



A non-PCP Approach to Succinct Quantum-Safe Zero-Knowledge [★]

Jonathan Bootle^{1,3}, Vadim Lyubashevsky¹, Ngoc Khanh Nguyen^{1,2}, and Gregor Seiler^{1,2}

¹ IBM Research – Zurich, Switzerland
nkn@zurich.ibm.com,  [0000-0001-8240-6167]
vad@zurich.ibm.com

² ETH Zurich, Switzerland
gseiler@inf.ethz.ch

³ UC Berkeley, USA
jonathan.bootle@berkeley.edu,  [0000-0003-3582-3368]

Abstract. Today’s most compact zero-knowledge arguments are based on the hardness of the discrete logarithm problem and related classical assumptions. If one is interested in quantum-safe solutions, then all of the known techniques stem from the PCP-based framework of Kilian (STOC 92) which can be instantiated based on the hardness of any collision-resistant hash function. Both approaches produce asymptotically logarithmic sized arguments but, by exploiting extra algebraic structure, the discrete logarithm arguments are a few orders of magnitude more compact in practice than the generic constructions.

In this work, we present the first (poly)-logarithmic, potentially *post-quantum* zero-knowledge arguments that deviate from the PCP approach. At the core of succinct zero-knowledge proofs are succinct commitment schemes (in which the commitment and the opening proof are sub-linear in the message size), and we propose two such constructions based on the hardness of the (Ring)-Short Integer Solution (Ring-SIS) problem, each having certain trade-offs. For commitments to N secret values, the communication complexity of our first scheme is $\tilde{O}(N^{1/c})$ for any positive integer c , and $O(\log^2 N)$ for the second. Both of these are a significant theoretical improvement over the previously best lattice construction by Bootle et al. (CRYPTO 2018) which gave $O(\sqrt{N})$ -sized proofs.

Keywords. Lattices, Zero-Knowledge Proofs, SNARKS

1 Introduction

Zero-knowledge proofs are a crucial component in many cryptographic protocols. They are essential to electronic voting, verifiable computation, cryptocurrencies, and for adding stronger security and privacy guarantees to digital signature

[★] This work was supported by the SNSF ERC Transfer Grant CRETP2-166734 – FELICITY. The work was done while the first author was at IBM Research – Zurich.

and encryption schemes. Across almost all applications, it is important to be able to prove in zero-knowledge that one knows how to open a cryptographic commitment, and to prove that the committed values have particular properties or satisfy some relations.

Recent years have seen an explosion of new zero-knowledge proof techniques, each with improvements in proof-size, proving time, or verification time. These new constructions are based on a variety of cryptographic assumptions, including the discrete logarithm assumption [14], various pairing-based assumptions in the Generic and Algebraic Group Models [26,27], collision-resistant hash functions [8,7], and lattice-based assumptions such as (R)SIS and (R)LWE [13,20,3,19].

Of these, only constructions from hash-functions and lattices stand any chance of being post-quantum secure. At this point in time, general-purpose lattice-based proof systems still lag far behind, both asymptotically and in practice, in proof-size and usability. This may seem somewhat surprising, since unlike hash-functions, lattices are endowed with algebraic structure that allows for constructions of rather efficient encryption [34], signature [17,37], and identity-based encryption schemes [24,18]. One could hope that the additional lattice structure can be also exploited for succinct zero-knowledge proofs as well.

1.1 Our Contribution

In this paper, we present two novel lattice-based commitment schemes with associated zero-knowledge opening protocols which prove knowledge of N secret integers with $\tilde{O}(N^{1/c})$ and $\tilde{O}(\log^2 N)$ communication complexity, for any constant c . For the former argument, we sketch out a method for constructing an argument of knowledge of a satisfying assignment for an arithmetic circuit with N gates, with $\tilde{O}(N^{1/c})$ communication complexity in the full version of this paper. Both arguments of knowledge follow the same basic methodology, which is to replace the homomorphic commitment schemes used in earlier classically-secure protocols with a commitment scheme based on the (Ring)-SIS problem, and adapt the security proofs to match.

Our constructions follow the usual framework of being interactive schemes converted to non-interactive ones using the Fiat-Shamir transform. As with many schemes constructed in this fashion, their security is proven in the ROM rather than the QROM. The latter would be very strong evidence of quantum security,⁴ but the former also appears to give strong evidence of quantum security in practice. To this day, there is no example of a practical scheme that has been proven secure in the ROM based on a quantum-safe computational assumption that has shown any weakness when the adversary was given additional quantum access to the random oracle. A recent line of works (e.g. [28,16,30,15]) that prove security in the QROM of schemes using the Fiat-Shamir transform which previously only known to be secure in the ROM give further evidence that security in the ROM based on a quantum-safe assumption is a meaningful security notion in practice.

⁴ Though still technically heuristic because of the assumption that a concrete hash function acts as a random oracle.

Our first construction extends the classical interactive argument of [25] in which the prover commits to message values using Pedersen commitments, and then commits to those commitments using a pairing-based commitment scheme. The two-level structure means that a clever commitment-opening procedure is possible, giving $\tilde{O}(N^{1/3})$ communication costs. With a d -level commitment scheme, one could hope to extend the technique and construct an argument with $\tilde{O}(N^{1/(d+1)})$ -sized proofs. However, in [25], Pedersen commitments map finite field elements to source group elements, which are mapped to target group elements by the second commitment scheme. In the classical setting, this is as far as the technique can take us, as no homomorphic, compressing commitment scheme for target group elements is known. In the lattice setting, however, the message space for SIS commitments are small integers and commitments are just made up of larger integers. So there is no fundamental reason not to continue. Using careful manipulation of matrices and moduli, our first new argument extends this technique to any constant number of levels.

