More Efficient Commitments from Structured Lattice Assumptions

Carsten Baum[1]⋆, Ivan Damgård⋆⋆[2], Vadim Lyubashevsky[3]⋆⋆⋆, Sabine Oechsner[2]†, and Chris Peikert‡[4]

[1] Department of Computer Science, Bar-Ilan University
carsten.baum@biu.ac.il
[2] Department of Computer Science, Aarhus University
{ivan, oechsner}@cs.au.dk
[3] IBM Research – Zurich, Switzerland
vad@zurich.ibm.com
[4] Department of Computer Science and Engineering, University of Michigan
cpeikert@umich.edu

Abstract. We present a practical construction of an additively homomorphic commitment scheme based on structured lattice assumptions, together with a zero-knowledge proof of opening knowledge. Our scheme is a design improvement over the previous work of Benhamouda et al. in that it is not restricted to being statistically binding. While it is still possible to instantiate our scheme to be either statistically binding or statistically hiding, it is most efficient when both hiding and binding properties are only computational. This results in approximately a factor of 4 reduction in the size of the proof and a factor of 6 reduction in the size of the commitment over the aforementioned scheme.

1 Introduction

Over the past several years, lattice-based cryptography has developed and matured rapidly. As this development continues, it is desirable to have a full suite of efficient lattice-based tools and protocols. This is particularly important since lattice problems are currently some of the promising "post-quantum" replacements for the discrete logarithm and factoring problems. Therefore, we want to construct standard cryptographic primitives such as encryption and commitment schemes, plus companion protocols, such as zero-knowledge proofs, in the lattice setting.

Commitment schemes [Blu82] are a key tool in the design of cryptographic protocols and have numerous applications (e.g. threshold encryption [DF89], electronic voting [CFSY96], etc.). In particular, when combined with zero-knowledge proofs, they can enforce "good" behavior by adversarial parties and make the design of protocols secure against malicious attacks easier. The main result of

⋆ Supported by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister's Office and COST Action IC1306.

⋆⋆ Supported by the European Research Council (ERC) under the European Unions's Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO)

⋆⋆⋆ Supported by the SNSF ERC Transfer Grant CRETP2-166734 – FELICITY.

† Supported by the European Research Council (ERC) under the European Unions's Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO); the Danish Independent Research Council under Grant-ID DFF-6108-00169 (FoCC); the European Union's Horizon 2020 research and innovation programme under grant agreement No 731583 (SODA)

‡ Supported by the National Science Foundation under CAREER Award CCF-1054495 and CNS-1606362, the Alfred P. Sloan Foundation, and by a Google Research Award. The views expressed are those of the authors and do not necessarily reflect the official policy or position of the National Science Foundation, the Sloan Foundation, or Google.

this work is the construction of an efficient commitment scheme and accompanying zero-knowledge proofs of knowledge for proving relations among committed values.

1.1 Related Work

There are several earlier works in this area: Kawachi et al.'s work on identification schemes [KTX08] presents a string commitment scheme based on the SIS assumption [Ajt96], where one commits to vectors over $\mathbb{Z}_q$. However, the message space is restricted to vectors of small norm; otherwise, the binding property is lost. This restriction causes problems in the applications we are interested in: for instance, if a player wants to prove (efficiently) that he has performed an encryption or decryption operation correctly in a cryptosystem that uses the ring $\mathbb{Z}_q$, one typically requires a commitment scheme that is linearly homomorphic and can commit to *arbitrary* vectors over $\mathbb{Z}_q$ rather that only short ones.

In [JKPT12], Jain et al. proposed a commitment scheme where the hiding property is based on the Learning Parity with Noise (LPN) assumption, a special case of the Learning With Errors (LWE) assumption [Reg05]. They also constructed zero-knowledge proofs to prove general relations on bit strings. A generalization of [JKPT12] was proposed by Xie et al. [XXW13]. Their work presents a commitment scheme that is based on Ring-LWE [LPR10] instead of LPN, and they build Σ-protocols from it. Further Σ-protocols based on (Ring-)LWE encryption schemes were presented by Asharov et al. [AJL$^+$12] and Benhamouda et al. [BCK$^+$14].

A main drawback of all these previous schemes is that the zero-knowledge proofs had a non-negligible soundness error, and hence one needs many iterations to have full security. In [BKLP15], a commitment scheme, as well as companion zero-knowledge protocols were constructed with much better efficiency: one can commit to a vector over $\mathbb{Z}_q$ resulting in a commitment that is only a constant factor larger than the committed vector. Furthermore, they gave protocols for proving knowledge of a committed string as well as proving linear and multiplicative relations on committed values. These are efficient in the sense that the soundness error is negligible already for a single iteration of the protocol. The commitments are unconditionally binding and computationally hiding, and the underlying assumption is Ring-LWE.

1.2 Our Contributions

We propose a commitment scheme that allows to commit to vectors over polynomial rings, as well as associated zero-knowledge proofs proofs of knowledge for proving knowledge of the commitment and relationships between committed values. In comparison to [BKLP15], which is the most closely related previous work, we achieve all the "different flavors" of commitments. While the technique in [BKLP15] only leads to a statistically binding commitment scheme, we show how to achieve statistically binding, statistically hiding, and a more efficient scheme that is only computationally hiding and binding. The latter construction gives rise to the currently most practical instantiation of a commitment scheme (that admits a zero-knowledge proof of opening) of arbitrary-sized messages based on the hardness of lattice problems. The binding property of our scheme relies on the Module-LWE assumption, while the hiding is based on the hardness of the Module-SIS problem.

Since the public appearance of a preliminary version of this work, the commitment scheme has already been crucially used for applications of constructing practical lattice-based voting schemes [dPLNS17] and privacy-preserving protocols [dPLS18].

1.3 Paper Organization

In Section 3, we introduce the problems (which are equivalent to Module-SIS and Module-LWE) upon which our commitment scheme will be based. We also show that for certain parameters these problems are information-theoretically hard – thus if one wants to build commitment schemes that are statistically hiding / binding, they need to have these properties based on the problems with those parameters. In Section 4, we present our commitment scheme along with the zero-knowledge proof of opening, and various relations between committed values. In Section 5, we first present a tighter analysis of the [BKLP15] scheme which allows one to instantiate it with smaller parameters, and then we compare it to the scheme from Section 4.

2 Preliminaries

2.1 The Setting

Let q be a prime and $r \in \mathbb{N}^+$. We set $N = 2^r$ and define the rings $R = \mathbb{Z}[X]/\langle X^N + 1\rangle, R_q = \mathbb{Z}_q[X]/\langle X^N + 1\rangle$. This is the setting that we will use throughout this work. We also define $\boldsymbol{I}_k \in R^{k\times k}$ to be an identity matrix of dimension (over R) k.

For each $f \in R$, let $f = \sum_i f_i X^i$, then we can define the following norms of f:

$$\begin{aligned}\ell_1 : ||f||_1 &= \sum\nolimits_i |f_i| \\ \ell_2 : ||f||_2 &= (\sum\nolimits_i |f_i|^2)^{1/2} \\ \ell_\infty : ||f||_\infty &= \max_i |f_i|.\end{aligned}$$

For $g \in R_q$ and $g = \sum_i \overline{g}_i X^i$, we identify each $\overline{g}_i$ with an element $g_i \in [-\frac{q-1}{2}, \frac{q-1}{2}]$ such that $\overline{g}_i = g_i \bmod q$. For a positive integer α, we write S_α to be the set of all elements in R with ℓ_∞-norm at most α.

For $\boldsymbol{f} \in R_q^k$ we then have the standard inequalities

$$||\boldsymbol{f}||_1 \le \sqrt{kN}||\boldsymbol{f}||_2 \le kN||\boldsymbol{f}||_\infty \text{ and } ||\boldsymbol{f}||_\infty \le ||\boldsymbol{f}||_1$$

The choice of the polynomial $X^N + 1$ allows to give tight bounds on the norms of product $f \cdot g$ of polynomials $f, g \in R_q$, based on their respective norms. In this work, we use the following two bounds (c.f. [Mic07], which are applicable to the polynomial modulus $X^N + 1$ and $X^N - 1$):

1. If $||f||_\infty \le \beta, ||g||_1 \le \gamma$ then $||f \cdot g||_\infty \le \beta \cdot \gamma$.
2. If $||f||_2 \le \beta, ||g||_2 \le \gamma$ then $||f \cdot g||_\infty \le \beta \cdot \gamma$.

2.2 Invertible Elements in R_q and the Challenge Space.

Of special importance in our work will be sets of elements of R_q that are both invertible and of small norm. The following Lemma shows that if one chooses the prime q in a particular way, then all elements with small norms (either ℓ_2 or ℓ_∞) will be invertible.

Lemma 1. *([LS17, Corollary 1.2]) Let $N \geq d > 1$ be powers of 2 and $q \equiv 2d + 1 (\bmod 4d)$ be a prime. Then $X^N + 1$ factors into d irreducible polynomials $X^{N/d} - r_j$ modulo q and any $y \in R_q \setminus \{0\}$ that satisfies*

$$\|y\|_\infty < \frac{1}{\sqrt{d}} \cdot q^{1/d} \quad \text{or} \quad \|y\|_2 < q^{1/d}$$

is invertible in R_q.

We will need invertibility of polynomials for two separate purposes. First, working with invertible polynomials will allow us to prove the universality of certain hash function families, which is important for establishing statistical binding and statistical hiding properties of our protocol.

More importantly, though, we will need the challenge space of our zero-knowledge proof to consist of short elements such that every difference of distinct elements is invertible in R_q. This property is crucial to the soundness of our zero-knowledge proof of commitment opening. For practical purposes, we would also like to define our sets so that they are easy to sample from.

The Challenge Space $\mathcal{C}$. One common way to define this challenge space is as

$$\mathcal{C} = \{c \in R_q \mid \|c\|_\infty = 1, \, \|c\|_1 = \kappa\}. \tag{1}$$

If we would like the size of $\mathcal{C}$ to be 2^λ, then we need to set κ such that $\binom{N}{\kappa} \cdot 2^\kappa > 2^\lambda$. For example, if $N = \lambda = 256$, then we can set $\kappa = 60$. Throughout the paper we will be assuming that the parameters of the ring R_q are set in such a way (as dictated by Lemma 1) that all non-zero elements of ℓ_∞-norm at most 2 are invertible in R_q. This implies that for any two distinct $c, c' \in \mathcal{C}$, the difference $c - c'$ is invertible in R_q. For convenience, we define this set of differences as $\bar{\mathcal{C}} = \{c - c' \mid c \neq c' \in \mathcal{C}\}$.

