Compressed Σ-protocol theory

Attema, T.

CHAPTER 5

5 Suitable Cryptographic Platforms

5.1 Introduction

Thus far, we have seen how to prove knowledge of homomorphism preimages. One of the main applications of this functionality is *opening linear forms on compactly committed vectors*. More precisely, proving knowledge of a (vector) commitment opening that satisfies some arbitrary linear constraint captured by a linear form. Our compressed Σ-protocols require the vector commitment scheme to be *homomorphic*. Moreover, since in every iteration of the compression mechanism the prover sends two commitments, the communication complexity is only reduced if the commitment scheme is *compact*, or at least *compressing*. Recall that the size of a compact vector commitment is constant in the dimension n of the committed vector and the size of a compressing commitment is merely sublinear in n.

It is easy to see that compact commitments can be at most *computationally* binding; the domain of the commitment function is much larger than its codomain. For this reason, compact and homomorphic commitment schemes are to be based on computational assumptions. In this chapter, we will present a number of cryptographic platforms in which commitment schemes with the desired properties, and their corresponding compressed Σ-protocols, can be instantiated. The instantiations of this chapter are based on the papers [AC20; ACK21; ACR21], co-authored by Ronald Cramer, Lisa Kohl and Matthieu Rambaud.

5.2 Discrete Logarithm Assumption

The most prominent example of a compact and homomorphic vector commitment scheme is the Pedersen vector commitment scheme [Ped91]. This scheme allows a prover to commit to n-dimensional vectors[1] of field elements $\mathbf{x} \in \mathbb{Z}_q^n$, where q is a prime. A commitment is a single group element, regardless of the dimension n, i.e., commitments are indeed compact. The commitment scheme is perfectly hiding and computationally binding under the discrete logarithm assumption. Its formal definition is given below.

[1] Actually, Pedersen only introduced a commitment scheme for single elements $x \in \mathbb{Z}_q$. The vector commitment scheme presented here is a natural generalization and is therefore typically referred to as the Pedersen vector commitment scheme.

Definition 5.1 (Pedersen Vector Commitment Scheme [Ped91])**.** The Pedersen vector commitment scheme is defined by the following setup algorithm and commitment function:

- $\mathsf{pk} = (q, \mathbb{H}, g_1, \ldots, g_n, h) \leftarrow \textsc{Setup}(1^\lambda, n)$, where $(q, \mathbb{H}, \cdot) \leftarrow \mathcal{G}(1^\lambda)$ for a prime order group generator $\mathcal{G}(\cdot)$, i.e., $q = |\mathbb{H}|$ is prime, and

$$(\mathbf{g}, h) = (g_1, \ldots, g_n, h) \leftarrow_R \mathbb{H}^{n+1}$$

  are sampled uniformly at random;

- $\textsc{com}_{\mathsf{pk}} \colon \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{H}, \quad (\mathbf{x}; \gamma) \mapsto \mathbf{g}^{\mathbf{x}} h^\gamma := h^\gamma \prod_{i=1}^n g_i^{x_i}$.

Recall that to commit to a vector $\mathbf{x} \in \mathbb{Z}_q^n$, the prover samples $\gamma \leftarrow_R \mathbb{Z}_q$ uniformly at random and outputs the commitment $\textsc{com}_{\mathsf{pk}}(\mathbf{x}; \gamma)$.

The Pedersen commitment function is a homomorphism, i.e., the compressed Σ-protocols of Chapter 3 apply. Since the randomness $\gamma \in \mathbb{Z}_q$ and coefficients $x_1, \ldots, x_n \in \mathbb{Z}_q$ of a Pedersen commitment $\textsc{com}_{\mathsf{pk}}(\mathbf{x}; \gamma)$ are all field elements, the randomness can be compressed too. More precisely, instead of compressing a vector $\mathbf{x}$ of dimension n, a vector $(\mathbf{x}; \gamma)$ of dimension $n+1$ will be compressed. This yields a minor improvement with respect to the abstract treatment of Section 3.4.

Theorem 5.1 now summarizes the main properties of the resulting compressed Σ-protocol for opening linear forms on Pedersen commitments. We immediately consider the most efficient variant of Theorem 3.11, where the linear form evaluation is incorporated into the commitment. More precisely, compression is applied to the homomorphism

$$\Psi \colon \mathbb{Z}_q^{n+1} \to \mathbb{H}, \quad (\mathbf{x}; \gamma) \mapsto \textsc{com}_{\mathsf{pk}}(\mathbf{x}, c \cdot L(\mathbf{x}); \gamma)\,,$$

for some challenge $c \leftarrow_R \mathbb{Z}_q$ sent by the verifier in the first round of the protocol. This variant has *computational* special-soundness. At the cost of increasing the communication costs by roughly a factor two, or by using the techniques from Section 3.4.3, this compressed Σ-protocol can be made *unconditionally* special-sound.

Theorem 5.1 (Compressed Σ-Protocol for Pedersen Commitments)**.** *Let $n + 1 = 2^\mu$ for some $\mu \in \mathbb{N}$, $\textsc{com}_{pk}$ the Pedersen vector commitment scheme instantiated with public key $\mathsf{pk} = (q, \mathbb{H}, g_1, \ldots, g_n, h)$ and $L \colon \mathbb{Z}_q^n \to \mathbb{Z}_q$ a linear form.*

Then the compressed Σ-protocol for relation

$$\mathfrak{R}_{Ped} = \{(P, y; \mathbf{x}, \gamma) : \mathbf{g}^{\mathbf{x}} h^\gamma = P \wedge L(\mathbf{x}) = y\}\,,$$

is perfectly complete, computationally $(2, 2, 3, \ldots, 3)$-out-of-$(q, \ldots, q)$ special-sound, under the discrete logarithm assumption, and special honest-verifier zero-knowledge (SHVZK). Moreover, it has $(2\mu + 2)$ communication rounds and the communication costs are:

- *$\mathcal{P} \to \mathcal{V}$: 2 elements of $\mathbb{Z}_q$ and $2\mu - 1$ elements of $\mathbb{H}$;*
- *$\mathcal{V} \to \mathcal{P}$: $\mu + 1$ elements of $\mathbb{Z}_q$.*

5.3 Pairing-Based Platform

In a pairing-based platform, the Pedersen commitment scheme has a straightforward adaptation to accommodate vectors of group, rather than field, elements [AFG+10]. More precisely, let $e\colon \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{H}$ be a (nondegenerate) bilinear mapping between groups $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ and $(\mathbb{H}, \cdot)$ of prime order q, i.e., e is a pairing and $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e)$ is a bilinear group. The adapted Pedersen vector commitment scheme allows a prover to commit to vectors $\mathbf{x}$ in $\mathbb{G}_1^n$ or $\mathbb{G}_2^n$. Lai et al. further extended this approach to commitments to mixed vectors in $\mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1}$ or, analogously, $\mathbb{Z}_q^{n_0} \times \mathbb{G}_2^{n_2}$ [LMR19]. Definition 5.2 formalizes this commitment scheme.

Definition 5.2 (Extended Pedersen Commitment Scheme [AFG+10; LMR19])**.** The following setup algorithm and commitment function define a pairing-based extension of the Pedersen Vector commitment scheme:

- Setup:
$$\mathsf{pk} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e, \mathbf{g}, \mathbf{h}, h) \leftarrow \textsc{Setup}(1^\lambda, n_0, n_1)\,,$$
where $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e) \leftarrow \mathcal{G}(1^\lambda)$ for a bilinear group generator $\mathcal{G}(\cdot)$ and $(\mathbf{g}, \mathbf{h}, h) \leftarrow_R \mathbb{G}_2^{n_1} \times \mathbb{H}^{n_0} \times \mathbb{H}$ are sampled uniformly at random.

- Commitment Function:
$$\textsc{com}_{\mathsf{pk}}\colon \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{Z}_q \to \mathbb{H}, \quad (\mathbf{x}, \mathbf{y}; \gamma) \mapsto \mathbf{h}^{\mathbf{x}} \cdot e(\mathbf{y}, \mathbf{g}) \cdot h^\gamma\,,$$
where $\mathbf{h}^{\mathbf{x}} := \prod_{i=1}^{n_0} h_i^{x_i}$ and $e(\mathbf{y}, \mathbf{g}) := \prod_{i=1}^{n_1} e(y_i, g_i)$.

This commitment scheme is perfectly hiding and computationally binding under the *double pairing assumption*. Informally, this assumption states that it is hard to find elements $r_1, r_2 \in \mathbb{G}_1$ such that $e(r_1, g_1)e(r_2, g_2) = 1$ for random $g_1, g_2 \in \mathbb{G}_2$. Abe et al. showed that the double pairing assumption is implied by the decisional Diffie-Hellman (DDH) assumption in $\mathbb{G}_2$ [AFG+10]. Therefore, the above commitment scheme is computationally binding under the DDH assumption in $\mathbb{G}_2$.

Note that the double pairing assumption does not hold in *symmetric* bilinear groups, i.e., when $\mathbb{G}_1 = \mathbb{G}_2$. Namely, in this case $e(-g_2, g_1)e(g_1, g_2) = 1$ for all $g_1, g_2 \in (\mathbb{G}_2, +)$. Similarly, it is easily seen that the DDH assumption does not hold in $\mathbb{G}_2$ if there exists a pairing $e\colon \mathbb{G}_2 \times \mathbb{G}_2 \to \mathbb{H}$. For this reason, we require the bilinear group to be asymmetrical. If the DDH assumption holds in both $\mathbb{G}_1$ and $\mathbb{G}_2$, we also say that the *symmetrical external Diffie-Hellman* (SXDH) assumption holds.