The second argument is based on the techniques in the Bulletproofs protocol [14], and an earlier protocol [12], which use Pedersen commitments to commit to long message vectors. The additional structure of the Pedersen commitment scheme allows a neat folding technique, which reduces the length of committed message vectors by a factor of two. The prover and verifier repeatedly employ the technique over logarithmically many rounds of interaction until message vectors are reduced to a single value, which the prover can then easily send to the verifier. This gives logarithmic proof sizes. Our new lattice protocol stems from the observation that a SIS-based commitment scheme has some structure similarity to the Pedersen commitment scheme, and thus can be made compatible with the same folding technique. A technical complication that is unique to the lattice setting is keeping the coefficients of the extracted values from growing (too much) during each fold, as a direct adaptation of the bulletproof technique would result in unconstrained growth for every fold which would make the proof meaningless.

Finally, we make a comparison of these two techniques in terms of commitment/proof sizes as well as sizes of the extracted solutions, alternatively called “slack”. Our conclusion is that the Bulletproofs folding argument offers smaller poly-logarithmic proof size at the cost much larger slack. Hence, if one does not necessarily need their extracted solution to be very small, then using Bulletproofs appears to be more suitable. However, in many applications, such as group signatures or verifiable encryption, zero-knowledge proofs are just one part of a more complex scheme. If the extracted witnesses are large, then we must adjust parameters not only for the zero-knowledge proof but also for other components of the scheme. Thus, we believe the leveled commitments can be applied in such scenarios at the cost of slightly larger proofs than lattice-based Bulletproofs.

Discussion and Open Problems. The ultimate goal of the line of research that this paper is pursuing is constructing zero-knowledge proofs with concrete parameters smaller than those that can be achieved following the PCP approach. Our current results achieve parameters that are essentially the same asymptot-

ically, but are larger in practice. The asymptotic equivalence comes from the fact that we succeeded in making the *dimension* of the vector(s) representing the proof be logarithmic in the message size. And while we have also somewhat restricted the *coefficient* growth of the proof vector, the coefficients still grow by some factor with each “folding” of the vector dimension. Finding a technique to even further restrict the coefficient growth is the main open problem stemming from this work.

From experience with other primitives, using the additional algebraic structure of concrete assumptions should (eventually) result in size advantages over the generic PCP-based approaches that have the implicit lower bounds (of around 100-200KB) posed by using Merkle-tree commitments. While lattice-based constructions may not achieve the extreme compactness of discrete logarithm based approaches (e.g. Bulletproofs, which have proofs sizes of a few kilobytes for reasonably-sized circuits), there is reason to hope that they can be shorter (and faster) than generic constructions. As an analogy, when lattice-based signatures first appeared [24,33], they were significantly larger than the generic quantum-safe signatures that one could construct using techniques, such as one-way functions and Merkle trees, dating back to the 1970s [29,35]. But expanding upon these early lattice constructions via novel algorithms and techniques exploiting the underlying mathematical structure of lattices, the current state-of-the-art lattice-based signatures [17,37] are currently an order of magnitude smaller and two orders of magnitude faster than those stemming from generic constructions [10]. We believe that the techniques of this paper can similarly be the beginning of the path to more practical succinct quantum-safe zero-knowledge.

1.2 Technical Overview

Levelled Commitments. The commitment scheme in [5] arranged N elements of \mathbb{Z}_q to which one wants to commit to, into an $m \times k$ matrix \mathbf{S} and created the commitment

$$\mathbf{A} \cdot \mathbf{S} = \mathbf{T} \text{ mod } q \tag{1}$$

where $\mathbf{A} \leftarrow \mathbb{Z}_p^{n \times m}$ is a random matrix and $p < q$. Notice that $\mathbf{T} \in \mathbb{Z}_q^{n \times k}$, and [5] showed that the proof of knowledge of an \mathbf{S} with small (but larger than the honest prover uses in the proof) coefficients satisfying (1) can be done with λm elements in \mathbb{Z}_q ⁵ where λ is a security parameter. The total size of the proof is therefore the size of \mathbf{T} and the size of the proof of (1), which is $nk + \lambda m$ elements in \mathbb{Z}_q . Since $n = \mathcal{O}(\lambda)$, the optimal way to commit to N elements in \mathbb{Z}_q is to arrange them into a matrix $\mathbf{S} \in \mathbb{Z}_q^{m \times k}$, where $m = k = \tilde{\mathcal{O}}(\sqrt{N})$. This makes the total proof size $\tilde{\mathcal{O}}(\sqrt{N})$.

To illustrate our levelled commitment technique, we will describe a commitment scheme and a protocol for achieving a proof size of $\tilde{\mathcal{O}}(N^{1/3})$. We will

⁵ We provide additional background in Section 2.3 for readers not familiar with previous work.

commit to $\mathbf{S} \in \mathbb{Z}^{m_1 \cdot m_2 \times m_3}$ as $\mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \mathbf{S} \bmod q_2) \bmod q_1 = \mathbf{T}$ where $\mathbf{A}_1 \leftarrow \mathbb{Z}_{q_1}^{n \times nm_1}$, $\mathbf{A}_2 \leftarrow \mathbb{Z}_{q_2}^{n \times m_2}$. Our proof will prove knowledge of an $\bar{\mathbf{S}}$ with somewhat larger coefficients than \mathbf{S} , and also an $\bar{\mathbf{R}} \in \mathbb{Z}^{n \cdot m_1 \times m_3}$ satisfying

$$\mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}} \bmod q_2 + \bar{\mathbf{R}} \cdot q_2) \bmod q_1 = \mathbf{T}, \quad (2)$$