2.3 Normal Distributions

The continuous normal distribution over $\mathbb{R}^N$ centered at $\boldsymbol{v} \in \mathbb{R}^N$ with standard deviation σ has probability density function

$$\rho^N_{\boldsymbol{v},\sigma}(\boldsymbol{x}) = \frac{1}{\sqrt{2\pi}\sigma} \cdot \exp\left(\frac{-||\boldsymbol{x} - \boldsymbol{v}||_2^2}{2\sigma^2}\right).$$

In this work we are more interested in a discrete version. The *discrete normal distribution* over R^k centered at $\boldsymbol{v} \in R^k$ with standard deviation σ is given by the distribution function (for all $\boldsymbol{x} \in R^k$)

$$\mathcal{N}^k_{\boldsymbol{v},\sigma}(\boldsymbol{x}) = \rho^{k \cdot N}_{\boldsymbol{v},\sigma}(\boldsymbol{x}) / \rho^{k \cdot N}_\sigma(R^k),$$

where we omit the subscript $\boldsymbol{v}$ when it is zero.

For the zero-knowledge proofs, we adapt the tail-bound from [Lyu12, Lemma 4.4] as

Remark 1. For any $\delta > 0$,

$$\Pr[||\boldsymbol{z}||_2 > \delta\sigma\sqrt{kN} \mid \boldsymbol{z} \xleftarrow{\$} \mathcal{N}^k_\sigma] < \delta^{kN} \cdot \exp\left(\frac{kN}{2}(1 - \delta^2)\right).$$

In our protocols, we set $\delta = 2$. This choice is sufficient for Remark 1 as we surely have $N = \Omega(\lambda)$, so the tail-bound holds with probability that is overwhelming in λ.

Moreover, the rejection sampling theorem from [Lyu12, Theorem 4.6] can be expressed in our setting as follows:

Lemma 2. *Let $V \subseteq R^k$ such that all elements have $||\cdot||_2$-norm less than T, $\sigma \in \mathbb{R}$ such that $\sigma = \omega(T\sqrt{\log(kN)})$ and $h : V \to \mathbb{R}$ be a probability distribution. Then there exists a $M = O(1)$ such that the distribution of the following two algorithms $\mathcal{A}, \mathcal{S}$ is within statistical distance $2^{-\omega(\log(kN))}/M$.*

$\mathcal{A}$**:**

1. $\boldsymbol{v} \overset{\$}{\leftarrow} h$
2. $\boldsymbol{z} \overset{\$}{\leftarrow} \mathcal{N}^k_{\boldsymbol{v},\sigma}$
3. Output (z, v) *with probability* $\min\left(\frac{\mathcal{N}^k_\sigma(\boldsymbol{z})}{M\mathcal{N}^k_{\boldsymbol{v},\sigma}(\boldsymbol{z})}, 1\right)$

$\mathcal{S}$**:**

1. $\boldsymbol{v} \overset{\$}{\leftarrow} h$
2. $\boldsymbol{z} \overset{\$}{\leftarrow} \mathcal{N}^k_\sigma$
3. Output (z, v) *with prob.* $1/M$

The probability that $\mathcal{A}$ outputs something is at least $\dfrac{1 - 2^{-\omega(\log(kN))}}{M}$.

As mentioned in [Lyu12], by setting $\sigma = \alpha T$ one obtains

$$M = \exp\left(12/\alpha + 1/(2\alpha^2)\right)$$

such that the statistical distance of the output of $\mathcal{A}, \mathcal{S}$ is at most $2^{-100}/M$ while $\mathcal{A}$ outputs a result with probability at least $(1 - 2^{-100})/M$. In practice one would choose $kN \gg 128$, but already for $kN = 128$ one obtains that $M \approx 4.5$, and it just decreases for larger choices.

2.4 Commitments & Zero-Knowledge Proofs

For completeness, we now give a formal definition of commitment schemes and zero-knowledge proofs. As we mainly care about zero-knowledge proofs of opening knowledge for commitments in this work, the definitions will be tailored to this setting.

Consider the following three algorithms $\mathsf{KeyGen}, \mathsf{Commit}, \mathsf{Open}$, which have 1^λ as implicit input:

KeyGen is a PPT algorithm that outputs the public parameters $\mathsf{PP} \in \{0,1\}^{poly(\lambda)}$ containing a definition of the message space $\mathcal{M}$.

Commit is a PPT algorithm that, on input the public parameters PP and a message $x \in \mathcal{M}$ outputs values $c, r \in \{0,1\}^{poly(\lambda)}$.

Open is a deterministic polynomial-time algorithm that, on input the public parameters PP, a message $x \in \mathcal{M}$ and values $c, r \in \{0,1\}^{poly(\lambda)}$ outputs a bit $b \in \{0,1\}$.

A scheme is ϵ-hiding if for all algorithms $\mathcal{A}$, the probability (over the randomness of KeyGen, Commit, and the algorithm $\mathcal{A}$) that $i' = i$ in the below experiment is less than ϵ:

1. $\mathcal{A}$ receives $\mathsf{PP} \leftarrow \mathsf{KeyGen}()$
2. $\mathcal{A}$ outputs $x_0, x_1 \in \mathcal{M}$
3. $\mathcal{A}$ receives c created as: $i \leftarrow \{0,1\}$, $(c, r) \leftarrow \mathsf{Commit}(\mathsf{PP}, x_i)$
4. $\mathcal{A}$ outputs $i' \in \{0,1\}$

If the algorithms $\mathcal{A}$ are restricted to polynomial-time algorithms, then the scheme is called *computationally hiding*. If there is no restriction on the running time of such algorithms, then the scheme is *statistically hiding*. In this paper, we will be proving that $\mathcal{A}$ in fact cannot distinguish between a commitment of a message of his choosing and a uniformly-random element in the space of commitments. This definition is stronger and implies the one above.

Similarly, the commitment scheme is called ϵ-binding if, for any $\mathcal{A}$,

$$\Pr\left[\begin{array}{c} \mathcal{A}(\mathsf{PP}) = (x, x', r, r', c) \\ \text{s.t. } x \neq x' \ \& \ \mathsf{Open}(\mathsf{PP}, x, c, r) = \mathsf{Open}(\mathsf{PP}, x', c, r') = 1 \end{array} \middle| \ \mathsf{PP} \leftarrow \mathsf{KeyGen}() \right] < \epsilon,$$

where the probability is taken over the randomness of $\mathcal{A}$ and KeyGen. If we restrict $\mathcal{A}$ to being polynomial-time, then the binding property is *computational*. If we allow for arbitrarily-powerful algorithms, then the property is *statistical*.

Zero-Knowledge Proofs of Knowledge of Opening. A Zero-Knowledge Proof of Knowledge for the Opening of a commitment c is an interactive protocol Π between two PPT algorithms $\mathcal{P}, \mathcal{V}$, such that $\mathcal{V}$ in the end of Π outputs a bit. We call $\mathcal{P}$ the *prover* and $\mathcal{V}$ the *verifier*. Assume that $\mathsf{PP} \leftarrow \mathsf{KeyGen}(), x \in \mathcal{M}, (c, r) \leftarrow \mathsf{Commit}(\mathsf{PP}, x)$, then the protocol Π will have the following three properties:

- Completeness: If $\mathcal{P}$ on input (PP, c, x, r) and $\mathcal{V}$ on input (PP, c) follow the protocol honestly, then $\mathcal{V}$ outputs 1 except with negligible probability.
- Soundness: If a PPT algorithm $\mathcal{A}$ on input (PP, c) makes the algorithm $\mathcal{V}$ output 1 in Π with polynomial probability p, then there exists an algorithm $\mathcal{E}$ which, given black-box access to $\mathcal{A}$, outputs (x', r') such that $\mathsf{Open}(\mathsf{PP}, x', c, r') = 1$ in time $\mathit{poly}(p, \lambda)$ with constant non-zero probability.
- Honest-Verifier Zero-Knowledge: There exists a PPT algorithm $\mathcal{S}$ whose output distribution on input (PP, c) is indistinguishable of the transcript of Π when running with $\mathcal{P}, \mathcal{V}$.

In this work, we deal with statistical zero-knowledge proofs: the statistical distance of the output distribution of $\mathcal{S}$ and the transcripts of Π is negligible in κ.

3 The Knapsack Problem over R_q and Lattice Problems

The security of our commitment scheme is based on the hardness of the Module-SIS and Module-LWE problems defined in [LS15]. These problems are generalizations of the usual SIS [Ajt96] and LWE [Reg05] problems to polynomial rings. At the other extreme, these problems become exactly Ring-SIS [PR06,LM06] and Ring-LWE [LPR10]. As with SIS and LWE, these problems can be defined over any norm (in practice, we do not know of any algorithms that are more successful at attacking these problems due to the norm that is being used). Because it is convenient for our scheme, we will be relying on the Module-SIS problem in the ℓ_2-norm, and on the Module-LWE problem in the ℓ_∞ norm.

Module-SIS and Module-LWE problems are essentially vector knapsack problems over a particular ring. For this reason, rather than working with Module-SIS and Module-LWE, we will directly work with knapsacks. We first define the Search Knapsack problem in the ℓ_2 norm (SKS^2) and define its security. The SKS^2 problem is exactly the Module-SIS problem (in Hermite Normal Form).

Definition 1. *The* $\mathsf{SKS}^2_{n,k,\beta}$ *problem asks to find a short vector* $\boldsymbol{y}$ *satisfying* $[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y} = \boldsymbol{0}^n$ *when given a random* $\boldsymbol{A}'$*. We say that an algorithm* $\mathcal{A}$ *has advantage* ϵ *in solving the* $\mathsf{SKS}^2_{n,k,\beta}$ *problem if*

$$\Pr\left[\|y_i\|_2 \le \beta \wedge [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y} = \boldsymbol{0}^n \mid \boldsymbol{A}' \overset{\$}{\leftarrow} R_q^{n\times(k-n)}; \boldsymbol{0} \neq \boldsymbol{y} = \begin{bmatrix} y_1 \\ \dots \\ y_k \end{bmatrix} \leftarrow \mathcal{A}(\boldsymbol{A}')\right] \ge \epsilon$$

We next define the Decisional Knapsack problem in the ℓ_∞ norm (DKS^∞), which is equivalent to the Module-LWE problem when the number of samples is limited.

