)*. Vol. 12826. Lecture Notes in Computer Science. Springer, 2021, pp. 549–579.
- [ACL+22] M. R. Albrecht, V. Cini, R. W. F. Lai, G. Malavolta, and S. A. K. Thyagarajan. “Lattice-Based SNARKs: Publicly Verifiable, Preprocessing, and Recursively Composable - (Extended Abstract)”. In: *CRYPTO (2)*. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 102–132.
- [ADPS16] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. “Post-quantum Key Exchange - A New Hope”. In: *USENIX Security Symposium*. USENIX Association, 2016, pp. 327–343.
- [AF22] T. Attema and S. Fehr. “Parallel Repetition of (k_1, \dots, k_μ) -Special-Sound Multi-Round Interactive Proofs”. In: *CRYPTO (1)*. Vol. 13507. Lecture Notes in Computer Science. Springer, 2022, pp. 415–443.
- [AFK22] T. Attema, S. Fehr, and M. Kloöß. *Fiat-Shamir Transformation of Multi-round Interactive Proofs*. 2022.
- [Ajt96] M. Ajtai. “Generating hard instances of lattice problems”. In: *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing*. STOC ’96. 1996, pp. 99–108.
- [AKSY22] S. Agrawal, E. Kirshanova, D. Stehlé, and A. Yadav. “Practical, Round-Optimal Lattice-Based Blind Signatures”. In: *CCS*. ACM, 2022, pp. 39–53.
- [AL21] M. R. Albrecht and R. W. F. Lai. “Subtractive Sets over Cyclotomic Rings - Limits of Schnorr-Like Arguments over Lattices”. In: *CRYPTO (2)*. Vol. 12826. Lecture Notes in Computer Science. Springer, 2021, pp. 519–548.
- [ALS20] T. Attema, V. Lyubashevsky, and G. Seiler. “Practical Product Proofs for Lattice Commitments”. In: *CRYPTO (2)*. Vol. 12171. Lecture Notes in Computer Science. Springer, 2020, pp. 470–499.
- [BBB+18] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. “Bulletproofs: Short Proofs for Confidential Transactions and More”. In: *IEEE Symposium on Security and Privacy*. 2018, pp. 315–334.
- [BBC+18] C. Baum, J. Bootle, A. Cerulli, R. del Pino, J. Groth, and V. Lyubashevsky. “Sub-linear Lattice-Based Zero-Knowledge Arguments for Arithmetic Circuits”. In: *CRYPTO*. 2018, pp. 669–699.
- [BBHR19] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev. “Scalable Zero Knowledge with No Trusted Setup”. In: *Proceedings of the 39th Annual International Cryptology Conference*. CRYPTO ’19. 2019, pp. 733–764.
- [BCC+16] J. Bootle, A. Cerulli, P. Chaidos, J. Groth, and C. Petit. “Efficient Zero-Knowledge Arguments for Arithmetic Circuits in the Discrete Log Setting”. In: *EUROCRYPT*. 2016, pp. 327–357.
- [BCFL22] D. Balbás, D. Catalano, D. Fiore, and R. W. F. Lai. *Functional Commitments for Circuits from Falsifiable Assumptions*. Cryptology ePrint Archive, Paper 2022/1365. <https://eprint.iacr.org/2022/1365>. URL: <https://eprint.iacr.org/2022/1365>.
- [BCHO22] J. Bootle, A. Chiesa, Y. Hu, and M. Orrù. “Gemini: Elastic SNARKs for Diverse Environments”. In: *Proceedings of the 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’22. 2022, pp. 427–457.
- [BCI+13] N. Bitansky, A. Chiesa, Y. Ishai, R. Ostrovsky, and O. Paneth. “Succinct Non-Interactive Arguments via Linear Interactive Proofs”. In: *Proceedings of the 10th Theory of Cryptography Conference*. TCC ’13. 2013, pp. 315–333.

- [BCK+14] F. Benhamouda, J. Camenisch, S. Krenn, V. Lyubashevsky, and G. Neven. “Better Zero-Knowledge Proofs for Lattice Encryption and Their Application to Group Signatures”. In: *ASIACRYPT*. 2014, pp. 551–572.