Let us first show that the above extracted commitment of $(\bar{\mathbf{S}}, \bar{\mathbf{R}})$ is binding based on the hardness of SIS when $\|\bar{\mathbf{S}}\| \ll q_2$, $\|\bar{\mathbf{R}}\| \ll q_1/q_2$ and $q_2 \ll q_1$. Suppose, for contradiction, there are two $(\bar{\mathbf{S}}, \bar{\mathbf{R}}) \neq (\bar{\mathbf{S}}', \bar{\mathbf{R}}')$ satisfying (2). In the first case, suppose that $\bar{\mathbf{R}} \neq \bar{\mathbf{R}}'$. By definition, the coefficients of $(\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}} \bmod q_2$ are smaller than q_2 , and thus $\bar{\mathbf{R}} \neq \bar{\mathbf{R}}'$ implies that

$$(\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}} \bmod q_2 + \bar{\mathbf{R}} \cdot q_2 \neq (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}}' \bmod q_2 + \bar{\mathbf{R}}' \cdot q_2. \quad (3)$$

If the parameters are set such that the coefficients of both sides of the above equation are less than q_1 , then this gives a solution to SIS for \mathbf{A}_1 . Now assume that $\bar{\mathbf{R}} = \bar{\mathbf{R}}'$, and so $\bar{\mathbf{S}} \neq \bar{\mathbf{S}}'$. If $(\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}} \equiv (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}}' \pmod{q_2}$, then there must be some $\bar{\mathbf{S}}_i \neq \bar{\mathbf{S}}'_i \in \mathbb{Z}^{m_2 \times m_3}$ such that $\mathbf{A}_2 \cdot \bar{\mathbf{S}}_i \equiv \mathbf{A}_2 \cdot \bar{\mathbf{S}}'_i \pmod{q_2}$, and so we have a SIS solution for \mathbf{A}_2 . If $(\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}} \not\equiv (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \bar{\mathbf{S}}' \pmod{q_2}$, then the inequality in (3) holds and we have a SIS solution for \mathbf{A}_1 .

We present the basic protocol in Fig. 1. The boxed text contains the parts necessary to make the protocol zero-knowledge. In this overview, we will ignore these and only show that the protocol is a proof of knowledge. First, let us show the correctness of the protocol. Because the coefficients of \mathbf{S} , \mathbf{C}_1 , and \mathbf{C}_2 are small, the coefficients of \mathbf{Z} are also small with respect to q_2 . Similarly, because the coefficients of \mathbf{V} consist of a product of a matrix with coefficients less than q_2 with a 0/1 matrix \mathbf{C}_1 , parameters can be set such that the coefficients of the product are less than q_1 . Thus $\|\mathbf{z}_i\| \leq \beta_z$ and $\|\mathbf{V}\| \leq \beta_v q_2$ can be satisfied with an appropriate choice of parameters. We now move on to showing that the verification equations hold. Note that

$$\mathbf{V} = \begin{bmatrix} \mathbf{V}_1 \\ \vdots \\ \mathbf{V}_{m_1} \end{bmatrix} = \left(\begin{bmatrix} \mathbf{A}_2 \mathbf{S}_1 \\ \vdots \\ \mathbf{A}_2 \mathbf{S}_{m_1} \end{bmatrix} \bmod q_2 \right) \cdot \mathbf{C}_1 \equiv \begin{bmatrix} \mathbf{A}_2 \mathbf{S}_1 \mathbf{C}_1 \\ \vdots \\ \mathbf{A}_2 \mathbf{S}_{m_1} \mathbf{C}_1 \end{bmatrix} \pmod{q_2}, \quad (4)$$

and so one can write $\mathbf{A}_2 \cdot [\mathbf{S}_1 \mathbf{C}_1 \cdots \mathbf{S}_{m_1} \mathbf{C}_1] \equiv [\mathbf{V}_1 \cdots \mathbf{V}_{m_1}] \pmod{q_2}$, and therefore $\mathbf{A}_2 \cdot \mathbf{Z} \equiv \mathbf{A}_2 \cdot [\mathbf{S}_1 \mathbf{C}_1 \cdots \mathbf{S}_{m_1} \mathbf{C}_1] \cdot \mathbf{C}_2 \pmod{q_2}$, which is the first verification equation. For the second verification equation, observe from (4) and (2) that

$$\mathbf{A}_1 \cdot \mathbf{V} = \mathbf{A}_1 \cdot \left(\begin{bmatrix} \mathbf{A}_2 \mathbf{S}_1 \\ \vdots \\ \mathbf{A}_2 \mathbf{S}_{m_1} \end{bmatrix} \bmod q_2 \right) \cdot \mathbf{C}_1 \equiv \mathbf{T} \cdot \mathbf{C}_1 \pmod{q_1}.$$

Finally, ignoring constant terms, the total communication cost (including the statement) can be bounded above by

$$m_3 \cdot (n \log q_1 + \lambda) + m_2 \cdot \lambda \log \beta_z + m_1 \cdot (n \lambda \log(\beta_v q_2) + \lambda^2) + n \lambda \log q_2.$$