Definition 2. *The* $\mathsf{DKS}^\infty_{n,k,\beta}$ *problem asks to distinguish the distribution* $[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y}$ *for a short* $\boldsymbol{y}$*, from the uniform distribution when given* $\boldsymbol{A}'$*. We say that an algorithm* $\mathcal{A}$ *has advantage* ϵ *in solving the* $\mathsf{DKS}^\infty_{n,k,\beta}$ *problem if*

$$\Big| \Pr[b = 1 \mid \boldsymbol{A}' \overset{\$}{\leftarrow} R_q^{n\times(k-n)}; \boldsymbol{y} \overset{\$}{\leftarrow} S_\beta^k; b \leftarrow \mathcal{A}(\boldsymbol{A}', [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y})] \\ - \Pr[b = 1 \mid \boldsymbol{A}' \overset{\$}{\leftarrow} R_q^{n\times(k-n)}; \boldsymbol{u} \overset{\$}{\leftarrow} R_q^n; b \leftarrow \mathcal{A}(\boldsymbol{A}', \boldsymbol{u})] \Big| \ge \epsilon$$

3.1 Unconditional Hardness of the Knapsack Problem.

In this section we will give ranges of parameters when the DKS^∞ and SKS^2 problems become unconditionally hard. This will be used in the next section to derive parameter sets for when the commitment scheme is statistically binding or statistically hiding.

Lemma 3. *If* $\boldsymbol{y} = [y_1, \dots, y_k] \in R_q^k$ *has the property that each non-zero* y_i *is invertible in* R_q*, then*

$$\Pr_{\boldsymbol{A}' \overset{\$}{\leftarrow} R_q^{n\times(k-n)}} \left[[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y} = \boldsymbol{0}^n \right] = q^{-n\cdot N}.$$

Proof. Notice that if the product $[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y}$ is to be $\boldsymbol{0}^n$, then one of $y_{n+1}, \dots, y_k$ must be non-zero. Without loss of generality, assume that $y_k \neq 0$, and write $\boldsymbol{a}_i \in R_q^n$ to be the i^{th} column of $\boldsymbol{A} = [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ]$. Then the above probability can be rewritten as

$$\Pr_{\boldsymbol{a}_k \overset{\$}{\leftarrow} R_q^n} \left[\boldsymbol{a}_k \cdot y_k = -\sum_{i=1}^{k-1} y_i \boldsymbol{a}_i \right] = \Pr_{\boldsymbol{a}_k \overset{\$}{\leftarrow} R_q^n} \left[\boldsymbol{a}_k = -y_k^{-1} \cdot \sum_{i=1}^{k-1} y_i \boldsymbol{a}_i \right] = q^{-n\cdot N},$$

where we used the fact that y_k is invertible in R_q. □

Lemma 4. *Let* $1 < d < N$ *be a power of* 2*. If* q *is a prime congruent to* $2d + 1 (\bmod 4d)$ *and*

$$q^{n/k} \cdot 2^{256/(k\cdot N)} \le 2\beta < \frac{1}{\sqrt{d}} \cdot q^{1/d}. \tag{2}$$

then any (all-powerful) algorithm $\mathcal{A}$ *has advantage at most* 2^{-128} *in solving* $\mathsf{DKS}^\infty_{n,k,\beta}$*.*

Proof. By the form of q and the fact that $2\beta < \frac{1}{\sqrt{d}} \cdot q^{1/d}$, Lemma 1 implies that all non-zero $y \in S_{2\beta}$ are invertible. This fact will be used in the sequel.

We now would like to show that the function family

$$\mathcal{H} = \{h_{\boldsymbol{A}'} : S_\beta^k \to R_q^n\}, \text{ where } h_{\boldsymbol{A}'}(\boldsymbol{y}) = [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y}$$

is universal. In other words, we want to show that for all $\boldsymbol{y} \neq \boldsymbol{y}' \in S_\beta^k$,

$$\Pr_{\boldsymbol{A}'}\left[[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y} = [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y}'\right] \leq q^{-n \cdot N}.$$

For any $\boldsymbol{y} \neq \boldsymbol{y}'$, the above probability is equivalent to

$$\Pr_{\boldsymbol{A}'}\left[[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y} = \boldsymbol{0}^n\right], \tag{3}$$

where $\boldsymbol{y} = (y_1, \ldots, y_k)$ is some non-zero vector in $S_{2\beta}^k$. Since all non-zero y_i in $S_{2\beta}$ are invertible, Lemma 3 immediately proves that the probability in (3) is exactly $q^{-n \cdot N}$.

Since we established that $\mathcal{H}$ is universal hash function family mapping onto R_q^n, and the min-entropy of $\boldsymbol{y} \overset{\$}{\leftarrow} S_\beta^k$ is greater than

$$k \cdot N \cdot \log(2\beta) > k \cdot N \cdot \left(\frac{n}{k} \cdot \log q + \frac{256}{k \cdot N}\right) = \log\left|R_q^n\right| + 256.$$

By the Leftover Hash Lemma, this implies that the distribution $(\boldsymbol{A}', h_{\boldsymbol{A}'}(\boldsymbol{y}))$ is within statistical distance 2^{-128} of $(\boldsymbol{A}', \boldsymbol{u})$ for a uniform $\boldsymbol{u}$. The Lemma for the restriction of β as in (2) follows directly from the definition of the $\mathsf{DKS}_{n,k,\beta}^\infty$ problem.

□

Remark. The upper bound in (2) was used in the proof of the lemma for showing the universality of $\mathcal{H}$ when the domain consists of elements with ℓ_∞ norms less than half the upper-bound. Intuitively, though, the hardness of the $\mathsf{DKS}_{n,k,\beta}^\infty$ problem should increase as β increases and so the problem should be hard for values that are greater than the upper bound. Indeed we have simple reductions from $\mathsf{DKS}_{n,k,\beta}^\infty$ to $\mathsf{DKS}_{n,k,\beta'}^\infty$ when $\beta \mid \beta'$ (c.f. [Lyu12, Lemma 3.6]), but we are not aware how to obtain such a reduction for all $\beta' > \beta$. The lack of a general reduction, however, is not important to practical applications since one would anyway want to use the smallest value of β that satisfies the lower bound in (2).

Lemma 5. *Let $1 < d < N$ be a power of 2. If q is a prime congruent to $2d + 1 (\bmod 4d)$ and*

$$\beta < q^{1/d}, \text{ and}$$
$$\beta < \sqrt{\frac{N}{2\pi e}} \cdot q^{n/k} \cdot 2^{-128/(k \cdot N)} - \sqrt{N}/2,$$

then any (all-powerful) algorithm $\mathcal{A}$ has advantage at most 2^{-128} in solving $\mathsf{SKS}_{n,k,\beta}^2$.

Proof. By the form of q and the fact that $\beta < q^{1/d}$, Lemma 1 implies that all non-zero y_i with $\|y_i\| < \beta$ are invertible in R_q. Lemma 3 therefore implies that,

$$\Pr_{\boldsymbol{A}'}\left[[\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y} = \boldsymbol{0}^n\right] = q^{-n \cdot N}.$$

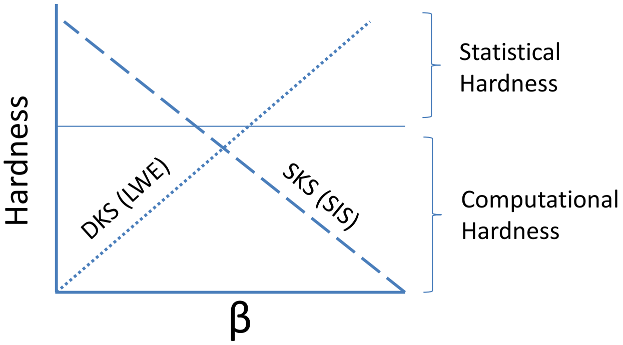

Fig. 1. As β increases, the DKS^∞ problem becomes harder, while the SKS^2 problem becomes easier.

If $V_N(r)$ is the volume of an N-dimensional ball of radius r, then it's simple to see that there are fewer than $V_N(\beta + \sqrt{N}/2)$ polynomials $y \in R$ such that $\|y\|_2 \leq \beta$.[5] By the union bound, we obtain

$$\begin{aligned}
&\Pr_{\boldsymbol{A}'}\left[\exists \boldsymbol{y} = \begin{bmatrix} y_1 \\ \dots \\ y_k \end{bmatrix} \text{ s.t. } \|y_i\| \leq \beta \text{ and } [\; \boldsymbol{I}_n \quad \boldsymbol{A}' \;] \cdot \boldsymbol{y} = \boldsymbol{0}^n \right] \\
&\leq V_N(\beta + \sqrt{N}/2)^k \cdot q^{-n \cdot N} \\
&< \left(\sqrt{\frac{2\pi e}{N}} \cdot (\beta + \sqrt{N}/2)\right)^{k \cdot N} \cdot q^{-n \cdot N}
\end{aligned}$$

□

3.2 Computational Hardness of the Knapsack Problem

For typical settings of parameters, the best attacks against the DKS^∞ and SKS^2 problems use lattice reduction algorithms. If we look at the $\mathsf{SKS}^2_{n,k,\beta}$ problem, then we can define the set

$$\Lambda = \{\boldsymbol{y} \in R^k \;:\; [\; \boldsymbol{I}_n \quad \boldsymbol{A}' \;] \cdot \boldsymbol{y} = \boldsymbol{0}^n \bmod q\}. \tag{4}$$

It's easy to see that Λ is an additive group over R^k. Finding a solution $\boldsymbol{y} = \begin{bmatrix} y_1 \\ \dots \\ y_k \end{bmatrix}$ such that $\|y_i\| \leq \beta$ is at least as hard as finding a $\boldsymbol{y}$ such that $\|\boldsymbol{y}\| \leq \beta \cdot \sqrt{k}$. Since Λ is also an additive group over $\mathbb{Z}^{k \cdot N}$, this is equivalent to finding a vector of norm $\beta \cdot \sqrt{k}$ in a random lattice of dimension kN. As we saw in Lemma 5, once β is small enough, such short vectors no longer exist and so even an all-powerful adversary cannot solve the $\mathsf{SKS}^2_{n,k,\beta}$ problem. But it is known that as β gets larger, the problem becomes easier.