- [BCS16] E. Ben-Sasson, A. Chiesa, and N. Spooner. “Interactive Oracle Proofs”. In: *Proceedings of the 14th Theory of Cryptography Conference*. TCC ’16-B. 2016, pp. 31–60.
- [BCS21] J. Bootle, A. Chiesa, and K. Sotiraki. “Sumcheck Arguments and Their Applications”. In: *CRYPTO (1)*. Vol. 12825. Lecture Notes in Computer Science. Springer, 2021, pp. 742–773.
- [BDGL16] A. Becker, L. Ducas, N. Gama, and T. Laarhoven. “New directions in nearest neighbor searching with applications to lattice sieving”. In: *SODA*. SIAM, 2016, pp. 10–24.
- [BDK+18] J. W. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé. “CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM”. In: *2018 IEEE European Symposium on Security and Privacy, EuroS&P*. 2018, pp. 353–367.
- [BDL+18] C. Baum, I. Damgård, V. Lyubashevsky, S. Oechsner, and C. Peikert. “More Efficient Commitments from Structured Lattice Assumptions”. In: *SCN*. 2018, pp. 368–385.
- [BF22] B. Bünz and B. Fisch. *Multilinear Schwartz-Zippel mod N with Applications to Succinct Arguments*. Cryptology ePrint Archive, Paper 2022/458. <https://eprint.iacr.org/2022/458>. 2022. URL: <https://eprint.iacr.org/2022/458>.
- [BFS20] B. Bünz, B. Fisch, and A. Szepieniec. “Transparent SNARKs from DARK Compilers”. In: *EUROCRYPT (1)*. Vol. 12105. Lecture Notes in Computer Science. Springer, 2020, pp. 677–706.
- [BLNS20] J. Bootle, V. Lyubashevsky, N. K. Nguyen, and G. Seiler. “A Non-PCP Approach to Succinct Quantum-Safe Zero-Knowledge”. In: *CRYPTO (2)*. Vol. 12171. Lecture Notes in Computer Science. Springer, 2020, pp. 441–469.
- [BLNS23] J. Bootle, V. Lyubashevsky, N. K. Nguyen, and A. Sorniotti. *A Framework for Practical Anonymous Credentials from Lattices*. To appear at CRYPTO 2023. <https://eprint.iacr.org/2023/560>. 2023. URL: <https://eprint.iacr.org/2023/560>.
- [BS22] W. Beullens and G. Seiler. “LaBRADOR: Compact Proofs for R1CS from Module-SIS”. In: (2022). <https://eprint.iacr.org/2022/1341>. URL: <https://eprint.iacr.org/2022/1341>.
- [BTT22] C. Boschini, A. Takahashi, and M. Tibouchi. “MuSig-L: Lattice-Based Multi-signature with Single-Round Online Phase”. In: *CRYPTO (2)*. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 276–305.
- [CHM+20] A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. Ward. “Marlin: Preprocessing zkSNARKs with Universal and Updatable SRS”. In: *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’20. 2020, pp. 738–768.
- [CN11] Y. Chen and P. Q. Nguyen. “BKZ 2.0: Better Lattice Security Estimates”. In: *ASIACRYPT*. Vol. 7073. Lecture Notes in Computer Science. Springer, 2011, pp. 1–20.
- [CP22] L. de Castro and C. Peikert. “Functional Commitments for All Functions, with Transparent Setup”. In: *IACR Cryptol. ePrint Arch.* (2022), p. 1368.
- [DFM20] J. Don, S. Fehr, and C. Majenz. “The Measure-and-Reprogram Technique 2.0: Multi-round Fiat-Shamir and More”. In: *CRYPTO (3)*. Vol. 12172. Lecture Notes in Computer Science. Springer, 2020, pp. 602–631.
- [DKL+18] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. “CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme”. In: *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2018.1 (2018), pp. 238–268.