Abe et al. observed that the commitment scheme of Definition 5.2 introduces an alternative for Pedersen commitments to vectors of field elements [AFG+10]. Namely, a commitment to n different n-dimensional Pedersen commitments is a commitment to an n^2-dimensional $\mathbb{Z}_q$-vector. This two-tiered commitment scheme only requires $2n + 1$ public group elements. By contrast, Pedersen's commitment scheme requires $n^2 + 1$ public group elements to commit to an n^2-dimensional $\mathbb{Z}_q$-vector. Replacing the Pedersen vector commitment scheme in Theorem 5.1 by this two-tiered approach results in a compressed Σ-protocol with exactly the same

communication costs, but with a square root improvement in the size of the public parameters.

In addition, Lai et al. show how this approach can be extended to construct a commitment scheme for vectors with coefficients in $\mathbb{Z}_q$, $\mathbb{G}_1$ and $\mathbb{G}_2$ [LMR19]. In contrast to previous schemes, a commitment to a vector $\mathbf{x} \in \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2}$ consists of two elements in the group $\mathbb{H}$. The reason is that $(x, y) = (g_1, -g_2)$ is a nontrivial solution for the equation $e(x, g_2)e(g_1, y) = 1$ for any $(g_1, g_2) \in \mathbb{G}_1 \times \mathbb{G}_2$. Such a solution would break the binding property of the naive generalization of Definition 5.2 in which commitments consist of only one target group element. However, with high probability, there does not exist a solution $(x, y) \in \mathbb{G}_1 \times \mathbb{G}_2$ to the system of equations $e(x, g_2)e(g_1, y) = 1$ and $e(x, g_2')e(g_1', y) = 1$, where $(g_1, g_2), (g_1', g_2') \in \mathbb{G}_1 \times \mathbb{G}_2$ are sampled uniformly at random. For this reason, the commitments consist of two target group elements and, under the SXDH assumption, breaking their binding property can be reduced to solving a similar system of equations. The resulting commitment scheme is described in Definition 5.3. It is computationally hiding under the DDH assumption in $\mathbb{G}_T$, and it is computationally binding under the SXDH assumption. The scheme can be made perfectly hiding by introducing an additional randomizer $\gamma_2 \in \mathbb{Z}_q$.

Definition 5.3 (Compact Commitments to $(\mathbb{Z}_q, \mathbb{G}_1, \mathbb{G}_2)$-Vectors [LMR19])**.** The following setup algorithm and commitment function define a pairing-based extension of the Pedersen Vector commitment scheme:

- Setup:

$$\mathsf{pk} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e, \mathbf{g}_1, \mathbf{g}_1', \mathbf{g}_2, \mathbf{g}_2', \mathbf{h}, \mathbf{h}', h, h') \leftarrow \textsc{Setup}(1^\lambda, n_0, n_1, n_2)\,,$$

  where $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e) \leftarrow \mathcal{G}(1^\lambda)$ for a bilinear group generator $\mathcal{G}(\cdot)$ and $(\mathbf{g}_1, \mathbf{g}_1', \mathbf{g}_2, \mathbf{g}_2', \mathbf{h}, \mathbf{h}', h, h') \leftarrow_R \mathbb{G}_1^{2n_2} \times \mathbb{G}_2^{2n_1} \times \mathbb{H}^{2n_0+2}$ are sampled uniformly at random.

- Commitment Function: $\textsc{com}_{\mathsf{pk}} \colon \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2} \times \mathbb{Z}_q^2 \to \mathbb{H}$,

$$(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2; \gamma) \mapsto \begin{pmatrix} \mathbf{h}^{\mathbf{x}_0} \cdot e(\mathbf{x}_1, \mathbf{g}_2) \cdot e(\mathbf{g}_1, \mathbf{x}_2) \cdot h_1^{\gamma} \\ \mathbf{h}'^{\mathbf{x}_0} \cdot e(\mathbf{x}_1, \mathbf{g}_2') \cdot e(\mathbf{g}_1', \mathbf{x}_2) \cdot h_1'^{\gamma} \end{pmatrix}.$$

The commitment scheme of Definition 5.3 is a homomorphic and compact commitment scheme for mixed vectors $\mathbf{x} \in \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2}$. However, it does not allow a prover to commit to elements of the target group $\mathbb{H}$ of the pairing $e \colon \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{H}$. Unfortunately, we do not know how to do this compactly while preserving the required homomorphic properties. For this reason, we introduce the homomorphic commitment scheme of Definition 5.4. This scheme is based on the ElGamal encryption scheme [ElG84]. The commitment scheme is unconditionally binding and computationally hiding under the DDH assumption in $\mathbb{G}_T$. However, in contrast to the previous commitment schemes, it is not compact. More precisely, an ElGamal commitment to a vector $\mathbf{x}_T \in \mathbb{H}^{n_T}$ contains $n_T + 1$ group elements.

Definition 5.4 (ElGamal Commitment Scheme)**.** The ElGamal vector commitment scheme is defined by the following setup algorithm and commitment function:

- $\mathsf{pk} = (q, \mathbb{H}, G_1, \ldots, G_{n_T}, H) \leftarrow \textsc{Setup}(1^\lambda, n_T)$, where $(q, \mathbb{H}, \cdot) \leftarrow \mathcal{G}(1^\lambda)$ for a prime order group generator $\mathcal{G}(\cdot)$ and $(\mathbf{G}, H) = (G_1, \ldots, G_{n_T}, H) \leftarrow_R \mathbb{H}^{n_T+1}$ are sampled uniformly at random;
- $\textsc{com}_{\mathsf{pk}} \colon \mathbb{H}^{n_T} \times \mathbb{Z}_q \to \mathbb{H}^{N_T+1}, \quad (\mathbf{x}_T; \rho) \mapsto \begin{pmatrix} H^\rho \\ \mathbf{G}^\rho * \mathbf{x}_T \end{pmatrix}$, where $\mathbf{G}^\rho := (G_1^\rho, \ldots, G_{n_T}^\rho)$ and $*$ denotes the component-wise product.

Combined, the commitment schemes of Definition 5.3 and Definition 5.4 provide a homomorphic commitment scheme for *bilinear group vectors*

$$\mathbf{x} \in \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2} \times \mathbb{H}^{n_T} .$$

This commitment scheme is only compact in the dimension n_0, n_1 and n_2; the size of a commitment is linear in the dimension n_T of the $\mathbb{H}$-component. For completeness we have included the definition of the resulting commitment scheme for bilinear group vectors.

Definition 5.5 (Bilinear Group Vector Commitment Scheme [LMR19])**.** The following setup algorithm and commitment function define a bilinear group vector commitment scheme:

- Setup:

$$\mathsf{pk} = \begin{pmatrix} q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e, \mathbf{g}_1, \mathbf{g}_1', \mathbf{g}_2, \\ \mathbf{g}_2', \mathbf{h}, \mathbf{h}', \mathbf{G}, h, h', H \end{pmatrix} \leftarrow \textsc{Setup}(1^\lambda, n_0, n_1, n_2, n_T) ,$$

  where $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e) \leftarrow \mathcal{G}(1^\lambda)$ for a bilinear group generator $\mathcal{G}(\cdot)$ and

$$(\mathbf{g}_1, \mathbf{g}_1', \mathbf{g}_2, \mathbf{g}_2', \mathbf{h}, \mathbf{h}', \mathbf{G}, h, h', H) \leftarrow_R \mathbb{G}_1^{2n_2} \times \mathbb{G}_2^{2n_1} \times \mathbb{H}^{2n_0+n_T+3}$$

  are sampled uniformly at random.

- Commitment Function: $\textsc{com}_{\mathsf{pk}} \colon \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2} \times \mathbb{H}^{n_T} \times \mathbb{Z}_q^2 \to \mathbb{H}^{n_T+3}$,

$$(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_T; \gamma, \rho) \mapsto \begin{pmatrix} \mathbf{h}^{\mathbf{x}_0} \cdot e(\mathbf{x}_1, \mathbf{g}_2) \cdot e(\mathbf{g}_1, \mathbf{x}_2) \cdot h_1^\gamma \\ \mathbf{h}'^{\mathbf{x}_0} \cdot e(\mathbf{x}_1, \mathbf{g}_2') \cdot e(\mathbf{g}_1', \mathbf{x}_2) \cdot h_1'^\gamma \\ H^\rho \\ \mathbf{G}^\rho * \mathbf{x} \end{pmatrix} .$$

A compressed Σ-protocol, instantiated with the above bilinear group vector commitment scheme, allows a prover to prove knowledge of a commitment opening satisfying a linear constraint $L(\mathbf{x}) = \mathbf{y}$ captured by a linear mapping

$$L \colon \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2} \times \mathbb{H}^{n_T} \to \mathbb{Z}_q \times \mathbb{G}_1 \times \mathbb{G}_2 \times \mathbb{H} .$$

As before, we apply the compressed Σ-protocol of Theorem 3.11, where the linear form evaluations are incorporated into the commitment. Note that, since the commitment scheme is only compact in its $(\mathbb{Z}_q, \mathbb{G}_1, \mathbb{G}_2)$-part, only the $(\mathbb{Z}_q, \mathbb{G}_1, \mathbb{G}_2)$-part of the L-evaluation should be incorporated into the commitment. For the

same reason, compression is only applied to the $(\mathbb{Z}_q, \mathbb{G}_1, \mathbb{G}_2)$-part of the committed vector. Theorem 5.2 summarizes the main properties of the compressed Σ-protocol for bilinear group vectors. For simplicity, we assume that $n_0 + 1 = n_1 = n_2$, but the result is easily extended to arbitrary input dimensions. Note that the communication complexity of this compressed Σ-protocol is logarithmic in n_0, n_1 and n_2, but linear in n_T.