[5] If one puts a box of side length 1 centered on every integer point, then the set of boxes put on all points a distance of $\leq \beta$ away from the origin is completely covered by a ball of radius $\beta + \sqrt{N}/2$. Thus the volume of the sphere is greater than the combined volume of the boxes; and the latter is equal to the number of points.

If we now look at the $\mathsf{DKS}^{\infty}_{n,k,\beta}$ problem, then the best current attack requires finding a close vector to a target in Λ. In case the input is of the form $(\boldsymbol{A}', \boldsymbol{t})$ for a uniform $\boldsymbol{t}$, then the target vector will be uniformly distributed in space. On the other hand, if $\boldsymbol{t} = [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y}$ for a $\boldsymbol{y}$ with small coefficients, then the target vector will be close to Λ. Deciding between the two distributions involves finding a lattice point close to the target and looking at the distance. Lemma 4 essentially states that if β becomes too big, then $\boldsymbol{t} = [\ \boldsymbol{I}_n \quad \boldsymbol{A}'\ ] \cdot \boldsymbol{y}$ will have the same distribution as a uniform $\boldsymbol{t}$, thus making the problem unsolvable. It is also known that as β becomes smaller, the problem becomes easier.

A visual representation of the above discussion is represented in Figure 1. Due to the fact that one problem is in the ℓ_2 norm, while the other is in the ℓ_∞ norm, the graph should only be seen as a visualization of the fact that as norm of the vector $\boldsymbol{y}$ increases, the DKS^{∞} problem becomes harder, while the SKS^{2} problem becomes easier. One should not infer anything about the actual hardness of these problems based on the slopes in the picture. The only important thing is that for some value, the hardness of the two problems becomes roughly the same. This rough visualization will be useful in the next section for explaining the strategy for the optimal setting of parameters.

4 The Commitment Scheme with a Proof of Opening

Parameter	Explanation
$R = \mathbb{Z}[X]/\langle X^N + 1\rangle$	The ring over which we define the norms of vectors
$R_q = \mathbb{Z}_q[X]/\langle X^N + 1\rangle$	The ring over which we do most of the computations
q	Prime modulus defining R_q
k	Width (over R_q) of the commitment matrices
n	Height (over R_q) of the commitment matrix $\boldsymbol{A}_1$
ℓ	Dimension (over R_q) of the message space
β	Norm bound for honest prover's randomness in ℓ_∞-norm
S_β	Set of all elements $x \in R$ with ℓ_∞-norm at most β
$\mathcal{C}$	A subset of S_1 from which challenges come from (see (1))
$\bar{\mathcal{C}}$	The set of differences $\mathcal{C} - \mathcal{C}$ excluding 0
κ	The maximum ℓ_1 norm of any element in $\mathcal{C}$
$\sigma = 11 \cdot \kappa \cdot \beta \cdot \sqrt{kN}$	Standard deviation used in the zero-knowledge proof

Table 1. Overview of Parameters and Notation.

4.1 The Commitment Scheme

Our commitment scheme can be seen as a particular instantiation of the scheme due to Damgård et al. [DPP93]. A "wrinkle" in our scheme is that the opening of the commitment does not simply involve producing the message with the randomness that was used in the commitment. The reason is that we do not have efficient zero-knowledge proofs that can prove knowledge of simply the message and the randomness that was used to commit. The zero-knowledge proof can prove something weaker, and therefore our commitment scheme should still be binding with such a relaxed opening.[6]

[6] This was also the property in the commitment scheme of [BKLP15].

One thing to notice is that the randomness is generated according to a distribution using the ℓ_∞ norm, whereas the opening is using the ℓ_2 norm. The reason for this "mismatch" is that the most efficient lattice-based zero-knowledge proofs prove the knowledge of small vectors in the ℓ_2 norm. On the other hand, when committing, it is simpler to just use the ℓ_∞ norm. If one wishes to use the ℓ_2 or the ℓ_∞ norm everywhere, the scheme is easily modifiable.

KeyGen*:* We will create public parameters that can be used to commit to messages $\boldsymbol{x} \in R_q^\ell$. Create $\boldsymbol{A}_1 \in R_q^{n \times k}$ and $\boldsymbol{A}_2 \in R_q^{\ell \times k}$ as

$$\boldsymbol{A}_1 = [\ \boldsymbol{I}_n \quad \boldsymbol{A}_1'\ ], \quad \text{where } \boldsymbol{A}_1' \overset{\$}{\leftarrow} R_q^{n \times (k-n)} \tag{5}$$

$$\boldsymbol{A}_2 = [\ \boldsymbol{0}^{\ell \times n} \quad \boldsymbol{I}_\ell \quad \boldsymbol{A}_2'\ ], \quad \text{where } \boldsymbol{A}_2' \overset{\$}{\leftarrow} R_q^{\ell \times (k-n-\ell)} \tag{6}$$

Commit*:* To commit to $\boldsymbol{x} \in R_q^\ell$, we choose a random polynomial vector $\boldsymbol{r} \overset{\$}{\leftarrow} S_\beta^k$ and output the commitment

$$Com(\boldsymbol{x}; \boldsymbol{r}) := \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix} \cdot \boldsymbol{r} + \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x} \end{bmatrix} \tag{7}$$

Open*:* A valid opening of a commitment $\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix}$ is a 3-tuple consisting of an $\boldsymbol{x} \in R_q^\ell$, $\boldsymbol{r} = \begin{bmatrix} r_1 \\ \dots \\ r_k \end{bmatrix} \in R_q^k$, and $f \in \bar{\mathcal{C}}$. The verifier checks that

$$f \cdot \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix} \cdot \boldsymbol{r} + f \cdot \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x} \end{bmatrix},$$

and that for all i, $\|r_i\|_2 \le 4\sigma\sqrt{N}$.

4.2 Hiding and Binding

The hiding property of the scheme is based on the $\mathsf{DKS}^\infty_{n+\ell,k,\beta}$ problem.

Lemma 6. *For any $\boldsymbol{x}, \boldsymbol{x}' \in R^\ell$, if there exists an algorithm $\mathcal{A}$ that has advantage ϵ in breaking the hiding property of the commitment scheme, then there exists another algorithm $\mathcal{A}'$ that runs in the same time and has advantage ϵ in solving the $\mathsf{DKS}^\infty_{n+\ell,k,\beta}$ problem.*

Proof. Given an instance $\boldsymbol{B} = [\ \boldsymbol{I}_{n+\ell} \mid \boldsymbol{B}'\ ], \boldsymbol{t}$ of the $\mathsf{DKS}^\infty_{n+\ell,k,\beta}$ problem, the algorithm $\mathcal{A}'$ obtains the messages $\boldsymbol{x}_0, \boldsymbol{x}_1 \in R_q^\ell$ from $\mathcal{A}$, generates a random matrix $\boldsymbol{R} \in R_q^{n \times \ell}$, generates a bit $b \overset{\$}{\leftarrow} \{0,1\}$ and outputs the public parameters of the scheme as

$$\begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{R} \\ \boldsymbol{0}^{\ell \times n} & \boldsymbol{I}_\ell \end{bmatrix} \cdot \boldsymbol{B},$$

and the commitment of $\boldsymbol{x}_b$ as

$$\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{R} \\ \boldsymbol{0}^{\ell \times n} & \boldsymbol{I}_\ell \end{bmatrix} \cdot \boldsymbol{t} + \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x}_b \end{bmatrix}. \tag{8}$$

If $\mathcal{A}$ returns $b = b'$, then $\mathcal{A}'$ outputs 1 (and outputs 0 if $b \neq b'$).

We first want to show that the public parameters $\boldsymbol{A}_1, \boldsymbol{A}_2$ are correctly distributed. If we rewrite

$$\boldsymbol{B} = \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{0}^{n\times\ell} & \boldsymbol{B}_1' \\ \boldsymbol{0}^{\ell\times n} & \boldsymbol{I}_\ell & \boldsymbol{B}_2' \end{bmatrix},$$

then we can see that

$$\begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{R} \\ \boldsymbol{0}^{\ell\times n} & \boldsymbol{I}_\ell \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{0}^{n\times\ell} & \boldsymbol{B}_1' \\ \boldsymbol{0}^{\ell\times n} & \boldsymbol{I}_\ell & \boldsymbol{B}_2' \end{bmatrix} = \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{R} & \boldsymbol{B}_1' + \boldsymbol{R}\cdot\boldsymbol{B}_2' \\ \boldsymbol{0}^{\ell\times n} & \boldsymbol{I}_\ell & \boldsymbol{B}_2' \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix}.$$

Since $\boldsymbol{B}_1', \boldsymbol{B}_2', \boldsymbol{R}$ are all uniform and independent, it's clear that the above distribution of $\boldsymbol{A}_1, \boldsymbol{A}_2$ is exactly as in the KeyGen algorithm.

It should also be clear that if $\boldsymbol{t}$ is uniformly random, then the "commitment" $\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix}$ in (8) is independent of $\boldsymbol{x}_b$, and in this case the probability that $\mathcal{A}$ can output a $b' = b$ is exactly $1/2$.

If, on the other hand, $\boldsymbol{t} = \boldsymbol{B}\cdot\boldsymbol{r}$, then the commitment in (8) is

$$\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{R} \\ \boldsymbol{0}^{\ell\times n} & \boldsymbol{I}_\ell \end{bmatrix} \cdot \boldsymbol{B}\cdot\boldsymbol{r} + \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x}_b \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix} \cdot \boldsymbol{r} + \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x}_b \end{bmatrix} = Com(\boldsymbol{x}_b; \boldsymbol{r}).$$

Thus $\mathcal{A}$ should output $b' = b$ with probability at least $1/2 + \epsilon$. Therefore the advantage of $\mathcal{A}'$ in solving the $\mathsf{DKS}^\infty_{n+\ell,k,\beta}$ problem is at least ϵ. □

The next lemma shows that the binding property of the scheme is based on the SKS^2 problem.