- [DLP14] L. Ducas, V. Lyubashevsky, and T. Prest. “Efficient Identity-Based Encryption over NTRU Lattices”. In: *ASIACRYPT*. 2014, pp. 22–41.
- [ENS20] M. F. Esgin, N. K. Nguyen, and G. Seiler. “Practical Exact Proofs from Lattices: New Techniques to Exploit Fully-Splitting Rings”. In: *ASIACRYPT (2)*. 2020, pp. 259–288.
- [EZS+19] M. F. Esgin, R. K. Zhao, R. Steinfeld, J. K. Liu, and D. Liu. “MatRiCT: Efficient, Scalable and Post-Quantum Blockchain Confidential Transactions Protocol”. In: *CCS*. ACM, 2019, pp. 567–584.
- [FHK+20] P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Prest, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. *Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU*. Tech. rep. <https://falcon-sign.info/falcon.pdf>. 2020.
- [FS86] A. Fiat and A. Shamir. “How to Prove Yourself: Practical Solutions to Identification and Signature Problems”. In: *CRYPTO*. 1986, pp. 186–194.
- [GLS+21] A. Golovnev, J. Lee, S. T. V. Setty, J. Thaler, and R. S. Wahby. “Brakedown: Linear-time and post-quantum SNARKs for R1CS”. In: *IACR Cryptol. ePrint Arch.* (2021), p. 1043.

- [GMNO18] R. Gennaro, M. Minelli, A. Nitulescu, and M. Orrù. “Lattice-Based zk-SNARKs from Square Span Programs”. In: *Proceedings of the 25th ACM Conference on Computer and Communications Security*. CCS ’18. 2018, pp. 556–573.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions”. In: *STOC*. 2008, pp. 197–206.
- [HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. “NTRU: A Ring-Based Public Key Cryptosystem”. In: *ANTS*. Vol. 1423. Lecture Notes in Computer Science. Springer, 1998, pp. 267–288.
- [ISW21] Y. Ishai, H. Su, and D. J. Wu. “Shorter and Faster Post-Quantum Designated-Verifier zkSNARKs from Lattices”. In: *CCS*. ACM, 2021, pp. 212–234.
- [JRLS22] C. Jeudy, A. Roux-Langlois, and O. Sanders. *Lattice Signature with Efficient Protocols, Application to Anonymous Credentials*. Cryptology ePrint Archive, Paper 2022/509. <https://eprint.iacr.org/2022/509>. 2022. URL: <https://eprint.iacr.org/2022/509>.
- [Kat21] S. Katsumata. “A New Simple Technique to Bootstrap Various Lattice Zero-Knowledge Proofs to QROM Secure NIZKs”. In: *CRYPTO (2)*. Vol. 12826. Lecture Notes in Computer Science. Springer, 2021, pp. 580–610.
- [KZG10] A. Kate, G. M. Zaverucha, and I. Goldberg. “Constant-Size Commitments to Polynomials and Their Applications”. In: *ASIACRYPT*. Vol. 6477. Lecture Notes in Computer Science. Springer, 2010, pp. 177–194.
- [LMS22] R. W. F. Lai, G. Malavolta, and N. Spooner. “Quantum Rewinding for Many-Round Protocols”. In: *TCC (1)*. Vol. 13747. Lecture Notes in Computer Science. Springer, 2022, pp. 80–109.
- [LN22] V. Lyubashevsky and N. K. Nguyen. “BLOOM: Bimodal Lattice One-out-of-Many Proofs and Applications”. In: *ASIACRYPT (4)*. Vol. 13794. Lecture Notes in Computer Science. Springer, 2022, pp. 95–125.
- [LNP22] V. Lyubashevsky, N. K. Nguyen, and M. Plançon. “Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General”. In: *CRYPTO (2)*. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 71–101.
- [LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. “A Toolkit for Ring-LWE Cryptography”. In: *EUROCRYPT*. 2013, pp. 35–54.
- [LRY16] B. Libert, S. C. Ramanna, and M. Yung. “Functional Commitment Schemes: From Polynomial Commitments to Pairing-Based Accumulators from Simple Assumptions”. In: *ICALP*. Vol. 55. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, 30:1–30:14.