Theorem 5.2 (Compressed Σ-Protocol for Bilinear Group Vectors)**.** *Let $n_0 + 1 = n_1 = n_2 = 2^\mu$ for some $\mu \in \mathbb{N}$, $n_T \in \mathbb{N}$, COM_{pk} the bilinear group vector commitment function instantiated with the bilinear group $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{H}, e)$ and $L\colon \mathbb{Z}_q^{n_0} \times \mathbb{G}_1^{n_1} \times \mathbb{G}_2^{n_2} \times \mathbb{H}^{n_T} \to \mathbb{Z}_q \times \mathbb{G}_1 \times \mathbb{G}_2 \times \mathbb{H}$ linear.*

Then the compressed Σ-protocol for relation

$$\mathfrak{R}_{Bil} = \{(P, \mathbf{y}; \mathbf{x}, \gamma, \rho) : \text{COM}_{pk}(\mathbf{x}; \gamma, \rho) = P \wedge L(\mathbf{x}) = \mathbf{y}\},$$

is perfectly complete, computationally $(2, 2, 3, \ldots, 3)$-out-of-$(q, \ldots, q)$ special-sound, under the symmetrical external Diffie-Hellman (SXDH) assumption, and special honest-verifier zero-knowledge (SHVZK). Moreover, it has $(2\mu + 2)$ communication rounds and the communication costs are:

- *$\mathcal{P} \to \mathcal{V}$: 3 elements of $\mathbb{Z}_q$, 2 elements of $\mathbb{G}_1$, 2 elements of $\mathbb{G}_2$ and $6\mu + 2n_T - 3$ elements of $\mathbb{H}$;*
- *$\mathcal{V} \to \mathcal{P}$: $\mu + 1$ elements of $\mathbb{Z}_q$.*

5.4 Knowledge of Exponent Assumption

If one desires, the functionality of opening linear forms on compactly committed vectors can also be achieved from the *Knowledge-of-Exponent Assumption* (KEA). In order to introduce this assumption, let $\mathbb{H}$ be a group of prime order q and let us consider the following problem: on input $g, h = g^a \in \mathbb{H}$ output a pair $G, H \in \mathbb{H}$ with $G = H^a$. A simple solution to this problem is to output $G = g^c$ and $H = h^c$ for an arbitrary $c \in \mathbb{Z}_q$. Informally, the KEA states that this is the *only* way to solve this problem. More precisely, for any adversary that successfully outputs a pair (G, H) there exists an extractor that outputs the exponent c such that $G = g^c$ and thus $H = h^c$. We stress that the KEA is of a different nature than the discrete logarithm or decisional Diffie-Hellman assumption. KEA is not an intractability assumption and it is unfalsifiable [Nao03; BCP+14]. For these reasons, its application is not completely without controversy.

Opening linear forms on compact commitments instantiated from the KEA does not proceed by the standard compression paradigm. Namely, the basic protocol for this functionality already has *constant* communication complexity, i.e., compression is not needed. However, since the techniques of Section 7.2 only require *black-box* access to a protocol for opening linear forms on compactly committed vectors, they are equally applicable to a KEA instantiation. For this reason, we present the KEA approach here, even though it is not an instantiation of the compressed Σ-protocols of Chapter 3. Basing the linear form openings on the KEA results in *constant* communication complexity instead of logarithmic. However,

the resulting protocol does require a trusted setup. Below, we will elaborate on this trusted setup requirement.

We now describe the KEA based vector commitment scheme together with its protocol for opening linear forms. Our approach uses the techniques of [Gro10] and only minor adaptations are required.

A compact commitment to a vector $\mathbf{x} \in \mathbb{Z}_q^n$ is, as before, a Pedersen vector commitment $P = \mathbf{g}^{(\gamma,\mathbf{x})} := g_0^\gamma \prod_{i=1}^n g_i^{x_i}$. To prove knowledge of a commitment opening of P, the prover simply sends another Pedersen commitment $Q = \mathbf{h}^{(\gamma,\mathbf{x})}$ to $\mathbf{x}$, under the same randomness γ, using a different vector of group elements $\mathbf{h} = \mathbf{g}^\alpha = (g_0^\alpha, \ldots, g_n^\alpha) \in \mathbb{H}^{n+1}$. The value $\alpha \in \mathbb{Z}_q$ is sampled uniformly at random by a *trusted* party and is only shared with a *designated* verifier. Both vectors of groups elements are public. The proof Q is verified by checking that $Q = P^\alpha$, i.e., only a designated verifier that knows the secret value α can verify a proof. It is crucial that the prover does not know α, otherwise it can simply forge a proof by computing $Q = P^\alpha$.

The knowledge-of-exponent assumption states that an adversary capable of computing pairs (P, Q) with $Q = P^\alpha$, either knows α or an opening to P. From this assumption knowledge soundness follows. Correctness follows immediately and zero-knowledge follows since the proof Q is uniquely determined by P and α. In fact, the verifier can compute the proof $Q = P^\alpha$ without knowledge of $\mathbf{x}$. Hence, the proof Q does not reveal any additional information about the witness $\mathbf{x}$. Note that the resulting protocol only has one round, i.e., it is non-interactive, and its communication costs are independent of the dimension n.

Given a bilinear pairing $e\colon \mathbb{H} \times \mathbb{H} \to \mathbb{H}_T$, the verification procedure can be made public, i.e., given e, even parties that do not know α can verify a proof. In this case verification amounts to checking that $e(P, h_0) = e(g_0, Q)$.

If during the setup phase, the prover is only given the group elements h_0 and h_i for $i \in S \subseteq \{1, \ldots, n\}$, then the proof Q can only be computed if $x_i = 0$ for all $i \notin S$. Groth [Gro10] refers to the resulting proof as a *restriction* proof, since it actually shows that the nonzero entries of the committed vector $\mathbf{x}$ are restricted to the subset S of indices. The restriction proof is an important building block of our KEA-based protocol for opening linear forms on Pedersen commitments. Therefore, it is described in Protocol 14.

To additionally prove that the committed vector $\mathbf{x}$ satisfies the linear constraint $L(\mathbf{x}) = y$ some adaptations are required. More precisely, in this case, the group elements are sampled under the condition that $\mathbf{g} = (g, g^\beta, \ldots, g^{\beta^n})$ for some secret $\beta \in \mathbb{Z}_q$. The KEA that takes this additional structure into account is called the n-power Knowledge-of-Exponent Assumption (n-PKEA).

Groth [Gro10] showed that, using this additional structure, efficient circuit zero-knowledge protocols exist, i.e., protocols for proving knowledge of a secret vector $\mathbf{x} \in \mathbb{Z}_q^n$ such that $C(\mathbf{x}) = 0$ for some arbitrary arithmetic circuit C. Note that an arithmetic circuit constraint $C(\mathbf{x}) = 0$ is not necessarily linear. Groth's protocols can be adapted to our situation, where we simply wish to prove the validity of a *linear* constraint $L(\mathbf{x}) = y$ for some linear form $L\colon \mathbb{Z}_q^n \to \mathbb{Z}_q$. In Section 7.2, we will show how to handle nonlinear instances.

The adaptation of Groth's protocol relies on the following observation. Suppose that $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{Z}_q^n$ is such that $L(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle$ for all $\mathbf{x} \in \mathbb{Z}_q^n$, and let us

Protocol 14 KEA Restriction Proof for Pedersen Commitments.

PARAMETERS:	$n \in \mathbb{N}$, group $(\mathbb{H}, \cdot)$ of prime order q, pairing $e\colon \mathbb{H} \times \mathbb{H} \to \mathbb{H}_T$, $\mathbf{g} \in \mathbb{H}^{n+1}$, $h_0 = g_0^{\alpha}$, $h_i = g_i^{\alpha}$ for $i \in S \subseteq \{1, \ldots, n\}$ for a (secret) $\alpha \in \mathbb{Z}_q$
PUBLIC INPUT:	$P \in \mathbb{H}$
PROVER'S PRIVATE INPUT:	$\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_q^n$
PROVER'S CLAIM:	$\mathbf{g}^{(\gamma, \mathbf{x})} = g_0^{\gamma} \prod_{i=1}^{n} g_i^{x_i} = P \wedge x_i = 0 \; \forall i \notin S$

Prover $\mathcal{P}$		Verifier $\mathcal{V}$
$Q = h_0^{\gamma} \prod_{i \in S} h_i^{x_i}$	$\xrightarrow{\quad Q \quad}$	$e(P, h_0) \stackrel{?}{=} e(g_0, Q)$

define the following polynomials:

$$F(Y) = \gamma + \sum_{i=1}^{n} x_i Y^i\,, \quad G(Y) = \sum_{i=0}^{n-1} a_{n-i} Y^i\,,$$
$$\text{and} \quad H(Y) = F(Y)G(Y) = \sum_{i=0}^{2n-1} c_i Y^i\,.$$

Then the $(n+1)$-th coefficient of $H(Y)$ equals $c_n = \langle \mathbf{x}, \mathbf{a} \rangle = L(\mathbf{x})$. Moreover, since $\mathbf{g} = (g, g^{\beta}, \ldots, g^{\beta^n})$ for some secret $\alpha, \beta \in \mathbb{Z}_q$,

$$P = \mathbf{g}^{(\gamma, \mathbf{x})} = g^{F(\beta)}, \quad R := \mathbf{g}^{(a_n \ldots, a_1, 0)} = g^{G(\beta)} \quad \text{and} \quad e(P, R) = e\big(g, g^{H(\beta)}\big)\,.$$

Since $\mathbf{a}$ is public, both the prover and the verifier can compute the group element R. Hence, to prove that $L(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle = y$, the prover must convince the verifier that

$$e(P, R) = e\big(g, \mathbf{g}^{(c_0, \ldots, c_{n-1}, y, c_{n+1}, \ldots, c_{2n-1})}\big)\,, \tag{5.1}$$

for some $c_i \in \mathbb{Z}_q$. Note that we make some abuse of notation by implicitly assuming the vector $\mathbf{g}$ to be long enough, i.e., $\mathbf{g} = (g, g^{\beta}, \ldots, g^{\beta^{2n-1}}) \in \mathbb{H}^{2n}$.