Lemma 7. *If there is an algorithm $\mathcal{A}$ who can break the binding of the commitment scheme with probability ϵ, then there is an algorithm $\mathcal{A}'$ who can solve the $\mathsf{SKS}^2_{n,k,16\sigma\sqrt{\kappa N}}$ problem with advantage ϵ.*

Proof. Given an instance $\boldsymbol{A}_1 = [\ \boldsymbol{I}_n \quad \boldsymbol{A}_1'\ ]$ of the $\mathsf{SKS}^2_{n,k,\gamma}$ problem, the algorithm $\mathcal{A}'$ creates a random $\boldsymbol{A}_2$ as in (6) and sets $\boldsymbol{A}_1, \boldsymbol{A}_2$ as the public parameters of the commitment scheme. If $\mathcal{A}$ is able to come up with a commitment $\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix}$ and valid openings $(\boldsymbol{x}; \boldsymbol{r}; f)$ and $(\boldsymbol{x}'; \boldsymbol{r}'; f')$ with $\boldsymbol{x} \neq \boldsymbol{x}'$ such that

$$f \cdot \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix} \cdot \boldsymbol{r} + f \cdot \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x} \end{bmatrix},$$

$$f' \cdot \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix} \cdot \boldsymbol{r}' + f' \cdot \begin{bmatrix} \boldsymbol{0}^n \\ \boldsymbol{x}' \end{bmatrix},$$

then multiplying the first equation by f', the second by f and subtracting, we would obtain

$$\boldsymbol{A}_1 \cdot (f'\cdot\boldsymbol{r} - f\cdot\boldsymbol{r}') = \boldsymbol{0}^n. \tag{9}$$

$$\boldsymbol{A}_2 \cdot (f'\cdot\boldsymbol{r} - f\cdot\boldsymbol{r}') + (f\cdot f'\cdot\bar{\boldsymbol{x}} - f\cdot f'\cdot\bar{\boldsymbol{x}}') = \boldsymbol{0}^\ell. \tag{10}$$

Because f and f' are invertible and $\boldsymbol{x} \neq \boldsymbol{x}'$, we have that $f\cdot f'\cdot(\boldsymbol{x} - \boldsymbol{x}') \neq \boldsymbol{0}^\ell$, and therefore (10) implies that $(f'\cdot\boldsymbol{r} - f\cdot\boldsymbol{r}')$ is also non-zero. Since $f \in \bar{\mathcal{C}}$, we know that $\|f\|_2 \leq 2\sqrt{\kappa}$, and since every polynomial r_i comprising $\boldsymbol{r}$ has ℓ_2 norm bounded by $4\sigma\sqrt{N}$, we have $\|f'\cdot r_i\|_2, \|f\cdot r_i'\|_2 \leq 8\sigma\sqrt{\kappa N}$, and thus the norm of the difference is at most $16\sigma\sqrt{\kappa N}$. And (9) therefore gives a solution to the $\mathsf{SKS}^2_{n,k,16\sigma\sqrt{\kappa N}}$ instance. □

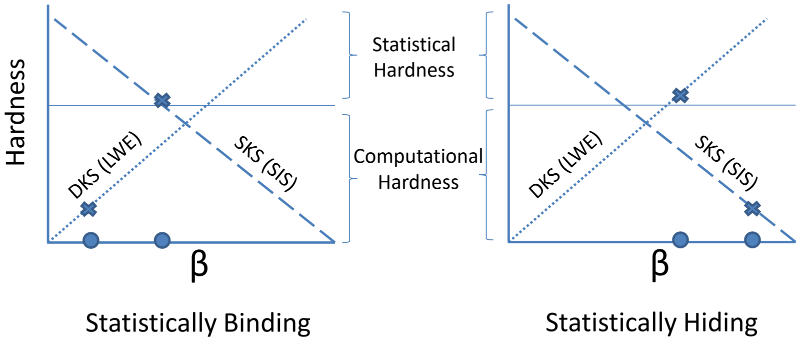

Fig. 2. Setting parameters for statistically binding and statistically hiding versions of the scheme. In each graph, the leftmost circle is the parameter β in the hardness of the DKS^∞ problem, and the right circle is the parameter β in the SKS^2 problem. The crosees correspond to the security of these problems with the particular parameters.

4.3 Instantiations

There are three "interesting" ways in which one can instantiate the commitment scheme. If we would like the scheme to be statistically-hiding, then Lemma 6 implies that the $\mathsf{DKS}^\infty_{n+\ell,k,\beta}$ problem should be difficult even for all-powerful adversaries. Lemma 4 then describes exactly how the parameters of the scheme should be set. Statistically-hiding schemes are therefore based entirely on the hardness of the SKS^2 (or equivalently, Module-SIS) problem.

On the other hand, if we would like the scheme to be statistically binding, then Lemma 7 states that it's enough for the $\mathsf{SKS}^2_{n,k,16\sigma\sqrt{\kappa N}}$ problem to be unconditionally hard. Lemma 5, in turn, dictates the setting of parameters. The statistically-binding variant of the scheme is therefore based entirely on the DKS^∞ (or equivalently, Module-LWE) problem.

Figure 2 presents a visualization of how one needs to choose the parameters of the commitment scheme in order to achieve statistical binding/hiding. In both instances, the left circle indicates the parameter β for the DKS^∞ problem which controls the hardness of breaking the hiding property of the scheme. The right circle is the value of $\beta = 16\sigma\sqrt{\kappa N}$ for the SKS^2 problem.

The third way in which we can instantiate the scheme is, from a practical perspective, the most notable. The ability to instantiate our scheme in this manner is the main advantage of this scheme over the construction in [BKLP15]. While the structure of the commitment scheme in [BKLP15] required the scheme to be statistically binding, our construction has the freedom to move the "circles" in Figure 2 arbitrarily along the horizontal axis (with the restriction that the distance between them is preserved). If one measures the security of the commitment scheme by the weakest of the hiding and binding (as is the natural way to measure security), then it makes sense to set the hardness of the two to be the same. A visual sketch of this is given in Figure 3. We point out that it does not matter what the exact "slopes" representing the hardness of the DKS^∞ and SKS^2 problems are. Since these two lines (or curves) intersect, the minimal hardness in either of the variants in Figure 2 can always be raised by shifting the "circles" to either the left or the right.

The above intuition for setting the parameters has been used for signature schemes since [Lyu12]. Signature schemes constructed in such fashion (where the hardness of recovering the secret key is based on LWE and the hardness of forgery is based on SIS) turn out to provide significant savings over those that are just based on LWE or SIS. To get a rough estimate of the savings, we can look at a recent work [KLS17] that instantiated the same scheme based on either entirely the (Module)-LWE assumption, or the optimal way based on both Module-LWE and Module-SIS. The size of the signature (which will correspond to the size of the zero-knowledge proof in our scheme) was around 5× shorter in the scheme that optimized its parameter settings using the intuition in Figure 3. Since the constructions of signatures and commitments with proofs of opening are fairly similar, one would expect a similar magnitude of savings for our commitment scheme over the one in [BKLP15] as well. We present this comparison in Section 5.

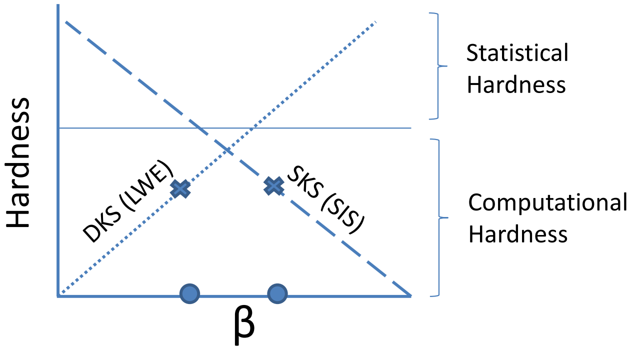

Fig. 3. Optimal setting of the parameters for the commitment scheme.

4.4 Zero-Knowledge Protocols

We will now give a zero-knowledge proof of knowledge of a valid opening. The protocol is almost identical to those underlying the constructions of digital signature schemes using the "Fiat-Shamir with Aborts" [Lyu09] approach. In particular, the proof is a 3-move Σ-protocol in which an honest prover sometimes needs to abort for security reasons. It can be shown that *non-aborting* interactions are honest-verifier zero-knowledge, and the protocol itself is a proof of knowledge. The fact that only non-aborting interactions are zero-knowledge does not cause a problem in practice. The interactive protocol is usually converted to a non-interactive one using the Fiat-Shamir transform, in which case the aborting transcripts are never seen. If one wishes to keep the protocol interactive, one can slightly change it by making the prover apply an auxiliary commitment to the first move, and opening the commitment in the last. The above-described transformation techniques are standard, and so we only present the underlying interactive protocol.

Proof for Opening a Commitment. Below, we will look at the properties of this protocol.

Lemma 8. *The protocol Π_{OPEN} has the following properties:*

- Completeness: *The verifier accepts with overwhelming probability when Π_{OPEN} does not abort. The probability of abort is at most $1 - \frac{1-2^{-100}}{M}$.*

$\underline{\mathit{\Pi}_{\text{Open}}}$

Public Instance-Specific Information: $\boldsymbol{A} = \begin{bmatrix}\boldsymbol{A}_1 \\ \boldsymbol{A}_2\end{bmatrix}$ as in $(5), (6)$ defining $Com(\cdot;\cdot)$.
Prover's Information: $\boldsymbol{r} \in S_\beta^k$
Commitment: $\boldsymbol{c} = \begin{bmatrix}\boldsymbol{c}_1 \\ \boldsymbol{c}_2\end{bmatrix} = Com(\boldsymbol{x}; \boldsymbol{r})$ as in (7).

$$\begin{array}{lcl}
\underline{\text{Prover}} & & \underline{\text{Verifier}} \\
\boldsymbol{y} \xleftarrow{\$} \mathcal{N}_\sigma^k & & \\
\boldsymbol{t} := \boldsymbol{A}_1 \cdot \boldsymbol{y} & & \\
 & \xrightarrow{\boldsymbol{t}} & \\
 & & d \xleftarrow{\$} \mathcal{C} \\
 & \xleftarrow{d} & \\
\boldsymbol{z} = \boldsymbol{y} + d \cdot \boldsymbol{r} & & \\
\text{Abort with probability} & & \\
\quad 1 - \min\left(1, \frac{\mathcal{N}_\sigma^k(\boldsymbol{z})}{M \cdot \mathcal{N}_{\mathsf{d}r,\sigma}^k(\boldsymbol{z})}\right) & & \\
 & \xrightarrow{\boldsymbol{z}} & \\
 & & \text{Write } \boldsymbol{z} = \begin{bmatrix} z_1 \\ \dots \\ z_k \end{bmatrix} \\
 & & \text{Accept iff } \forall i,\ \|z_i\|_2 \le 2\sigma\sqrt{N} \text{ and} \\
 & & \boldsymbol{A}_1 \cdot \boldsymbol{z} = \boldsymbol{t} + d \cdot \boldsymbol{c}_1
\end{array}$$

Fig. 4. Zero-Knowledge Proof of Opening.