- [LS15] A. Langlois and D. Stehlé. “Worst-case to average-case reductions for module lattices”. In: *Des. Codes Cryptogr.* 75.3 (2015), pp. 565–599.
- [LS18] V. Lyubashevsky and G. Seiler. “Short, Invertible Elements in Partially Splitting Cyclotomic Rings and Applications to Lattice-Based Zero-Knowledge Proofs”. In: *EUROCRYPT (1)*. Springer, 2018, pp. 204–224.
- [LS19] V. Lyubashevsky and G. Seiler. “NTTRU: Truly Fast NTRU Using NTT”. In: *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2019.3 (2019), pp. 180–201.
- [LTV12] A. López-Alt, E. Tromer, and V. Vaikuntanathan. “On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption”. In: *STOC*. 2012, pp. 1219–1234.
- [Lyu09] V. Lyubashevsky. “Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures”. In: *ASIACRYPT*. 2009, pp. 598–616.
- [Lyu12] V. Lyubashevsky. “Lattice Signatures Without Trapdoors”. In: *EUROCRYPT*. 2012, pp. 738–755.
- [MP12] D. Micciancio and C. Peikert. “Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller”. In: *EUROCRYPT*. 2012, pp. 700–718.
- [MR07] D. Micciancio and O. Regev. “Worst-Case to Average-Case Reductions Based on Gaussian Measures”. In: *SIAM Journal on Computing* 37 (1 2007), pp. 267–302.
- [MR09] D. Micciancio and O. Regev. “Lattice-based Cryptography”. In: *Post-Quantum Cryptography*. Ed. by D. J. Bernstein, J. Buchmann, and E. Dahmen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 147–191. ISBN: 978-3-540-88702-7. DOI: 10.1007/978-3-540-88702-7_5. URL: https://doi.org/10.1007/978-3-540-88702-7_5.
- [NS22] N. K. Nguyen and G. Seiler. “Practical Sublinear Proofs for R1CS from Lattices”. In: *CRYPTO (2)*. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 133–162.
- [PPS21] C. Peikert, Z. Pepin, and C. Sharp. “Vector and Functional Commitments from Lattices”. In: *TCC (3)*. Vol. 13044. Lecture Notes in Computer Science. Springer, 2021, pp. 480–511.

- [SE94] C.-P. Schnorr and M. Euchner. “Lattice basis reduction: Improved practical algorithms and solving subset sum problems”. In: *Math. Program.* 66 (1994), pp. 181–199.
- [Sei18] G. Seiler. “Faster AVX2 optimized NTT multiplication for Ring-LWE lattice cryptography”. In: *IACR Cryptology ePrint Archive 2018* (2018). <http://eprint.iacr.org/2018/039>, p. 39.
- [Set20] S. Setty. “Spartan: Efficient and general-purpose zkSNARKs without trusted setup”. In: *Proceedings of the 40th Annual International Cryptology Conference*. CRYPTO ’20. Referencing Cryptology ePrint Archive, Report 2019/550, revision from 2020.02.28. 2020, pp. 704–737.
- [SS13] D. Stehlé and R. Steinfeld. “Making NTRUEncrypt and NTRUSign as Secure as Standard Worst-Case Problems over Ideal Lattices”. In: *IACR Cryptol. ePrint Arch.* (2013), p. 4.
- [SSEK22] R. Steinfeld, A. Sakzad, M. F. Esgin, and V. Kuchta. *Private Re-Randomization for Module LWE and Applications to Quasi-Optimal ZK-SNARKs*. Cryptology ePrint Archive, Paper 2022/1690. <https://eprint.iacr.org/2022/1690>. 2022. URL: <https://eprint.iacr.org/2022/1690>.
- [WW23] H. Wee and D. J. Wu. “Succinct Vector, Polynomial, and Functional Commitments from Lattices”. In: *EUROCRYPT (3)*. Vol. 14006. Lecture Notes in Computer Science. Full version: <https://eprint.iacr.org/2022/1515>. Springer, 2023, pp. 385–416.