To prove the validity of Equation 5.1, the prover sends the group element $S = \prod_{i \neq n} g_i^{c_i}$, where $c_0, \ldots, c_{2n-1}$ are the coefficients of the polynomial $H(Y)$. Subsequently, the verifier checks that

$$e(P, R) = e\big(g, S \cdot g_n^{y}\big)\,.$$

The proof is completed by adding the following group elements:

- A commitment opening proof Q for P, proving knowledge of an opening $(\mathbf{x}; \gamma) \in \mathbb{Z}_q^{n+1}$ of P;
- A restriction proof T for S, showing that the exponent vector $(c_0, \ldots, c_{2n-1})$ of S is zero in its $(n+1)$-th coordinate.

Note that the element Q is in fact a restriction proof; it shows that $\mathbf{x}$ is an n-dimensional vector, i.e., the commitment P does not make use of the public group elements

$$g^{\beta^{n+1}}, \ldots, g^{\beta^{2n-1}} \in \mathbb{H}\,.$$

The KEA based non-interactive proof for opening linear forms on Pedersen commitments is described in Protocol 15. It is crucial that the prover does not know the secret values $\alpha, \alpha', \beta \in \mathbb{Z}_q$. Therefore, these values must be generated in a trusted setup phase. The size of the proof is independent of the dimension n of the committed vector. This non-interactive proof is an adaptation of Groth's product argument [Gro10, Section 6]. Its (security) analysis requires somewhat different techniques and formalization then the ones used before. For this reason, we refer to [Gro10] for a more formal analysis.

Protocol 15 KEA Protocol for Opening Linear Forms.

Parameters:	$n \in \mathbb{N}$, group $(\mathbb{H}, \cdot)$ of prime order q, pairing $e\colon \mathbb{H} \times \mathbb{H} \to \mathbb{H}_T$ and vectors of $\mathbb{H}$-elements $\mathbf{g} = (g_0, \ldots, g_{2n-1}) = (g, g^{\beta}, \ldots, g^{\beta^{2n-1}})$, $\mathbf{k} = (g_0^{\alpha'}, \ldots, g_{n-1}^{\alpha'}, 1, g_{n+1}^{\alpha'}, \ldots, g_{2n-1}^{\alpha'})$ and $\mathbf{h} = (g_0^{\alpha}, \ldots, g_n^{\alpha})$ for (secret) $\alpha, \alpha', \beta \in \mathbb{Z}_q$
Public Input:	$P \in \mathbb{H}$, $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{Z}_q$ and $y \in \mathbb{Z}_q$
Prover's Private Input:	$\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_q^n$
Prover's Claim:	$g_0^{\gamma} \prod_{i=1}^{n} g_i^{x_i} = P \wedge L(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle = y$

Prover $\mathcal{P}$		Verifier $\mathcal{V}$
$F(Y) = \gamma + x_1 Y + \cdots + x_n Y^n$		
$G(Y) = a_n + a_{n-1} Y + \cdots + a_1 Y^{n-1}$		
$F(Y)G(Y) = c_0 + \cdots + c_{2n-1} Y^{2n-1}$		
$Q = \mathbf{h}^{(\gamma, \mathbf{x})}$		
$S = \prod_{i \neq n} g_i^{c_i}$		$R = \mathbf{g}^{(a_n, \ldots, a_1, 0)}$
$T = \prod_{i \neq n} k_i^{c_i}$		
	$\xrightarrow{Q,S,T}$	
		$e(P, h_0) \stackrel{?}{=} e(g_0, Q)$
		$e(S, k_0) \stackrel{?}{=} e(g_0, T)$
		$e(P, R) \stackrel{?}{=} e(g, S \cdot g_n^y)$

5.5 Strong-RSA Assumption

Let us now move to a compressed Σ-protocol instantiation based on the assumption that a dishonest prover does not know the order of some given group. More precisely, its security is based on the strong-RSA assumption (Definition 2.16). This instantiation is inspired by the strong-RSA based polynomial commitment scheme DARK [BFS20]. A polynomial commitment scheme allows a prover to commit to a polynomial $f \in \mathbb{Z}_q[X]$ of arbitrary degree and admits a protocol for "opening polynomial evaluations," i.e., a protocol for proving that a committed polynomial f satisfies $f(x) = y$ for some public $x, y \in \mathbb{Z}_q$. DARK is a strong-RSA based adaptation of the Bulletproof protocol [BCC+16; BBB+18], and it allows a prover to open polynomial evaluations with logarithmic communication complexity. However, Block et al. [BHR+21] identified a gap in the security analysis of DARK. Fortunately, they also proposed an adaptation of DARK, solving the aforementioned security gap at the cost of increasing the communication complexity from logarithmic to polylogarithmic.

Note that a polynomial $f(X) = \sum_{i=0}^{n} a_i X^i$ is uniquely defined by its coefficient vector and an evaluation of a polynomial is a special type of linear form evaluation, i.e.,

$$f(x) = \langle (a_0, \ldots, a_n), (1, x, \ldots, x^n) \rangle \,.$$

Hence, the functionality of a polynomial commitment scheme is strictly weaker than "opening linear forms on compactly committed vectors." Some of the techniques introduced in DARK [BFS20] and its adaptation [BHR+21] crucially depend on the structure of linear forms corresponding to polynomial evaluations, and are therefore not applicable to opening arbitrary linear forms. For this reason, we must modify the aforementioned approaches.

An important building block in these strong-RSA based interactive proofs is the following *integer* commitment scheme. To simplify the exposition, and in order to focus on the important aspects, we consider a *non-hiding* variant. For a statistically hiding variant of this commitment scheme we refer the reader to [DF02].

Definition 5.6 (Non-Hiding Integer Commitment Scheme [FO97; DF02])**.** The following setup algorithm and commitment function define a non-hiding integer commitment scheme:

- $\mathsf{pk} = (\mathbb{H}, g) \leftarrow \textsc{Setup}(1^\lambda)$, where $(\mathbb{H}, \cdot) \leftarrow \mathcal{G}(1^\lambda)$ for a hidden-order group generator $\mathcal{G}(\cdot)$ and $g \leftarrow_R \mathbb{H}$ is sampled uniformly at random;
- $\textsc{com}_{\mathsf{pk}} \colon \mathbb{Z} \to \mathbb{H}, \quad x \mapsto g^x$.

The commitment scheme of Definition 5.6 is homomorphic and computationally binding under the hidden order assumption (Definition 2.17), which is implied by the Strong-RSA assumption.

By appropriately encoding vectors of integers $\mathbf{x} \in \mathbb{Z}^n$, this commitment scheme allows a prover to commit to vectors of *bounded* integers. More precisely, let $\mathbb{Z}(\alpha) = \{x \in \mathbb{Z} : |x| < \alpha\}$ and

$$\mathrm{Enc}_Q \colon \mathbb{Z}(\alpha)^n \to \mathbb{Z}, \quad (x_1, \ldots, x_n) \mapsto \sum_{i=1}^{n} x_i Q^{i-1} \,,$$

for some integer $Q \geq 2\alpha$. Since $Q \geq 2\alpha$, this encoding is injective. Moreover, base Q decomposition provides an efficient decoding algorithm Dec_Q. A commitment to a bounded integer vector $\mathbf{x} \in \mathbb{Z}(\alpha)^n$ is simply an integer commitment to $\mathrm{Enc}_Q(\mathbf{x}) \in \mathbb{Z}$. Definition 5.7 formalizes this commitment scheme. By the injectivity of Enc_Q, this commitment scheme is computationally binding under the strong-RSA assumption.

Definition 5.7 (Non-Hiding Bounded Integer Vector Commitment Scheme)**.** The following setup algorithm and commitment function define a non-hiding bounded integer vector commitment scheme:

- $\mathsf{pk} = (\mathbb{H}, g) \leftarrow \textsc{Setup}(1^\lambda)$, where $(\mathbb{H}, \cdot) \leftarrow \mathcal{G}(1^\lambda)$ for a hidden-order group generator $\mathcal{G}(\cdot)$ and $g \leftarrow_R \mathbb{H}$ is sampled uniformly at random;
- $\textsc{com}_{\mathsf{pk}} \colon \mathbb{Z}(\alpha)^n \to \mathbb{H}, \quad \mathbf{x} \mapsto g^{\mathrm{Enc}_Q(\mathbf{x})}$, where $\mathrm{Enc}_Q(\mathbf{x}) = \sum_{i=1}^n x_i Q^{i-1}$ for some $Q \geq 2\alpha$.