- Special Soundness: *Given a commitment* $\boldsymbol{c}$ *and a pair of transcripts for* $\mathit{\Pi}_{\text{Open}}$ $(\boldsymbol{t}, d, \boldsymbol{z})$, $(\boldsymbol{t}, d', \boldsymbol{z}')$ *where* $d \neq d'$, *we can extract a valid opening* $\left(\boldsymbol{x}, \boldsymbol{r} = \begin{bmatrix} r_1 \\ \dots \\ r_k \end{bmatrix}, f\right)$ *of* $\boldsymbol{c}$ *with* $||r_i||_2 \le 4\sigma\sqrt{N}$, *and* $f \in \bar{\mathcal{C}}$.
- Honest-Verifier Zero-Knowledge: *Non-aborting transcripts of* $\mathit{\Pi}_{\text{Open}}$ *with an honest verifier can be simulated with statistically indistinguishable distribution.*

Proof. Completeness: An honest prover can clearly answer correctly for any challenge d and by Lemma 2, the abort probability of the prover for our choice of parameters is at most $1 - \frac{1-2^{-100}}{M}$. For the verifier, by Remark 1 the bound on the ℓ_2-norm of every polynomial z_i comprising $\boldsymbol{z}$ is $2\sigma \cdot \sqrt{N}$ except with negligible probability.

Special Soundness: Notice that two valid transcripts for different challenges d, d' allows the computation of an $f = (d - d') \in \bar{\mathcal{C}}$ and an $\boldsymbol{r} = \begin{bmatrix} r_1 \\ \dots \\ r_k \end{bmatrix} = \boldsymbol{z} - \boldsymbol{z}'$ such that $\boldsymbol{A}_1 \cdot \boldsymbol{r} = f \cdot \boldsymbol{c}_1$. We define the message contained in $\boldsymbol{c}$ as $\boldsymbol{x} = \boldsymbol{c}_2 - f^{-1} \cdot \boldsymbol{A}_2 \cdot \boldsymbol{r}$. Since $\|r_i\|_2 \le \|z_i\|_2 + \|z_i'\|_2 \le 4\sigma\sqrt{N}$ and $\begin{bmatrix}\boldsymbol{A}_1 \\ \boldsymbol{A}_2\end{bmatrix} \cdot \boldsymbol{r} + f \cdot \begin{bmatrix}\boldsymbol{0}^n \\ \boldsymbol{x}\end{bmatrix} = f \cdot \begin{bmatrix}\boldsymbol{c}_1 \\ \boldsymbol{c}_2\end{bmatrix}$, the opening $(\boldsymbol{x}, \boldsymbol{r}, f)$ is valid.

Honest-Verifier Zero-Knowledge:

To simulate an accepting conversation, draw a random d from $\mathcal{C}$ and a random $\boldsymbol{z}$ from $\mathcal{N}_\sigma^k$. Set $\boldsymbol{t} = \boldsymbol{A}_1\boldsymbol{z} - d\boldsymbol{c}_1$. This distribution is statistically indistinguishable from the real non-aborting transcript as the simulator simply acts as $\mathcal{S}$ as in Lemma 2. □

In addition to the zero-knowledge protocol described above, we can also give protocols that prove knowledge of various other properties of the commitment. Most of these protocols are fairly straight-forward to construct using the additive-homomorphic property of the commitment scheme. We only provide brief sketches here.

Proof for Opening to a Specific Message. The protocol Π_{OPEN} demonstrates that the prover knows how to open a commitment, without revealing either the randomness or the message. An easy variant, which we will call $\Pi_{\text{OPEN-X}}$, can be used to show that the prover can open $\boldsymbol{c}$ to a specific message $\boldsymbol{x}$: it is enough to show that a commitment can be opened to 0, since one can use that protocol on input $\boldsymbol{c} - \mathit{Com}(\boldsymbol{x}; \boldsymbol{0})$. Now, to prove that a commitment can be opens to 0, the verifier makes an additional check in Π_{OPEN} to make sure that $\boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{t} + d \cdot \boldsymbol{c}_2$.

Proof for Linear Relation. Suppose that the prover has published two commitments $\boldsymbol{c}_1 = \mathit{Com}(\boldsymbol{x}_1; \boldsymbol{r}_1), \boldsymbol{c}_2 = \mathit{Com}(\boldsymbol{x}_2; \boldsymbol{r}_2)$ and claims that $\boldsymbol{x}_2 = g \cdot \boldsymbol{x}_1$ for for some $g \in R_q$. The protocol Π_{LIN} for proving this relation is similar to running Π_{OPEN} on two separate commitments, but the prover's first message and the verifier's check also contains the relationship between the two. The protocol is given in Figure 5.

From two valid transcripts, we can recover $\boldsymbol{r}, \boldsymbol{r}', f$ such that

$$\boldsymbol{A}_1 \cdot \boldsymbol{r} = f \cdot \boldsymbol{c}_1 \tag{11}$$

$$\boldsymbol{A}_1 \cdot \boldsymbol{r}' = f \cdot \boldsymbol{c}_1' \tag{12}$$

$$g \cdot \boldsymbol{A}_2 \cdot \boldsymbol{r} - \boldsymbol{A}_2 \cdot \boldsymbol{r}' = f \cdot (g \cdot \boldsymbol{c}_2 - \boldsymbol{c}_2') \tag{13}$$

and define

$$\boldsymbol{x} = \boldsymbol{c}_2 - f^{-1} \cdot \boldsymbol{A}_2 \cdot r \tag{14}$$

$$\boldsymbol{x}' = \boldsymbol{c}_2' - f^{-1} \cdot \boldsymbol{A}_2 \cdot r' \tag{15}$$

as in the proof of Lemma 8. As in that proof, this implies that $(\boldsymbol{x}, \boldsymbol{r}, f)$ is a valid opening for $\begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix}$ (and analogously for $\boldsymbol{x}', \boldsymbol{r}', f$). The relationship $\boldsymbol{x}' = g \cdot \boldsymbol{x}$ is derived from plugging in the values of $\boldsymbol{c}_2, \boldsymbol{c}_2'$ from (14) and (15) into (13).

Proof for Sum. Suppose that the prover has published three commitments $\boldsymbol{c}_1 = \mathit{Com}(\boldsymbol{x}_1; \boldsymbol{r}_1)$, $\boldsymbol{c}_2 = \mathit{Com}(\boldsymbol{x}_2; \boldsymbol{r}_2)$, $\boldsymbol{c}_3 = \mathit{Com}(\boldsymbol{x}_3; \boldsymbol{r}_3)$ and claims that $\boldsymbol{x}_3 = \alpha_1 \cdot \boldsymbol{x}_1 + \alpha_2 \cdot \boldsymbol{x}_2$ where $\alpha_1, \alpha_2 \in R_q$ are public constants. The protocol Π_{SUM} is very similar to the previous protocol.

Achieving Zero-Knowledge for Dishonest Verifiers. One easy way to have our protocols be zero-knowledge against dishonest verifiers is if a trusted source of random bits is available (which can be implemented via a coin-flipping protocol). One gets the challenge from this source and then clearly honest-verifier zero-knowledge is sufficient.

$\underline{\mathit{\Pi}_{\text{LIN}}}$

Public Instance-Specific Information: $\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \end{bmatrix}$ as in (5), (6) defining $Com(\cdot;\cdot)$.

Prover's Information: $\boldsymbol{r}, \boldsymbol{r}' \in S_\beta^k, \boldsymbol{x} \in R_q^k, g \in R_q$

Commitment: $\boldsymbol{c} = \begin{bmatrix} \boldsymbol{c}_1 \\ \boldsymbol{c}_2 \end{bmatrix} = Com(\boldsymbol{x}; \boldsymbol{r}), \boldsymbol{c}' = \begin{bmatrix} \boldsymbol{c}_1' \\ \boldsymbol{c}_2' \end{bmatrix} = Com(g \cdot \boldsymbol{x}; \boldsymbol{r}')$, as in (7).

Prover | | Verifier

$\boldsymbol{y}, \boldsymbol{y}' \xleftarrow{\$} \mathcal{N}_\sigma^k$

$\boldsymbol{t} := \boldsymbol{A}_1 \cdot \boldsymbol{y}$, $\boldsymbol{t}' := \boldsymbol{A}_1 \cdot \boldsymbol{y}'$

$\boldsymbol{u} := g \cdot \boldsymbol{A}_2 \cdot \boldsymbol{y} - \boldsymbol{A}_2 \cdot \boldsymbol{y}'$ $\xrightarrow{\boldsymbol{t}, \boldsymbol{t}', \boldsymbol{u}}$

$d \xleftarrow{\$} \mathcal{C}$

$\xleftarrow{d}$

$\boldsymbol{z} = \boldsymbol{y} + d \cdot \boldsymbol{r}$, $\boldsymbol{z}' = \boldsymbol{y}' + d \cdot \boldsymbol{r}'$

Abort with probability

$1 - \min\left(1, \frac{\mathcal{N}_\sigma^k(\boldsymbol{z})}{M \cdot \mathcal{N}_{\mathrm{d}r,\sigma}^k(\boldsymbol{z})}\right)$

Abort with probability

$1 - \min\left(1, \frac{\mathcal{N}_\sigma^k(\boldsymbol{z}')}{M \cdot \mathcal{N}_{\mathrm{d}r',\sigma}^k(\boldsymbol{z}')}\right)$

$\xrightarrow{\boldsymbol{z}, \boldsymbol{z}'}$

Write $\boldsymbol{z} = \begin{bmatrix} z_1 \\ \ldots \\ z_k \end{bmatrix}$, $\boldsymbol{z}' = \begin{bmatrix} z_1' \\ \ldots \\ z_k' \end{bmatrix}$

Accept iff $\forall i$, $\|\boldsymbol{z}_i\|_2, \|\boldsymbol{z}_i'\|_2 \leq 2\sigma\sqrt{N}$ and

$\boldsymbol{A}_1 \cdot \boldsymbol{z} = \boldsymbol{t} + d \cdot \boldsymbol{c}_1$, $\boldsymbol{A}_1 \cdot \boldsymbol{z}' = \boldsymbol{t}' + d \cdot \boldsymbol{c}_1'$ and

$g \cdot \boldsymbol{A}_2 \cdot \boldsymbol{z} - \boldsymbol{A}_2 \cdot \boldsymbol{z}' = (g \cdot \boldsymbol{c}_2 - \boldsymbol{c}_2') \cdot d + \boldsymbol{u}$

Fig. 5. Zero-Knowledge Proof of a Linear Relation.