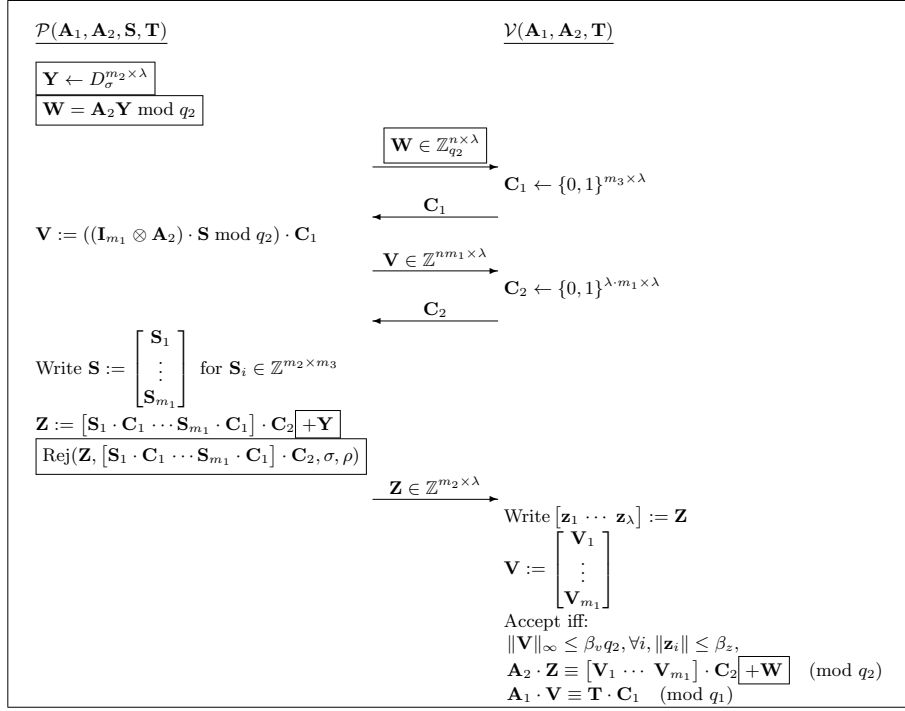


Fig. 1. Levelled commitment with two levels. Here, β_v and β_z are parameters which satisfy $\beta_z \ll q_2, \beta_v \ll q_1/q_2$.

Note that the last term does not depend on m_1, m_2, m_3 hence we ignore it for now. Therefore, in order to minimise the expression above, we want to set m_1, m_2, m_3 such that all three corresponding terms are (almost) equal. We select appropriate n, q_1, q_2 such that both \mathbf{A}_1 and \mathbf{A}_2 are binding (7) and we get $\log q_2 < \log q_1 = \mathcal{O}(\log N)$ and $n = \mathcal{O}(\lambda)$. Similarly, $\log \beta_v = \mathcal{O}(\log N)$ and $\log \beta_z = \mathcal{O}(\log N)$. Therefore, the total communication cost is approximately $\tilde{\mathcal{O}}(\sqrt[3]{N})$.

In Section 3, we extend this approach to more than two levels. Generally, we propose a proof of knowledge for $d \geq 1$ levels with total communication size equal to $\mathcal{O}\left(N^{\frac{1}{d+1}} \cdot (d^3 \lambda \log^2 N + d\lambda^2)\right)$. In Section 3.3, we also show how to apply techniques similar to previous work [14,12,5] in order to extract a relatively short solution to the relaxed equation (e.g. (2) for $d = 2$). Due to space limitations, we skip the details in this overview.

Bulletproofs folding. Our starting point is the lattice equation:

$$\mathbf{A}\mathbf{s} = \mathbf{t} \tag{5}$$

where $\mathbf{A} \in R^{1 \times k}$, $\mathbf{s} \in R^k$ and $R = \mathbb{Z}_q[X]/(X^n + 1)$. Thus, the number of secrets is $N = kn$. In the same vein as [14,12], we are interested in constructing a protocol where proving knowledge of pre-image \mathbf{s} of \mathbf{t} comes down to proving knowledge of some other pre-image, say \mathbf{s}' , whose length $k/2$ is half that of \mathbf{s} . By recursively applying this argument $\log k$ times, we obtain poly-logarithmic proof size. Concretely, we fold the initial statement (5) as follows. Let us write $\mathbf{A} = [\mathbf{A}_1 \ \mathbf{A}_2]$ and $\mathbf{s} = \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \end{bmatrix}$ where $\mathbf{s}_1, \mathbf{s}_2 \in R^{k/2}$.

Let $\mathbf{l} := \mathbf{A}_1 \mathbf{s}_2 \in R$ and $\mathbf{r} := \mathbf{A}_2 \mathbf{s}_1 \in R$. Then, for all $c \in R$, $(c\mathbf{A}_1 + \mathbf{A}_2)(\mathbf{s}_1 + c\mathbf{s}_2) = c^2 \mathbf{l} + c\mathbf{t} + \mathbf{r}$. We observe that $\mathbf{s}_1 + c\mathbf{s}_2$ has length $k/2$, suggesting the following protocol for proving knowledge of \mathbf{s} . First, the prover \mathcal{P} sends \mathbf{l}, \mathbf{r} to the verifier \mathcal{V} . Then, \mathcal{V} samples a challenge c uniformly at random from a challenge space $\mathcal{C} \subseteq R$ and sends it to \mathcal{P} . Finally, \mathcal{P} sends $\mathbf{z} = \mathbf{s}_1 + c\mathbf{s}_2$ to \mathcal{V} . Note that if \mathcal{P} is honest then \mathbf{z} satisfies the new lattice equation $\mathbf{B}\mathbf{z} = \mathbf{t}'$ where $\mathbf{B} = c\mathbf{A}_1 + \mathbf{A}_2$ and $\mathbf{t}' = c^2 \mathbf{l} + c\mathbf{t} + \mathbf{r}$. Therefore, instead of sending \mathbf{z} directly, the prover might repeat the protocol and treat \mathbf{z} as a new secret vector. By folding all the way down to vectors of length 1, we get communication costs of order $\log k$.