Our goal is to construct an interactive proof for proving knowledge of an opening $\mathbf{x} \in \mathbb{Z}(\alpha)^n$ of the commitment $P \in \mathbb{H}$ satisfying the linear constraint $L(\mathbf{x}) = y$ for some linear form $L \colon \mathbb{Z}^n \to \mathbb{Z}$, i.e., for proving knowledge of a short Ψ_Q-preimage, where

$$\Psi_Q \colon \mathbb{Z}^n \to \mathbb{H} \times \mathbb{Z}, \quad \mathbf{x} \mapsto \left(g^{\mathrm{Enc}_Q(\mathbf{x})}, L(\mathbf{x})\right).$$

The interactive proofs of Section 3.3 have soundness slack τ and approximation factor ζ, i.e., they allow a prover to prove knowledge of a Ψ_Q-preimage $\mathbf{x}'$ of $(P^\zeta, \zeta y) \in \mathbb{H} \times \mathbb{Z}$ with $\|\mathbf{x}'\|_\infty \leq \tau\alpha$. Oftentimes, this relaxation is acceptable as long as the commitment scheme is also binding with respect to (τ, ζ)-*relaxed* openings. More precisely, given a commitment P, it should be hard for a prover to find distinct openings $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}(\tau\alpha)^n$ of P^ζ. Concretely, this means that the encoding should be instantiated such that $Q \geq 2\tau\alpha$ instead of $Q \geq 2\alpha$. Hence, the soundness slack τ directly influences the efficiency of the interactive proof and should thus be kept to a minimum.

Further, the mapping Ψ_Q is a $\mathbb{Z}$-module homomorphism. For this reason, instantiating the compression mechanism of Section 3.3.2 directly, requires a challenge set $\mathcal{C} \subseteq \mathbb{Z}$. This either leaves us with a small challenge set, e.g., $\mathcal{C} = \{-1, 0, 1\}$, or with a large soundness slack τ. For instance, challenge sets of the form $\mathcal{C} = \mathbb{Z}(B) = \{x \in \mathbb{Z} : |x| < B\}$ result in a soundness slack that grows *exponentially* in B. For this reason, we first apply the base extension techniques of Section 3.3.4. More precisely, we extend the base $\mathbb{Z}$ of the $\mathbb{Z}$-module homomorphism Ψ_Q to the $2d$-th cyclotomic number ring $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ for $d = 2^{d'}$ a power of two, i.e., we consider the $\mathcal{R}$-module homomorphism

$$\Psi_{Q,\mathcal{R}} \colon \mathcal{R} \otimes_{\mathbb{Z}} \mathbb{Z}^n \to \mathcal{R} \otimes_{\mathbb{Z}} (\mathbb{H} \times \mathbb{Z}), \quad \text{such that} \quad r \otimes \mathbf{x} \mapsto r \otimes \Psi_Q(\mathbf{x}).$$

Moreover, via the $\mathbb{Z}$-basis $\{1, \ldots, X^{d-1}\}$ of $\mathcal{R}$, we define the following ℓ_∞-norm on $\mathcal{R} \otimes_{\mathbb{Z}} \mathbb{Z} \cong \mathcal{R}$:

$$\left\|1 \otimes x_1 + X \otimes x_2 + \cdots + X^{d-1} \otimes x_d\right\|_\infty = \max_{1 \leq i \leq d} |x_i| \,,$$

where $x_i \in \mathbb{Z}$ for all i. This norm has a natural extension to $\mathcal{R} \otimes_{\mathbb{Z}} \mathbb{Z}^n$, i.e.,

$$\left\|1 \otimes \mathbf{x}_1 + X \otimes \mathbf{x}_2 + \cdots + X^{d-1} \otimes \mathbf{x}_d\right\|_\infty = \max_{1 \le i \le d} \|\mathbf{x}_i\|_\infty \,.$$

The reason for extending the base to a power-of-two cyclotomic number ring is that these rings contain challenge sets resulting in small soundness slack. To see this, we recall the following lemma by Benhamouda et al. [BCK+14].

Lemma 5.1 (Lemma 3.1 of [BCK+14])**.** *Let $d = 2^{d'} \in \mathbb{N}$ be a power of two and let $\mathbb{Z}[X]/(X^d+1)$ be the $2d$-th cyclotomic number ring. Then, for all $i \neq j$*

$$\frac{2}{X^i - X^j} \in \mathbb{Z}[X]/(X^d+1)\,.$$

Moreover, this polynomial only has coefficients in $\{-1, 0, 1\}$.

Proof. Without loss of generality, we may assume that $i = 0$ and $1 < j < 2d$. Now let k be the smallest positive integer such that $kj = 0 \mod d$. Since d is a power of two and $j \neq 0 \mod 2d$, it holds that $kj \neq 0 \mod 2d$ and $X^{kj} = -1$. Therefore,

$$\frac{2}{X^i - X^j} = \frac{2}{1 - X^j} = \frac{2}{1 - X^{kj}} \cdot (1 + X^j + X^{2j} + \cdots X^{(k-1)j}) = 1 + X^j + \cdots X^{(k-1)j}\,,$$

which proves the first claim of the lemma.

What remains to show is that no two exponents ℓj and $\ell' j$, for $0 \le \ell < \ell' < k$, are the same modulo d. Assuming the contrary, it follows that $(\ell' - \ell)j = 0 \mod d$ with $0 < \ell' - \ell < k$. This contradicts the assumption that k is the smallest positive integer such that $kj = 0 \mod d$ and completes the proof. □

Lemma 5.1 shows that the challenge set $\mathcal{C} = \{0, \pm 1, \pm X, \ldots, \pm X^{d-1}\}$ is a 2-exceptional subset of the power-of-two cyclotomic number ring $\mathcal{R} = \mathbb{Z}[X]/(X^d+1)$. Moreover, it immediately implies the following corollary.

Corollary 5.1. *Let $d = 2^{d'} \in \mathbb{N}$ be a power of two and let $\mathcal{R} = \mathbb{Z}[X]/(X^d+1)$ be the $2d$-th cyclotomic number ring. Further, let $\mathcal{C} = \{0, \pm 1, \pm X, \ldots, \pm X^{d-1}\} \subset \mathcal{R}$. Then,*

$$w(\mathcal{C}) = \max_{c \in \mathcal{C}, x \in \mathcal{R} \otimes_{\mathbb{Z}} \mathbb{Z} \setminus \{0\}} \frac{\|cx\|_\infty}{\|x\|_\infty} = 1\,,$$

$$\overline{w}(\mathcal{C}, 2) = \max_{c \neq c' \in \mathcal{C}, x \in \mathcal{R} \otimes_{\mathbb{Z}} \mathbb{Z} \setminus \{0\}} \frac{\left\|2(c - c')^{-1} x\right\|_\infty}{\|x\|_\infty} = d\,.$$

The following theorem summarizes the properties of the compression mechanism of Section 3.3.2 instantiated for the homomorphism

$$\Psi_{Q,\mathcal{R}} \colon \mathcal{R} \otimes_{\mathbb{Z}} \mathbb{Z}^n \to \mathcal{R} \otimes_{\mathbb{Z}} (\mathbb{H} \times \mathbb{Z})\,, \quad \text{such that} \quad r \otimes \mathbf{x} \mapsto r \otimes \Psi_Q(\mathbf{x})\,,$$

with challenge set $\mathcal{C} = \{0, \pm 1, \pm X, \ldots, \pm X^{d-1}\} \subset \mathcal{R} = \mathbb{Z}[X]/(X^d+1)$. The theorem is a direct consequence of Theorem 3.8 and Corollary 5.1. It is valid for

all values of $Q \in \mathbb{N}$. However, the compression mechanism has soundness slack $12d^3$ and approximation factor 8. More precisely, while the prover claims to know a Ψ_Q-preimage of (P, y) with ℓ_∞-norm at most α, it is only capable of proving knowledge of a Ψ_Q-preimage of $(P^8, 8 \cdot y)$ with ℓ_∞-norm at most $12d^3\alpha$. Therefore, the vector commitment scheme should be instantiated with $Q \geq 24d^3\alpha$.