A different approach is possible if a trapdoor commitment scheme Com_{trap} is available, where commitments in this scheme can be equivocated if the trapdoor is known. Then we can transform each of our protocols to a new one that is zero-knowledge: the prover commits to the first message $\boldsymbol{t}$ using Com_{trap}, gets the challenge d, then opens $Com_{trap}(\boldsymbol{t})$ and answers d. If the simulator knows the trapdoor, it can make a fake commitment first. Once d arrives, it runs the simulation and equivocates the initial commitment to the value of $\boldsymbol{t}'$ that it wants.

5 Instantiation and Comparison

In this section, we give concrete parameters for which our scheme can be instantiated and give a comparison to the previous commitment scheme from [BKLP15]. In fact, we provide an improved analysis of the binding property in [BKLP15] which allows it to be instantiated with smaller parameters. We then compare the scheme in this paper for concrete parameter settings to the one from [BKLP15].

5.1 The commitment scheme from [BKLP15].

We will now restate the scheme from [BKLP15] and then compare its instantiation to the scheme introduced in this paper.

Commitment and Opening.

KeyGen*:* We will create public parameters that can be used to commit to messages $\boldsymbol{x} \in R_q^\ell$. Create $\boldsymbol{A}_1 \in R_q^{n\times k}$ and $\boldsymbol{A}_2 \in R_q^{n\times \ell}$ as

$$\boldsymbol{A}_1 = [\ \boldsymbol{I}_n \quad \boldsymbol{A}_1'\ ], \quad \text{where } \boldsymbol{A}_1' \xleftarrow{\$} R_q^{n\times(k-n)} \tag{16}$$

$$\boldsymbol{A}_2 \xleftarrow{\$} R_q^{n\times\ell} \tag{17}$$

Commit*:* To commit to $\boldsymbol{x} \in R_q^\ell$, we choose a random polynomial vector $\boldsymbol{r} \xleftarrow{\$} S_\beta^k$ and output the commitment

$$Com(\boldsymbol{x};\boldsymbol{r}) := \boldsymbol{c} = \boldsymbol{A}_1 \cdot \boldsymbol{r} + \boldsymbol{A}_2 \cdot \boldsymbol{x} \tag{18}$$

Open*:* A valid opening of a commitment $\boldsymbol{c}$ is a 3-tuple consisting of an $\boldsymbol{x} \in R_q^\ell$, $\boldsymbol{r} = \begin{bmatrix} r_1 \\ \dots \\ r_k \end{bmatrix} \in R_q^k$, and $f \in \bar{\mathcal{C}}$. The verifier checks that

$$f \cdot \boldsymbol{c} = \boldsymbol{A}_1 \cdot \boldsymbol{r} + f \cdot \boldsymbol{A}_2 \cdot \boldsymbol{x},$$

and that for all i, $\|r_i\|_2 \leq 4\sigma\sqrt{N}$.

Hiding and Binding. The hiding property of the scheme is based on the $\mathsf{DKS}^\infty_{n,k,\beta}$ problem.

Lemma 9. *For any $\boldsymbol{x}, \boldsymbol{x}' \in R^\ell$, if there exists an algorithm $\mathcal{A}$ that has advantage ϵ in breaking the hiding property of the commitment scheme, then there exists another algorithm $\mathcal{A}'$ that runs in the same time and has advantage ϵ in solving the $\mathsf{DKS}^\infty_{n,k,\beta}$ problem.*

Because the coefficients of the vector $\boldsymbol{x}$ are not bounded, we do not know how to base the binding property on a computational assumption, since the latter generally involves finding a short solution to a linear equation. Instead, it was shown in [BKLP15, Theorem 3.1] how one can set the parameters so that binding is statistical. The below lemma is an improvement over this result. The improvement comes from the fact that we can use Lemma 1 to argue that elements with small norms are invertible in the ring. Without this improvement, the very last term in the below lemma would be $q^{(\ell-n/d)\cdot N}$ instead of $q^{(\ell-n)\cdot N}$. Even for the smallest possible value of $d = 2$, this would require significantly larger parameters to satisfy the condition in which the term appears.

Lemma 10. *Let $q = 2d + 1 (\mathrm{mod}\ 4d)$ for some d that's a power of 2, and $16\sigma\sqrt{\kappa N} < q^{1/d}$. Except with probability*

$$q^{(\ell-n/d)\cdot N} + 2^{1024} \cdot \left(\sqrt{2\pi e} \cdot (4\sigma + 1/2)\right)^{2k\cdot N} \cdot q^{(\ell-n)\cdot N}$$

over the random choices of $\boldsymbol{A}_1, \boldsymbol{A}_2$, there does not exist a commitment $\boldsymbol{c}$ that can be opened with $(\boldsymbol{x}, \boldsymbol{r}, f)$ and $(\boldsymbol{x}', \boldsymbol{r}', f')$ for distinct $\boldsymbol{x} \neq \boldsymbol{x}'$.

Proof. The existence of such two openings implies that

$$f \cdot \boldsymbol{c} = \boldsymbol{A}_1 \cdot \boldsymbol{r} + f \cdot \boldsymbol{A}_2 \cdot \boldsymbol{x}$$
$$f' \cdot \boldsymbol{c} = \boldsymbol{A}_1 \cdot \boldsymbol{r}' + f' \cdot \boldsymbol{A}_2 \cdot \boldsymbol{x}',$$

which can be rewritten as

$$\boldsymbol{A}_1 \cdot (f'\boldsymbol{r} - f\boldsymbol{r}') + ff' \cdot \boldsymbol{A}_2 \cdot (\boldsymbol{x} - \boldsymbol{x}') = \boldsymbol{0}, \tag{19}$$

and by renaming variables as

$$\boldsymbol{I}_n \cdot \boldsymbol{y}_0 + \boldsymbol{A}_1' \cdot \boldsymbol{y}_1 + \boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{0}. \tag{20}$$

As in the proof of Lemma 7, for all polynomials y_i comprising $\boldsymbol{y}_0$ and $\boldsymbol{y}_1$, we have the bound $\|y_i\|_2 \le 16\sigma\sqrt{\kappa N}$. Since a condition in the Lemma states that this quantity is less than $q^{1/d}$, we know that all non-zero y_i are invertible in R_q as per Lemma 1. There is no norm restriction on the coefficients of $\boldsymbol{z}$. We now want to compute the probability over the choice of $\boldsymbol{A}_1, \boldsymbol{A}_2$ that there exists a valid solution $(\boldsymbol{y}_0, \boldsymbol{y}_1, \boldsymbol{z})$ satisfying (20). We will handle this in separate cases.

If some polynomial of $\boldsymbol{y}_1$ is non-zero, then it's easy to see that $\Pr[(20)] = q^{-n \cdot N}$ using exactly the same reasoning as in the proof of Lemma 3.

Now suppose that all coefficients of $\boldsymbol{y}_1$ are 0. Then $\Pr[(20)]$ becomes $\Pr_{\boldsymbol{A}_2}[\boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{y}_0]$. If all the polynomials comprising $\boldsymbol{y}_0$ are 0, then the preceding probability is most $q^{-n \cdot N/d}$. This is because some polynomial comprising $\boldsymbol{z}$ must be non-zero, which implies that it is non-zero modulo one of the factors of $X^N + 1$. The conditions of the Lemma combined with Lemma 1 stipulates that $X^N + 1 = \prod_{i=1}^{d} p_i(X) \bmod q$, where the degree of every $p_i(X)$ in N/d. The probability bound then follows identically to Lemma 3 except working modulo $p_i(X)$ rather than $X^N + 1$.

The last possibility involves computing $\Pr_{\boldsymbol{A}_2}[\boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{y}_0]$ where $\boldsymbol{y}_0$ contains a non-zero polynomial. Since this polynomial has a "small" norm, it is invertible in R_q. We rewrite

$$\begin{aligned}
\Pr_{\boldsymbol{A}_2}[\boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{y}_0] &= \Pr_{\boldsymbol{A}_2}\left[\bigwedge_{i=1}^{d} (\boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{y}_0 \bmod p_i(X))\right] & (21)\\
&= \prod_{i=1}^{d} \left(\Pr_{\boldsymbol{A}_2}[\boldsymbol{A}_2 \cdot \boldsymbol{z} = \boldsymbol{y}_0 \bmod p_i(X)] \right), & (22)
\end{aligned}$$

where the last equality is true due to the fact that all computation modulo $p_i(X)$ are disjoint for all i. Since for all i, $\boldsymbol{y}_0 \bmod p_i(X)$ is non-zero in at least one place, there must be a non-zero element of $\boldsymbol{z} \bmod p_i(X)$ as well. Then by the same argument as above, we will have the probability of each multiplicand in (22) is at most $q^{-n \cdot N/d}$, making the whole product $q^{-n \cdot N}$.