Using similar techniques to [14,12], one can extract a solution $\bar{\mathbf{z}}$ to the original equation $\mathbf{A}\bar{\mathbf{z}} = \mathbf{t}$. The problem is that unless we define the challenge space properly, we do not have any information on the size of $\|\bar{\mathbf{z}}\|$. Hence, we let \mathcal{C} be the set of monomials of R , i.e. $\mathcal{C} = \{X^i : i \in \mathbb{Z}\}$. Then, using the fact that polynomials of the form $2/(X^i - X^j) \in R$ have coefficients in $\{-1, 0, 1\}$ [9], we bound the size of an extracted solution. The only drawback of using this approach is that we only obtain a solution for a relaxed equation. Concretely, if we apply the folding technique d times, then we only manage to find a small solution $\bar{\mathbf{z}}$ for the equation $\mathbf{A}\bar{\mathbf{z}} = 8^d \mathbf{t}$ such that $\|\bar{\mathbf{z}}\| = \mathcal{O}(n^{3d} \cdot 12^d \cdot p)$ where $p = \|\mathbf{s}\|_\infty$. For $d = \log k$, the relaxation factor becomes k^3 . The communication cost for the $(2d + 1)$ -round version of the protocol is equal to $N \log(2^d p)/2^d + 2dn \log q$.

Then, we would just pick q which is a little bit larger than the slack. It is worth mentioning that the protocol in its current state gives us soundness error of order $1/n$, hence we would need to repeat it $\lambda/\log n$ times in order to achieve soundness error $2^{-\lambda}$. Therefore, the total proof size can be bounded by $\mathcal{O}\left(\frac{\lambda N \log(2^d p)}{2^d \log n} + \lambda d^2 n\right)$.

Comparison. We investigate in which applications one technique offers asymptotically smaller proof size than the other (see Section 4 for more details). First of all, consider the case when we do not require the extracted solution to be “very small”. Then, levelled commitments for $d = \log N - 1$ levels provide proof size of order $\mathcal{O}(\lambda \log^5 N + \lambda^2 \log N)$.

On the other hand, by applying Bulletproofs folding $d = \log k$ times, we obtain proof size⁶ $\mathcal{O}(\lambda n \log N \cdot (\log N + 1))$. Consequently, the Bulletproofs approach achieves smaller proof size.

⁶ We note that more concrete bounds could be computed. However, this non-tight bound already shows that Bulletproofs folding offers smaller proof size.

Next, consider the case when one could only afford limited slack, i.e. the extracted solution is smaller than some set value $B = N^\alpha > N^2$. First, suppose that $N = \lambda^r$ for some $r \geq 3$ (we expect N to be much bigger than λ). Then, we show that levelled commitments and Bulletproofs provide $\tilde{\mathcal{O}}(N^u)$ and $\tilde{\mathcal{O}}(N^v)$ proof sizes respectively, where $u \approx \frac{1}{(\alpha-2)r}$ and $v \approx 1 - \frac{\alpha-1/2}{3 \log n + 4}$ ⁷.

Assume the allowed slack is small enough that both u and v are larger than $1/\log N$. Then, we just check which one of u, v is bigger⁸. Since $\log n \geq 1$ and for all $r \geq 3$, the function $f_r(x) := (15 - 2x)(x - 2)r - 14$ is positive for $3 \leq x \leq 7$, we deduce that u is smaller than v when $\alpha \leq 7$. This suggests that one should use the levelled commitments protocol when one can only tolerate a limited amount of slack.

1.3 Related Work

In this paper, we investigate techniques from [14,12] and [25] in the lattice world. These papers are the most closely related prior works, along with [14], which forms a key component of the argument in Section 3.

We review proof systems which can prove knowledge of a secret with N elements, or prove knowledge of a satisfying assignment to an arithmetic circuit with N gates.

Lattice-Based Arguments. The zero-knowledge argument given in [5] is based on the SIS assumption, and is capable of proving knowledge of commitment openings with proof size $\mathcal{O}(\sqrt{N})$. It was the first and only standard zero-knowledge protocol based on lattice assumptions to achieve a sublinear communication complexity. Previously, the only other lattice-based arguments of knowledge with better asymptotic proof size were lattice-based SNARKs [11,23,36]. Although they offer highly succinct, $\mathcal{O}(1)$ -sized proofs, the proofs are only checkable by a designated verifier, and soundness is based on strong, non-falsifiable assumptions.

Hash-Based Arguments STARKs [6] and Aurora [8] are non-interactive argument systems. Both achieve $\mathcal{O}(\log^2 N)$ -sized proofs, with Aurora more efficient by an order of magnitude due to better constants. Ligerio [2] achieves $\mathcal{O}(\sqrt{N})$ -sized proofs, but is highly efficient in practice.

Classically-Secure Arguments. In the discrete-logarithm setting, Bulletproofs [14] and a related argument [12] give $\mathcal{O}(\log N)$ communication complexity, using Pedersen commitments and the same recursive folding technique that inspired the argument described in Section 4. The protocol of [25] gives $\mathcal{O}(N^{1/3})$ proof sizes, and uses a two-tiered commitment scheme, based on Pedersen commitments and a related commitment scheme based on pairings. We extend the same idea to a multi-levelled lattice-based commitment scheme in Section 3.