Theorem 5.3 (Strong-RSA Based Compression Mechanism)**.** *Let $\alpha, Q \in \mathbb{N}$, $n \in \mathbb{N}$ even, $d \in \mathbb{N}$ a power of two and $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$. Then the compression mechanism Π_c, described in Protocol 7, instantiated for the base-$\mathcal{R}$ extension $\Psi_{Q,\mathcal{R}}$ of the strong-RSA homomorphism*

$$\Psi_Q \colon \mathbb{Z}^n \to \mathbb{H} \times \mathbb{Z}, \quad \mathbf{x} \mapsto \left(g^{\mathrm{Enc}_Q(\mathbf{x})}, L(\mathbf{x})\right),$$

with challenge set $\mathcal{C} = \{0, \pm 1, \pm X, \ldots, \pm X^{d-1}\}$, is an interactive proof for relation

$$\mathfrak{R}_{\mathsf{RSA}} = \left\{(P, y, \alpha; \mathbf{x}) : \Psi_{Q,\mathcal{R}}(\mathbf{x}) = (P, y) \wedge \|\mathbf{x}\|_\infty \leq \alpha\right\}.$$

It is perfectly complete and 3-out-of-$(2d + 1)$ special-sound with soundness slack $12d^3$ and approximation factor 8. Moreover, the communication costs are:

- *$\mathcal{P} \to \mathcal{V}$: $\frac{dn}{2} + 2d$ elements of $\mathbb{Z}$ and $2d$ elements of $\mathbb{H}$;*
- *$\mathcal{V} \to \mathcal{P}$: 1 element of $\mathcal{C} \subseteq \mathcal{R}$.*

The following theorem now summarizes the properties of the μ-fold recursive composition of the strong-RSA based interactive proof of Theorem 5.3. As before this theorem holds for any Q, but to account for the soundness slack it should be instantiated with $Q \geq 2 \cdot 12^\mu \cdot d^{3\mu} \cdot \alpha$.

Theorem 5.4 (Recursive Strong-RSA Based Compression Mechanism)**.** *Let $n = 2^\mu \in \mathbb{N}$ be a power of two. Then, the μ-fold recursive composition of the strong-RSA compression mechanism of Theorem 5.3 is a $(2\mu+1)$-round interactive proof for relation $\mathfrak{R}_{\mathsf{RSA}}$. It is perfectly complete and $(3, \ldots, 3)$-out-of-$(2d+1, \ldots, 2d+1)$ special-sound with soundness slack $12^\mu \cdot d^{3\mu}$ and approximation factor 8^μ. Moreover, the communication costs are:*

- *$\mathcal{P} \to \mathcal{V}$: $d + 2d\log_2 n$ elements of $\mathbb{Z}$ and $2d\log_2 n$ elements of $\mathbb{H}$;*
- *$\mathcal{V} \to \mathcal{P}$: μ element of $\mathcal{C} = \{0, \pm 1, \pm X, \ldots, \pm X^{d-1}\} \subset \mathcal{R}$.*

If the interactive proof of Theorem 5.4 is instantiated with the degree d of the base extension equal to $\log_2 n$, its knowledge error is constant in n (see Section 3.3.4). Therefore, to reduce the knowledge error down to $2^{-\lambda}$, $t = \mathcal{O}(\lambda)$ parallel repetitions are required. The communication complexity of the resulting protocol, measured in the number of elements, is $\mathcal{O}(\lambda \cdot \log_2^2 n)$, i.e., it is polylogarithmic in n. Moreover, the soundness slack equals

$$12^\mu d^{3\mu} = n^{\log_2(\log_2 n)+2+\log_2 3},$$

i.e., it is subexponential in n. Taking $\alpha = (q-1)/2$ and $L \colon \mathbb{Z}^n \to \mathbb{Z}_q$ for some odd prime q, shows that this protocol allows a prover to commit to a vector $\mathbf{x} \in \mathbb{Z}_q^n$ and proves that it satisfies an arbitrary $\mathbb{Z}_q$-linear constraint.

An advantage of this strong-RSA based interactive proof, over the discrete logarithm instantiation of Section 5.2, is that the public key size of the underlying commitment scheme is constant in n. By contrast, a Pedersen commitment to an n-dimensional vector requires $n+1$ group elements, i.e., there the public key size is linear in n. However, note that this improvement comes at the cost of increasing the communication complexity from logarithmic to polylogarithmic.

Our approach differs from the polynomial commitment schemes of [BFS20] and [BHR+21]. Restricting to polynomial commitment schemes allows for an adaptation that reduces the verification complexity, measured in the number of group exponentiations, from quasilinear down to polylogarithmic in n. However, this adaptation requires the use of *proofs of exponentiation* [Wes19]. Moreover, our instantiation is unconditionally sound, whereas the aforementioned polynomial commitment schemes have conditional soundness based on the strong-RSA assumption.

5.6 A Lattice Assumption: Short Integer Solutions

The final compressed Σ-protocol instantiation that we shall discuss is based on a lattice assumption and therefore plausibly secure against quantum adversaries. More precisely, its security is based on the hardness of the *Module Short Integer Solution* (MSIS) problem (Definition 2.20). As before, our goal is to construct an efficient protocol for opening linear forms on compactly committed vectors.

Before we describe the underlying MSIS-based commitment scheme, we introduce some notation. Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(x) \in \mathbb{Z}[X]$ of degree d. For any $p \in \mathbb{Z}$, we write $\mathcal{R}_p = \mathcal{R}/p\mathcal{R}$. Moreover, we equip $\mathcal{R}$ with the following ℓ_∞-norm:

$$\left\| \sum_{i=0}^{d-1} a_i X^i \right\|_\infty = \max_{0 \le i \le d-1} |a_i| \,, \quad \text{for all } \sum_{i=0}^{d-1} a_i X^i \in \mathcal{R} \,.$$

This norm has a natural extension to $\mathcal{R}^n$, i.e., $\|\mathbf{x}\|_\infty = \max_{1 \le i \le n} \|x_i\|_\infty$ for all vectors $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{R}^n$. Further, for $\alpha \in \mathbb{R}_{\ge 0}$, we write

$$\mathcal{R}(\alpha) = \{x \in \mathcal{R} : \|x\|_\infty \le \alpha\} \,.$$

The MSIS-based commitment scheme, described in Definition 5.8, allows a prover to commit to vectors $\mathbf{x} \in \mathcal{R}^n$ of bounded ℓ_∞-norm, i.e., $\mathbf{x} \in \mathcal{R}(\alpha)^n$ for some $\alpha \in \mathbb{R}_{\ge 0}$. It is based on Ajtai's seminal work [Ajt96] and different variants of this commitment scheme have been presented in prior works, e.g., in [BKL+15; BBC+18; BDL+18]. This commitment scheme is oftentimes instantiated with norm bound $\alpha = \lceil (p-1)/2 \rceil$ for some $p \in \mathbb{N}$. This instantiation allows a prover to commit to vectors in $\mathcal{R}_p^n$.

Definition 5.8 (Lattice-Based Commitment Scheme)**.** Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(x) \in \mathbb{Z}[X]$ of degree d and let $\alpha \in \mathbb{R}_{\ge 0}$. Then, the following setup algorithm and commitment function define a lattice-based vector commitment scheme:

- $\mathsf{pk} = (A_1, A_2) \leftarrow \textsc{Setup}(1^\lambda, \mathcal{R}, q, k, r, \alpha, n)$, where $q > 2\alpha$ is a rational prime, $k, r \in \mathbb{N}$ and $(A_1, A_2) \leftarrow_R \mathcal{R}_q^{k\times(n+r)}$ is sampled uniformly at random;
- $\textsc{com}_{\mathsf{pk}} \colon \mathcal{R}(\alpha)^n \times \mathcal{R}(\alpha)^r \to \mathcal{R}_q^k, \quad (\mathbf{x}; \gamma) \mapsto A_1\mathbf{x} + A_2\gamma \mod q.$

When considered as a function on $\mathcal{R}^n \times \mathcal{R}^k$, the commitment function $\textsc{com}_{\mathsf{pk}}$ is an $\mathcal{R}$-module *homomorphism*. Moreover, the following lemma shows that the commitment scheme is computationally *binding* under the MSIS assumption. Note that, for large enough $n + r$, the hardness of the $\mathrm{MSIS}^\infty_{k,n+r,2\alpha}$ problem is independent of $n + r$ (see Equation 2.2). Therefore, this vector commitment scheme is *compact*, i.e., the size of a commitment is constant in the input dimension n.

Lemma 5.2 (Binding)**.** *The commitment scheme of Definition 5.8 is binding, conditioned on the hardness of the* $\mathrm{MSIS}^\infty_{q,k,n+r,2\alpha}$*-problem over* $\mathcal{R}$*.*

Proof. Suppose that $(\mathbf{x}; \gamma) \neq (\mathbf{x}'; \gamma')$ are two distinct openings of the same commitment P. Then $\mathbf{s} = (\mathbf{x} - \mathbf{x}'; \gamma - \gamma') \neq 0$ satisfies $\|\mathbf{s}\|_\infty \leq 2\alpha$ and $[A_1, A_2]\mathbf{s} = 0$, i.e., $\mathbf{s}$ is a solution of the $\mathrm{MSIS}^\infty_{k,n+r,2\alpha}$ problem, which completes the proof. □

The following lemma shows that if q is chosen to be inert in $\mathcal{R}$, i.e., if $\mathcal{R}_q$ is a field, and the randomness dimension r is large enough, then the commitment scheme is statistically *hiding*. The assumption that q is inert in $\mathcal{R}$ is only made to simplify the exposition. In this case, $\mathcal{R}_q$ is a field and it is easily seen that

$$\Pr\big(A\mathbf{x} = A\mathbf{y} : A \leftarrow_R \mathcal{R}_q^{k\times r}\big) \leq \frac{1}{\big|\mathcal{R}_q^k\big|} = \frac{1}{q^{dk}} \quad \forall \mathbf{x} \neq \mathbf{y} \in \mathcal{R}_q^r\,,$$

i.e., the family of hash functions $h_A \colon \mathcal{R}_q^r \to \mathcal{R}_q^k,\ \mathbf{x} \mapsto A\mathbf{x}$ is universal. By contrast, if $\mathcal{R}_q$ is not a field and contains zero-divisors, this family of hash functions is not universal. Based on [LS18], Baum et al. [BDL+18] show how this lemma can be generalized to arbitrary (not necessarily inert) primes q. The results of [LS18], and thus the generalization of [BDL+18], are only applicable to cyclotomic number rings $\mathcal{R}$. Fortunately, the generalization [ACX21] of [LS18] allows one to handle arbitrary number rings $\mathcal{R} = \mathbb{Z}[X]/f(X)$.