Summarizing, if both $\boldsymbol{y}_0$ and $\boldsymbol{y}_1$ are zero, then $\Pr[(20)] = q^{-n \cdot N/d}$. Otherwise, $\Pr[(20)] = q^{-n \cdot N}$. There are at most $q^{\ell \cdot N}$ possible $\boldsymbol{z}$, and so using the union bound, we have

$$\Pr_{\boldsymbol{A}_1, \boldsymbol{A}_2}[\exists \boldsymbol{x}, \boldsymbol{x}' \text{ s.t. } (19) \mid f'\boldsymbol{r} - f\boldsymbol{r}' = \boldsymbol{0}] \le \Pr_{\boldsymbol{A}_2}[\exists \boldsymbol{z} \text{ s.t. } (20) \mid \boldsymbol{y}_0, \boldsymbol{y}_1 = \boldsymbol{0}] \le q^{N(\ell - n/d)}. \tag{23}$$

From the proof of Lemma 5, we know that there are fewer than

$$V_N(\beta + \sqrt{N}/2)^k < \left(\sqrt{\frac{2\pi e}{N}} \cdot (\beta + \sqrt{N}/2) \right)^{k \cdot N}$$

elements $\boldsymbol{r} \in R^k$ such that $\forall i, \|r_i\|_2 \le \beta = 4\sigma \cdot \sqrt{N}$, where $V_N(\alpha)$ is the volume of an N-dimensional ball of radius α. Furthermore, since the challenge set $\mathcal{C}$ has size 2^{256}, the size of $\bar{\mathcal{C}}$ is at most 2^{512}. Therefore

$$\Pr_{\boldsymbol{A}_1,\boldsymbol{A}_2}[\exists f, f', \boldsymbol{r}, \boldsymbol{r}', \boldsymbol{x}, \boldsymbol{x}' \text{ s.t. } (19) \mid f'\boldsymbol{r} - f\boldsymbol{r}' \neq \boldsymbol{0}] \tag{24}$$

$$< 2^{1024} \cdot \left(\sqrt{\frac{2\pi e}{N}} \cdot (\beta + \sqrt{N}/2)\right)^{2k\cdot N} \cdot q^{\ell \cdot N} \cdot q^{-n\cdot N}. \tag{25}$$

The claim in the lemma follows by summing (23) and (25). □

Zero-Knowledge Protocol. The zero-knowledge protocol from [BKLP15] for proving knowledge of an opening $(\boldsymbol{x}; \boldsymbol{r}; f)$ in presented in Figure 6. The main difference with the protocol from Figure 4 is that it also needs to mask the message $\boldsymbol{x}$ in addition to the randomness $\boldsymbol{r}$. Proving zero-knowledge and special soundness of this protocol is very similar as to that of the protocol in Figure 4 and is proved in [BKLP15].

$\underline{\mathit{\Pi}_{\text{OPEN}}}$

Public Instance-Specific Information: $\boldsymbol{A} = \begin{bmatrix}\boldsymbol{A}_1 & \boldsymbol{A}_2\end{bmatrix}$ as in (16), (17) defining $Com(\cdot;\cdot)$.
Prover's Information: $\boldsymbol{r} \in S_\beta^k$, $\boldsymbol{x} \in R_q^\ell$
Commitment: $\boldsymbol{c} = Com(\boldsymbol{x}; \boldsymbol{r})$ as in (18).

<u>Prover</u>		<u>Verifier</u>
$\boldsymbol{y}_r \overset{\$}{\leftarrow} \mathcal{N}_\sigma^k$, $\boldsymbol{y}_x \overset{\$}{\leftarrow} R_q^\ell$ $\boldsymbol{t} := \boldsymbol{A}_1 \cdot \boldsymbol{y}_r + \boldsymbol{A}_2 \cdot \boldsymbol{y}_x$		
	$\xrightarrow{\boldsymbol{t}}$	
		$d \overset{\$}{\leftarrow} \mathcal{C}$
	$\xleftarrow{d}$	
$\boldsymbol{z}_r = \boldsymbol{y}_r + d \cdot \boldsymbol{r}$ $\boldsymbol{z}_x = \boldsymbol{y}_x + d \cdot \boldsymbol{x}$ Abort with probability $1 - \min\left(1, \frac{\mathcal{N}_\sigma^k(\boldsymbol{z_r})}{M \cdot \mathcal{N}_{\mathrm{d}r,\sigma}^k(\boldsymbol{z_r})}\right)$		
	$\xrightarrow{\boldsymbol{z}_r,\ \boldsymbol{z}_x}$	
		Write $\boldsymbol{z} = \begin{bmatrix} z_1 \\ \dots \\ z_k \end{bmatrix}$ Accept iff $\forall i$, $\|\boldsymbol{z}_i\|_2 \leq 2\sigma\sqrt{N}$ and $\boldsymbol{A}_1 \cdot \boldsymbol{z}_r + \boldsymbol{A}_2 \cdot \boldsymbol{z}_x = \boldsymbol{t} + d \cdot \boldsymbol{c}$

Fig. 6. Zero-Knowledge Proof of Opening for the [BKLP15] protocol.

5.2 Concrete Instantiations

We now describe the parameters for instantiating the non-interactive versions of the protocols in Figure 4 and Figure 6. For the protocol in Figure 4, the commitment size is $N \cdot (n+\ell) \cdot \log q$ bits. The output of the non-interactive proof consists of $(d, \boldsymbol{z})$. Every coefficient of $\boldsymbol{z}$ is chosen according to a discrete Gaussian with standard deviation σ, and so for simplicity we can assume that all coefficients are smaller than 6σ (each coefficient is less than 6σ with probability at least $1 - 2 \cdot \exp(-18)$, and

parameter	This Paper (Figure 4)	[BKLP15] (Figure 6)
q	$\approx 2^{32}$	$\approx 2^{71}$
N	1024	1024
ℓ	1	1
d	2 or 4 or 8	2
n	1	6
k	3	9
κ	44	44
β (in S_β)	1	1
σ	26800	46464
hermite factor	≈ 1.0035	≈ 1.0035
commit. size	8.1 KB	54.5 KB
proof size	6.6 KB	30 KB

Table 2. Parameter settings for our scheme (from Figure 4) and the one from [BKLP15] (Figure 6)

.

the prover can simply try again in the unlikely event that some coefficient is larger). Therefore the size of the proof is approximately $N \cdot k \cdot \log 6\sigma$ bits. For the protocol in Figure 6, the commitment size is $N \cdot n \cdot \log q$ bits, while the proof length (of $d, \boldsymbol{z}_r, \boldsymbol{z}_r$)) is $N \cdot k \cdot \log 6\sigma + N \cdot \ell \cdot \log q$ bits.

In Table 2, we give a possible set of parameters for our scheme (based on the optimal setting as discussed in Section 4.3) and the one from [BKLP15]. The exact parameters that one would use in practice depends on the application in which one would use zero-knowledge proofs of commitments. The main purpose of the table is to give a comparison between the two techniques for similar security levels. For simplicity, we use the Hermite factor approach from [GN08] to compute the hardness of the schemes. We set the target Hermite factor to be around 1.0035. Even though this approach for setting parameters does not take into account the recent lattice reduction approaches that use sieving (that take theoretically less time at the expense of a lot more memory) in addition to enumeration, we believe that this methodology is adequate for comparison purposes.

5.3 Further Improvements

It is possible to reduce the size of the proofs of the protocols in Figures 4 and 6 by using the compression techniques in [GLP12,BG14]. The main idea in those works is that the prover does not need to send the part of the proof that gets multiplied by the identity matrix part of $\boldsymbol{A}_1$, and the verifier only checks an approximate equality. A very rough calculation shows that one could reduce the proof size in the protocol in Figure 4 to 4.4KB and to 15.8KB in the protocol from [BKLP15].

One can also use technique from [DKL+18] to reduce the size of the commitment in our protocol in Figure 4 (it's unclear if this can also be applied to [BKLP15]). The idea is that one can drop the low-order bits of the commitment $\boldsymbol{c}_1$ and the SKS^2/ SIS problems would still remain almost as hard as before. This technique can reduce the commitment size to around 6KB.

References

AJL+12. Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold

FHE. In *Advances in Cryptology - EUROCRYPT 2012*, pages 483–501, 2012.

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 99–108. ACM, 1996.

BCK$^+$14. Fabrice Benhamouda, Jan Camenisch, Stephan Krenn, Vadim Lyubashevsky, and Gregory Neven. Better zero-knowledge proofs for lattice encryption and their application to group signatures. In *Advances in Cryptology - ASIACRYPT 2014*, pages 551–572, 2014.

BG14. Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In *Topics in Cryptology - CT-RSA 2014 - The Cryptographer's Track at the RSA Conference 2014, San Francisco, CA, USA, February 25-28, 2014. Proceedings*, pages 28–47, 2014.

BKLP15. Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In *Computer Security - ESORICS 2015*, pages 305–325, 2015.

Blu82. Manuel Blum. Coin flipping by telephone - A protocol for solving impossible problems. In *COMPCON'82, Digest of Papers, Twenty-Fourth IEEE Computer Society International Conference, San Francisco, California, USA, February 22-25, 1982*, pages 133–137, 1982.

CFSY96. Ronald Cramer, Matthew K. Franklin, Berry Schoenmakers, and Moti Yung. Multi-autority secret-ballot elections with linear work. In *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*, pages 72–83, 1996.

DF89. Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, pages 307–315, 1989.

DKL$^+$18. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.

dPLNS17. Rafaël del Pino, Vadim Lyubashevsky, Gregory Neven, and Gregor Seiler. Practical quantum-safe voting from lattices. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 1565–1581, 2017.

dPLS18. Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. Zero-knowledge proofs of automorphism stability and applications to lattice-based privacy protocols. In Submission., 2018.

DPP93. Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. In *Advances in Cryptology - CRYPTO '93*, pages 250–265, 1993.

GLP12. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, pages 530–547, 2012.

GN08. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, pages 31–51, 2008.

JKPT12. Abhishek Jain, Stephan Krenn, Krzysztof Pietrzak, and Aris Tentes. Commitments and efficient zero-knowledge proofs from learning parity with noise. In *Advances in Cryptology - ASIACRYPT 2012*, pages 663–680, 2012.

KLS17. Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. Cryptology ePrint Archive, Report 2017/916, 2017. `http://eprint.iacr.org/2017/916`. To appear in Eurocrypt 2018.

KTX08. Akinori Kawachi, Keisuke Tanaka, and Keita Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In *Advances in Cryptology - ASIACRYPT 2008, 14th International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, Australia, December 7-11, 2008. Proceedings*, pages 372–389, 2008.

LM06. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, pages 144–155, 2006.

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, pages 1–23, 2010.

LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.

LS17. Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. Cryptology ePrint Archive, Report 2017/523, 2017. `https://eprint.iacr.org/2017/523`. To appear in Eurocrypt 2018.

Lyu09. Vadim Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, pages 598–616, 2009.

Lyu12. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 738–755, 2012.

Mic07. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007.

PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 145–166, 2006.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93, 2005.

XXW13. Xiang Xie, Rui Xue, and Minqian Wang. Zero knowledge proofs from ring-lwe. In *Cryptology and Network Security - 12th International Conference, CANS 2013, Paraty, Brazil, November 20-22. 2013. Proceedings*, pages 57–73, 2013.