⁷ Here, n denotes the degree of the underlying cyclotomic polynomial $X^n + 1$.

⁸ This would only *asymptotically* tell us which method offers smaller proof size.

There is also a long line of works on succinct non-interactive arguments based on pairings, culminating in protocols including [26] and [27] which have $\mathcal{O}(1)$ proof size, but rely on strong, non-falsifiable assumptions like the Knowledge-of-Exponent assumptions, or have security proofs in idealised models like the Generic Group Model [38] or Algebraic Group Model [21].

2 Preliminaries

Algorithms in our schemes receive a security parameter λ as input (sometimes implicitly) written in unary. Unless stated otherwise, we assume all our algorithms to be probabilistic. We denote by $\mathcal{A}(x)$ the probabilistic computation of the algorithm \mathcal{A} on input x . If \mathcal{A} is deterministic, we write $y := \mathcal{A}(x)$. We write PPT (resp. DPT) for probabilistic (resp. deterministic) polynomial time algorithms. The notation $y \leftarrow \mathcal{A}(x)$ denotes the event that \mathcal{A} on input x returns y . Given two functions $f, g : \mathbb{N} \rightarrow [0, 1]$ we write $f(\lambda) \approx g(\lambda)$ when $|f(\lambda) - g(\lambda)| = \lambda^{-\omega(1)}$. We say that f is *negligible* when $f(\lambda) \approx 0$ and that f is *overwhelming* when $f(\lambda) \approx 1$. For $n \in \mathbb{N}$, we write $[n] := \{1, \dots, n\}$. Regular font letters denote elements in \mathbb{Z} or \mathbb{Z}_q , for a prime q , and bold lower-case letters represent column vectors with coefficients in \mathbb{Z} or \mathbb{Z}_q . Bold upper-case letters denote matrices. By default, all vectors are column vectors. Let $\mathbf{I}_n \in \mathbb{Z}_q^{n \times n}$ be the $n \times n$ identity matrix. We write a list of objects with square brackets, e.g. $[a_1, \dots, a_k]$ is a list of k objects: a_1, \dots, a_k . Also, we denote by $[\]$ the empty list. For any statement st , we define $\llbracket st \rrbracket$ to be equal to 1 if st is true and 0 otherwise.

Sizes of elements. For an even (resp. odd) positive integer α , we define $r' = r \bmod \alpha$ to be the unique element r' in the range $-\frac{\alpha}{2} < r' \leq \frac{\alpha}{2}$ (resp. $-\frac{\alpha-1}{2} \leq r' \leq \frac{\alpha-1}{2}$) such that $r' = r \bmod \alpha$. For an element $w \in \mathbb{Z}_q$, we write $\|w\|_\infty$ to mean $|w \bmod q|$. Define the ℓ_∞ and ℓ_2 norms for $\mathbf{w} = (w_1, \dots, w_k) \in \mathbb{Z}_q^k$ as follows:

$$\|\mathbf{w}\|_\infty = \max_i \|w_i\|_\infty, \quad \|\mathbf{w}\| = \sqrt{\|w_1\|_\infty^2 + \dots + \|w_k\|_\infty^2}.$$

However, if we do not state explicitly that $\mathbf{w} \in \mathbb{Z}_q^k$ but rather treat \mathbf{w} as a vector of integers then the standard notions of L_2 and L_∞ norms apply. We will also consider the operator norm of matrices over \mathbb{Z} defined by $s_1(\mathbf{A}) = \max_{\|\mathbf{x}\| \neq 0} \left(\frac{\|\mathbf{A}\mathbf{x}\|}{\|\mathbf{x}\|} \right)$.

Probability Distributions. Let \mathcal{D} denote a distribution over some set S . Then, $d \leftarrow \mathcal{D}$ means that d was sampled from the distribution \mathcal{D} . If we write $d \stackrel{\$}{\leftarrow} S$ for some finite set S without a specified distribution this means that d was sampled uniformly random from S . We let $\Delta(X, Y)$ indicate the statistical distance between two distributions X, Y . Define the function $\rho_\sigma(x) = \exp\left(\frac{-x^2}{2\sigma^2}\right)$ and the discrete Gaussian distribution over the integers, D_σ , as $D_\sigma(x) = \frac{\rho(x)}{\rho(\mathbb{Z})}$ where $\rho(\mathbb{Z}) = \sum_{v \in \mathbb{Z}} \rho(v)$.

We will write $\mathbf{A} \leftarrow D_\sigma^{k \times \ell}$ to mean that every coefficient of the matrix \mathbf{A} is distributed according to D_σ .

Using the tail bounds for the 0-centered discrete Gaussian distribution (cf. [4]), we can show that for any $\sigma > 0$ the norm of $x \leftarrow D_\sigma$ can be upper-bounded using σ . Namely, for any $t > 0$, we have $\Pr_{x \leftarrow D_\sigma}[|x| > t\sigma] \leq 2e^{-t^2/2}$, and when \mathbf{x} is drawn from D_σ^m , we have

$$\Pr_{\mathbf{x} \leftarrow D_\sigma^m}[\|\mathbf{x}\| > \sqrt{2m} \cdot \sigma] < 2^{-m/4}. \quad (6)$$