Lemma 5.3 (Hiding)**.** *Let* $\mathcal{R} = \mathbb{Z}[X]/f(X)$ *for a monic and irreducible polynomial* $f(x) \in \mathbb{Z}[X]$ *of degree* $d \in \mathbb{N}$*, and let* λ *denote the security parameter. If* q *is inert in* $\mathcal{R}$ *and* $r \in \mathbb{N}$ *is such that*

$$r \geq \frac{2\lambda + dk\log_2 q}{d\log_2(2\alpha + 1)}\,,$$

then the commitment scheme of Definition 5.8 is statistically hiding.

Proof. Since q is inert in $\mathcal{R}$, it follows that $\mathcal{R}_q$ is a field and the family of functions $h_A \colon \mathcal{R}_q^r \to \mathcal{R}_q^k,\ \mathbf{x} \mapsto A\mathbf{x}$, indexed by $A \in \mathcal{R}_q^{k\times r}$, is a universal hash family. Further, the min-entropy of the uniform distribution over $\mathcal{R}(\alpha)^r$ equals

$$dr\log_2(2\alpha + 1) \geq 2\lambda + dk\log_2 q\,.$$

Since $q > 2\alpha$ and by the leftover hash lemma [ILL89], it therefore follows that the statistical distance between the distribution

$$\mathcal{X} = \{(A, A\gamma) : A \leftarrow_R \mathcal{R}_q^{k\times r}, \gamma \leftarrow_R \mathcal{R}(\alpha)^r\}$$

and the uniform distribution $\mathcal{U}$ over $\mathcal{R}_q^{k\times r} \times \mathcal{R}_q^k$ is at most $2^{-\lambda}$, which proves the lemma. □

As in the strong-RSA instantiation, due to the soundness slack and approximation factor, our compressed Σ-protocols only allow a prover to prove knowledge of a *relaxed* opening. The following definition formalizes the notion of a relaxed commitment opening for the lattice-based commitment scheme of Definition 5.8.

Definition 5.9 ((τ, ζ)-Relaxed Commitment Opening)**.** Let $\tau \in \mathbb{R}_{\geq 0}$, $\zeta \in \mathcal{R}$ and let P be a commitment for the commitment scheme of Definition 5.8. A (τ, ζ)-relaxed opening of P is a pair $(\mathbf{x}; \gamma) \in \mathcal{R}^{n+r}$, such that $\text{COM}(\mathbf{x}; \gamma) = \zeta \cdot P \in \mathcal{R}_q^k$ and $\|(\mathbf{x}; \gamma)\|_\infty \leq \tau\alpha$.

A (τ, ζ)-relaxed opening of a commitment P differs in two ways from a standard opening. First, it contains an approximation factor ζ, such that the relaxed opening gives a short preimage for $\zeta \cdot P \in \mathcal{R}_q^k$ instead of P. Second, the norm-bound $\tau\alpha$ of relaxed openings differs from the norm bound α on honestly committed vectors (typically $\tau > 1$).

As long as it is infeasible to find two distinct (τ, ζ)-relaxed openings $(\mathbf{x}; \gamma)$ and $(\mathbf{x}'; \gamma')$ of a commitment P with $(\mathbf{x}; \gamma) \neq (\mathbf{x}'; \gamma')$, proving knowledge of relaxed opening is sufficient in most practical scenarios. In this case, we say the commitment scheme is binding with respect to (τ, ζ)-relaxed openings. The following lemma reduces breaking the "binding with respect to relaxed openings" property to solving the MSIS-problem. Note that the hardness of the corresponding MSIS-problem does not depend on the approximation factor ζ.

Lemma 5.4 (Binding with respect to (τ, ζ)-Relaxed Openings)**.** *Let $\tau \in \mathbb{R}_{\geq 0}$ and $\zeta \in \mathcal{R}$. The commitment scheme of Definition 5.8 is binding with respect to (τ, ζ)-relaxed openings, conditioned on the hardness of the* $\text{MSIS}^\infty_{q,k,n+r,2\tau\alpha}$*-problem over $\mathcal{R}$.*

Proof. Suppose that $(\mathbf{x}; \gamma)$ and $(\mathbf{x}'; \gamma')$ are distinct (τ, ζ)-relaxed openings of a commitment P. Then $\mathbf{s} = (\mathbf{x} - \mathbf{x}'; \gamma - \gamma') \neq 0$ satisfies $\|\mathbf{s}\|_\infty \leq 2\tau\alpha$ and $[A_1, A_2]\mathbf{s} = 0$, i.e., $\mathbf{s}$ is a solution of the $\text{MSIS}^\infty_{k,n+r,2\tau\alpha}$ problem, which completes the proof. □

Our goal is to prove knowledge of a (τ, ζ)-relaxed commitment opening $(\mathbf{x}; \gamma)$, for appropriate $\tau \in \mathbb{R}_{\geq 0}$ and $\zeta \in \mathcal{R}$, that satisfies the constraint $L(\mathbf{x}) = \zeta \cdot y$, where $L\colon \mathcal{R}^n \to \mathcal{R}'$ is an $\mathcal{R}$-module homomorphism for some arbitrary $\mathcal{R}'$. To this end, we consider the following $\mathcal{R}$-module homomorphism:

$$\Psi\colon \mathcal{R}^n \times \mathcal{R}^r \to \mathcal{R}_q^k \times \mathcal{R}', \quad (\mathbf{x}; \gamma) \mapsto \left(A_1\mathbf{x} + A_2\gamma, L(\mathbf{x})\right).$$

Typically, $\mathcal{R}' = \mathcal{R}_p$ for some rational prime $p \neq q$. Note that, if the approximation factor ζ is invertible in $\mathcal{R}'$, then $L(\mathbf{x}) = \zeta \cdot y$ implies that $L(\zeta^{-1} \cdot \mathbf{x}) = y$.

For this reason, in most practical scenarios, the approximation factor is required to be invertible in $\mathcal{R}'$.

The Ψ-instantiation of the compressed Σ-protocol of Section 3.3, with some challenge set $\mathcal{C} \subseteq \mathcal{R}$, requires rejection sampling in order to be special honest-verifier zero-knowledge (SHVZK). More precisely, it requires a distribution-algorithm pair $(\mathcal{D}, \mathcal{F})$ that is V-hiding, for $V = \{c\mathbf{x} : \mathbf{x} \in \mathcal{R}(\alpha)^{n+r} \wedge c \in \mathcal{C}\}$, and β-bounded for some reasonably small $\beta \in \mathbb{R}_{\geq 0}$ (Definition 3.2). In our instantiation, we let $\mathcal{D}$ be the uniform distribution over an appropriate subset of $\mathcal{R}^{n+r}$. The following lemma shows that this approach gives the required properties.

Lemma 5.5 (Uniform Rejection Sampling)**.** *Let $\mathcal{R} = \mathbb{Z}[X]/f(X)$ for a monic and irreducible polynomial $f(X) \in \mathbb{Z}[X]$ of degree d, $\mathcal{C} \subseteq \mathcal{R}$ and $n, r \in \mathbb{N}$. Recall that $\mathcal{R}(\alpha) = \{\mathbf{x} \in \mathcal{R} : \|\mathbf{x}\|_\infty \leq \alpha\}$ and*

$$w(\mathcal{C}) = \max_{c \in \mathcal{C}, x \in \mathcal{R} \setminus \{0\}} \frac{\|cx\|_\infty}{\|x\|_\infty} .$$

Further, let $V = \{c\mathbf{x} \in \mathcal{R}^{n+r} : \mathbf{x} \in \mathcal{R}(\alpha)^{n+r} \wedge c \in \mathcal{C}\}$, $\gamma > w(\mathcal{C})\alpha$, $\mathcal{D}$ the uniform distribution over $\mathcal{R}(\gamma)^{n+r}$ and

$$\mathcal{F}(\mathbf{r}, \mathbf{v}) = \begin{cases} \perp, & \text{if } \|\mathbf{v} + \mathbf{r}\|_\infty > \gamma - w(\mathcal{C})\alpha, \\ \mathbf{v} + \mathbf{r}, & \text{otherwise.} \end{cases}$$

Then $(\mathcal{D}, \mathcal{F})$ is perfectly V-hiding and $(\gamma - w(\mathcal{C})\alpha)$-bounded, with abort probability

$$\delta \leq (n + r)d \cdot \frac{2w(\mathcal{C})\alpha + 2}{2\gamma + 1} .$$

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

$$\begin{aligned} \delta &= 1 - \left(\frac{2\lfloor \gamma - w(\mathcal{C})\alpha \rfloor + 1}{2\lfloor \gamma \rfloor + 1} \right)^{(n+r)d} \\ &\leq 1 - \left(1 - \frac{2w(\mathcal{C})\alpha + 2}{2\gamma + 1} \right)^{(n+r)d} \\ &\leq (n + r)d \cdot \frac{2w(\mathcal{C})\alpha + 2}{2\gamma + 1} . \end{aligned}$$

where the final step follows from Bernoulli's inequality.