2.1 Lattice-based Commitment Schemes

A non-interactive commitment scheme is a pair of PPT algorithms (Gen, Com). The setup algorithm $ck \leftarrow \text{Gen}(1^\lambda)$ generates a commitment key ck , which specifies message, randomness and commitment spaces $\mathbf{M}_{ck}, \mathbf{R}_{ck}, \mathbf{C}_{ck}$. It also specifies an efficiently sampleable probability distribution $D_{\mathbf{R}_{ck}}$ over \mathbf{R}_{ck} and a binding set $\mathbf{B}_{ck} \subset \mathbf{M}_{ck} \times \mathbf{R}_{ck}$. The commitment key also specifies a deterministic polynomial-time commitment function $\text{Com}_{ck} : \mathbf{M}_{ck} \times \mathbf{R}_{ck} \rightarrow \mathbf{C}_{ck}$. We define $\text{Com}_{ck}(\mathbf{m})$ to be the probabilistic algorithm that given $\mathbf{m} \in \mathbf{M}_{ck}$ samples $\mathbf{r} \leftarrow D_{\mathbf{R}_{ck}}$ and returns $\mathbf{c} = \text{Com}_{ck}(\mathbf{m}; \mathbf{r})$.

The commitment scheme is homomorphic, if the message, randomness and commitment spaces are abelian groups (written additively) and we have for all $\lambda \in \mathbb{N}$, and for all $ck \leftarrow \text{Gen}(1^\lambda)$, for all $\mathbf{m}_0, \mathbf{m}_1 \in \mathbf{M}_{ck}$ and for all $\mathbf{r}_0, \mathbf{r}_1 \in \mathbf{R}_{ck}$:

$$\text{Com}_{ck}(\mathbf{m}_0; \mathbf{r}_0) + \text{Com}_{ck}(\mathbf{m}_1; \mathbf{r}_1) = \text{Com}_{ck}(\mathbf{m}_0 + \mathbf{m}_1; \mathbf{r}_0 + \mathbf{r}_1).$$

Definition 2.1 (Hiding). *The commitment scheme is hiding if for all PPT stateful interactive adversaries \mathcal{A}*

$$\Pr \left[ck \leftarrow \text{Gen}(1^\lambda); (\mathbf{m}_0, \mathbf{m}_1) \leftarrow \mathcal{A}(ck); b \leftarrow \{0, 1\}; \right. \\ \left. \mathbf{r} \leftarrow D_{\mathbf{R}_{ck}}; \mathbf{c} \leftarrow \text{Com}_{ck}(\mathbf{m}_b; \mathbf{r}) : \mathcal{A}(\mathbf{c}) = b \right] \approx \frac{1}{2},$$

where \mathcal{A} outputs $\mathbf{m}_0, \mathbf{m}_1 \in \mathbf{M}_{ck}$.

Definition 2.2 (Binding). *The commitment scheme is computationally binding if a commitment can only be opened to one value within the binding set \mathbf{B}_{ck} . For all PPT adversaries \mathcal{A}*

$$\Pr \left[ck \leftarrow \text{Gen}(1^\lambda); (\mathbf{m}_0, \mathbf{r}_0, \mathbf{m}_1, \mathbf{r}_1) \leftarrow \mathcal{A}(ck) : \right. \\ \left. \mathbf{m}_0 \neq \mathbf{m}_1 \text{ and } \text{Com}_{ck}(\mathbf{m}_0; \mathbf{r}_0) = \text{Com}_{ck}(\mathbf{m}_1; \mathbf{r}_1) \right] \approx 0,$$

where \mathcal{A} outputs $(\mathbf{m}_0, \mathbf{r}_0), (\mathbf{m}_1, \mathbf{r}_1) \in \mathbf{B}_{ck}$.

The commitment scheme is compressing if the sizes of commitments are smaller than the sizes of the committed values.

Compressing Commitments Based on SIS. We work with the standard SIS (shortest integer solution) commitment scheme, which was already implicit

in the aforementioned work of Ajtai [1] and uses uniformly random matrices $\mathbf{A}_1 \in \mathbb{Z}_q^{r \times 2r \log_p q}$ and $\mathbf{A}_2 \in \mathbb{Z}_q^{r \times n}$ as a commitment key, where n is the number of elements that one wishes to commit to and $p < q$. A commitment to a vector $\mathbf{m} \in \mathbb{Z}_p^n$ involves choosing a random vector $\mathbf{r} \in \mathbb{Z}_p^{2r \log_p q}$ and outputting the commitment vector $\mathbf{v} = \mathbf{A}_1 \mathbf{r} + \mathbf{A}_2 \mathbf{m} \pmod q$. By the leftover hash lemma, $(\mathbf{A}_1, \mathbf{A}_1 \mathbf{r} \pmod q)$ is statistically close to uniform, and so the commitment scheme is statistically hiding.⁹ To prove binding, note that if there are two different $(\mathbf{r}, \mathbf{m}) \neq (\mathbf{r}', \mathbf{m}')$ such that $\mathbf{v} = \mathbf{A}_1 \mathbf{r} + \mathbf{A}_2 \mathbf{m} \equiv \mathbf{A}_1 \mathbf{r}' + \mathbf{A}_2 \mathbf{m}' \pmod q$, then $\mathbf{A}_1(\mathbf{r} - \mathbf{r}') + \mathbf{A}_2(\mathbf{m} - \mathbf{m}') \equiv \mathbf{0} \pmod q$ and the non-zero vector $\mathbf{s} = \begin{bmatrix} \mathbf{r} - \mathbf{r}' \\ \mathbf{m} - \mathbf{m}' \end{bmatrix}$ is a solution to the SIS problem for the matrix $\mathbf{A} = [\mathbf{A}_1 \ \mathbf{A}_2]$, i.e. $\mathbf{A} \mathbf{s} \equiv \mathbf{0} \pmod q$. As long as the parameters are set such that $\|\mathbf{s}\|$ is smaller than

$$\min\{q, 2^{2\sqrt{r \log q \log \delta}}\}^{10}, \quad (7)$$

the binding property of the commitment is based on an intractable version of the SIS problem [22].