Now let $\mathcal{F}'$ be the algorithm that aborts with probability δ and otherwise outputs $\mathbf{z} \leftarrow_R \mathcal{R}\big(\gamma - w(\mathcal{C})\alpha\big)^{n+r}$ sampled uniformly at random. Then it is easily seen that $\{\mathcal{F}(\mathbf{r}, \mathbf{v}) \mid \mathbf{r} \leftarrow_R \mathcal{D}\}$ and $\{\mathcal{F}'\}$ have exactly the same output distributions, i.e., $(\mathcal{D}, \mathcal{F})$ is perfectly V-hiding.

Finally, $(\mathcal{D}, \mathcal{F})$ is $(\gamma - w(\mathcal{C})\alpha)$-bounded, which completes the proof. □

Remark 5.1. The smallest lattice-based signatures actually take $\mathcal{D}$ to be a Gaussian distribution. Namely, when the secrets have a bounded ℓ_2-norm, the Gaussian

distribution results in better protocol parameters. In our instantiation, this is not the case; our secrets are bounded with respect to the ℓ_∞-norm. For this reason, it is beneficial to resort to a uniform distribution over an appropriate subset of $\mathcal{R}^{n+r}$. An additional benefit is that uniform sampling is less prone to side-channel attacks. This is the reason that the lattice-based digital signature scheme Dilithium also deploys a uniform rejection sampling approach [DKL+18].

Let now Π_{MSIS} be the compressed Σ-protocol of Section 3.3.3 instantiated for homomorphism Ψ, with the rejection sampling approach of Lemma 5.5 and some arbitrary ζ-exceptional challenge set $\mathcal{C} \subseteq \mathcal{R}$. The properties of Π_{MSIS}, summarized in the following theorem, follow immediately from Theorem 3.9 and Lemma 5.5.

Theorem 5.5 (MSIS-based Compressed Σ-Protocol)**.** *Let $n + r = 2^\mu$ for some $\mu \in \mathbb{N}$. Let Π_{MSIS} be the compressed Σ-protocol Π_{comp}, described in Protocol 8, instantiated with a ζ-exceptional challenge set $\mathcal{C} \subseteq \mathcal{R}$ of cardinality at least 3, the distribution-algorithm pair $(\mathcal{D}, \mathcal{F})$ of Lemma 5.5 and MSIS-based homomorphism*

$$\Psi \colon \mathcal{R}^n \times \mathcal{R}^r \to \mathcal{R}_q^k \times \mathcal{R}', \quad (\mathbf{x}; \gamma) \mapsto \big(A_1\mathbf{x} + A_2\gamma, L(\mathbf{x})\big)\,.$$

Then Π_{MSIS} is an interactive proof for relation

$$\mathfrak{R}_{\mathrm{MSIS}} = \{(P, y, \alpha; \mathbf{x}) : \Psi(\mathbf{x}) = (P, y) \wedge \|\mathbf{x}\|_\infty \leq \alpha\}\,.$$

It is complete with completeness error

$$\delta \leq (n + r)d \cdot \frac{2w(\mathcal{C})\alpha + 2}{2\gamma + 1}\,,$$

$(2, 3, \ldots, 3)$-out-of-$(|\mathcal{C}|, \ldots, |\mathcal{C}|)$ special-sound with soundness slack

$$\tau = 2 \cdot 6^\mu \cdot \overline{w}(\mathcal{C}, \zeta)^{3\mu+1} \cdot \big(w(\mathcal{C})^2 + w(\mathcal{C})^3\big)^\mu \cdot w(\mathcal{C}) \cdot \frac{\gamma - w(\mathcal{C})\alpha}{\alpha}$$

and approximation factor $\zeta^{3\mu+1}$, and it is non-abort special honest-verifier zero-knowledge (NA-SHVZK).

Moreover, it has $2\mu+3$ communication rounds and the communication costs are:

- *$\mathcal{P} \to \mathcal{V}$: 1 element of $\mathcal{R}$ with norm at most $\big(1+w(\mathcal{C})\big)^\mu \cdot \big(\gamma - w(\mathcal{C})\alpha\big)$, $2\mu+1$ elements of $\mathcal{R}'$ and $2\mu+1$ elements of $\mathcal{R}_q^k$;*
- *$\mathcal{V} \to \mathcal{P}$: $\mu + 1$ elements of $\mathcal{C} \subseteq \mathcal{R}$.*

As a concrete example of the compressed Σ-protocol Π_{MSIS}, let us consider the cyclotomic number ring $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ with $d = 2^{d'}$ a power-of-two and challenge set $\mathcal{C} = \{0, \pm 1, \pm X, \ldots, \pm X^{d-1}\} \subset \mathcal{R}$. Further, by taking norm-bound $\alpha = (p - 1)/2$ for some odd prime p and $\mathcal{R}' = \mathcal{R}_p$, we consider a prover that wishes to commit to vectors $\mathbf{x} \in \mathcal{R}_p^n$ and prove that $L(\mathbf{x}) = y$ for some linear form $L \colon \mathcal{R}_p^n \to \mathcal{R}_p$. In Section 5.5, we used exactly the same ring for the base-extension of our strong-RSA instantiation. Moreover, in Lemma 5.1 and Corollary 5.1, we showed that $\mathcal{C}$ is a 2-exceptional subset with $w(\mathcal{C}) = 1$ and $\overline{w}(\mathcal{C}, 2) = d$. Note that, since p is odd, the approximation factor $\zeta = 2$ is invertible in $\mathcal{R}_p$. Let us now analyze the communication complexity of this example.

This instantiation of Π_{MSIS} is $(2,3,\dots,3)$-out-of-$(2d+1,\dots,2d+1)$ special-sound. In Chapter 6, we will see that it therefore has knowledge error

$$1-\left(1-\frac{1}{2d+1}\right)\left(1-\frac{2}{2d+1}\right)^{\mu}\leq 1-\left(1-\frac{2}{2d+1}\right)^{\mu+1}\leq\frac{2\mu+2}{2d+1}\,,$$

where $\mu=\log_2(n+r)$. For simplicity, let us assume that $d\geq 2\mu+2$. Then, this compressed Σ-protocol has knowledge error at most $1/2$, and $t\leq\lambda$ parallel repetitions are required to reduce the knowledge error down to $2^{-\lambda}$. If $d<2\log_2(n+r)+2$, the base extension techniques of Section 3.3.4 can be deployed to increase the size of the challenge set.

Further, we let $\gamma=\Theta((n+r)td\alpha)=\Theta((n+r)tdp)$. By Theorem 5.5, this is enough to achieve a constant completeness error. Altogether, this instantiation allows a prover to prove knowledge of $(\tau,2^{3\mu+1})$-relaxed commitment openings, where

$$\tau=2d\cdot(12d^3)^{\mu}\cdot\frac{\gamma-\alpha}{\alpha}=\Theta\big(t\cdot d^2\cdot(n+r)^{3+\log_2 3+3\log_2 d}\big)\,.$$

Hence, in practice, the commitment scheme must be instantiated to be binding with respect to $(\tau,2^{3\mu+1})$-relaxed openings, i.e., the $\mathrm{MSIS}^{\infty}_{q,k,n+r,2\tau\alpha}$-problem over $\mathcal{R}$ must be computationally hard (Lemma 5.4). From the Micciancio-Regev bound (Equation 2.2) it follows that this problem is hard if

$$dk\log_2 q\geq\frac{\log_2^2(2\tau\alpha\sqrt{n+r})}{4\log_2\delta}=\Theta\left(\frac{\log^2 d\cdot\log^2\big(tdp\cdot(n+r)\big)}{\log\delta}\right),\tag{5.2}$$

where δ is the root Hermite factor.

By Theorem 5.5 and the fact that $t=\mathcal{O}(\lambda)$, it therefore follows that the resulting t-fold parallel repetition of Π_{MSIS} has communication complexity

$$\mathcal{O}\left(\frac{\lambda\cdot\log(n+r)\cdot\log^2 d\cdot\log^2\big(\lambda dp\cdot(n+r)\big)}{\log\delta}\right).$$

Finally, by Lemma 5.3 and Equation 5.2, we observe that $r=\mathcal{O}(\lambda+\frac{\log\lambda pn}{\log\delta})$. Hence, the resulting protocol has *polylogarithmic* communication complexity.

Note that, in the discrete logarithm instantiation over the group $\mathbb{G}$, the secret vector $\mathbf{x}$ has coefficients in the finite field $\mathbb{Z}_q$, where q is the exponent of $\mathbb{G}$. For the discrete logarithm problem to be hard in $\mathbb{G}$, the size of the prime q must therefore be exponential in the security parameter. The discrete logarithm instantiation does not allow a prover to directly prove relations over fields $\mathbb{Z}_p$ of small characteristic p. By contrast, the above lattice-based instantiation does not suffer from this limitation. In fact, smaller primes p correspond to harder MSIS-problem instantiations.

Remark 5.2. For simplicity, we have deployed a standard parallel repetition approach to reduce the knowledge error down to $2^{-\lambda}$. More precisely, in the considered t-fold parallel repetition, the verifier only accepts if the prover succeeds in all t parallel instances. However, while decreasing the knowledge error, this approach also increases the completeness error. To account for this effect, we have

chosen the protocol parameter γ to increase linearly in the number of parallel repetitions t. In Section 6.5.4, we describe a threshold parallel repetition approach that decreases both the completeness and knowledge error simultaneously. This approach would therefore allow for a further improvement of the above lattice instantiation.