In this paper, we will use the following lattice commitment scheme.

$\text{Gen}(1^\lambda) \rightarrow ck$: Select parameters p, q, r, v, N, B, σ . Pick uniformly random matrices $\mathbf{A}_1 \xleftarrow{\$} \mathbb{Z}_q^{r \times r \log_p q}$ and $\mathbf{A}_2 \xleftarrow{\$} \mathbb{Z}_q^{r \times n}$. Return $ck = (p, q, r, v, \ell, N, \beta, \mathbb{Z}_q, \mathbf{A}_1, \mathbf{A}_2)$. The commitment key defines the message space $\mathcal{M}_{ck} = \mathcal{R}_q^n$, randomness space $\mathcal{R}_{ck} = \mathcal{R}_q^{2r \log_p q}$, commitment space $\mathcal{C}_{ck} = \mathbb{Z}_q^r$, randomness distribution $D_{\mathcal{R}_{ck}} = D_\sigma^r$ and binding space

$$\mathcal{B}_{ck} = \left\{ \mathbf{s} = \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \in \mathcal{R}_q^{n+2r \log_p q} \mid \|\mathbf{s}\| < B \right\}.$$

$\text{Com}_{ck}(\mathbf{m}; \mathbf{r})$: Given $\mathbf{m} \in \mathbb{Z}_q^n$ and $\mathbf{r} \in \mathbb{Z}_q^{2r \log_p q}$ return $\mathbf{c} = \mathbf{A}_1 \mathbf{r} + \mathbf{A}_2 \mathbf{m}$.

In the following, when we make multiple commitments to vectors $\mathbf{m}_1, \dots, \mathbf{m}_\ell \in \mathcal{M}_{ck}$ we write $\mathbf{C} = \text{Com}_{ck}(\mathbf{M}; \mathbf{R})$ when concatenating the commitment vectors as $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_\ell]$. This corresponds to computing $\mathbf{C} = \mathbf{A}_1 \mathbf{R} + \mathbf{A}_2 \mathbf{M}$ with $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_\ell]$ and randomness $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_\ell]$.

2.2 Arguments of Knowledge

We will now formally define arguments of knowledge. Let R be a polynomial-time-decidable ternary relation. The first input will contain public parameters (a.k.a. common reference string) pp . We define the corresponding language L_{pp}

⁹ For improved efficiency, one could reduce the number of columns in \mathbf{A}_1 and make the commitment scheme computationally-hiding based on the hardness of the LWE problem.

¹⁰ This constant δ is related to the optimal block-size in BKZ reduction [22], which is the currently best way of solving the SIS problem. Presently, the optimal lattice reductions set $\delta \approx 1.005$.

indexed by pp that consists of statement u with a witness w such that $(pp, u, w) \in R$. This is a natural generalisation of standard NP languages, which can be cast as the special case of relations that ignore the first input.

A proof system consists of a PPT parameter generator \mathcal{K} , and interactive and stateful PPT algorithms \mathcal{P} and \mathcal{V} used by the prover and verifier. We write $(tr, b) \leftarrow \langle \mathcal{P}(pp), \mathcal{V}(pp, t) \rangle$ for running \mathcal{P} and \mathcal{V} on inputs pp , s , and t and getting communication transcript tr and the verifier's decision bit b . We use the convention that $b = 0$ means reject and $b = 1$ means accept.

Definition 2.3. *Proof system $(\mathcal{K}, \mathcal{P}, \mathcal{V})$ is called an argument of knowledge for the relation R if it is complete and knowledge sound as defined below.*

Definition 2.4. *$(\mathcal{K}, \mathcal{P}, \mathcal{V})$ has statistical completeness with completeness error $\rho : \mathbb{N} \rightarrow [0; 1]$ if for all adversaries \mathcal{A}*

$$\Pr \left[\begin{array}{l} pp \leftarrow \mathcal{K}(1^\lambda); (u, w) \leftarrow \mathcal{A}(pp); (tr, b) \leftarrow \langle \mathcal{P}(pp, u, w), \mathcal{V}(pp, u) \rangle \\ (pp, u, w) \in R \text{ and } b = 0 \end{array} \right] \leq \rho(\lambda).$$

Definition 2.5. *$(\mathcal{K}, \mathcal{P}, \mathcal{V})$ is knowledge sound with knowledge error $\epsilon : \mathbb{N} \rightarrow [0; 1]$ if for all DPT \mathcal{P}^* there exists an expected polynomial time extractor \mathcal{E} such that for all PPT adversaries \mathcal{A}*

$$\Pr \left[\begin{array}{l} pp \leftarrow \mathcal{K}(1^\lambda); (u, s) \leftarrow \mathcal{A}(pp); (tr, b) \leftarrow \langle \mathcal{P}^*(pp, u, s), \mathcal{V}(pp, u) \rangle; \\ w \leftarrow \mathcal{E}^{\mathcal{P}^*(pp, u, s)}(pp, u, tr, b) : (pp, u, w) \notin R \text{ and } b = 1 \end{array} \right] \leq \epsilon(\lambda).$$

It is sometimes useful to relax the definition of knowledge soundness by replacing R with a relation \bar{R} such that $R \subset \bar{R}$. For instance, in this work, our zero-knowledge proofs of pre-images will have “slack”. Thus, even though \mathbf{v} is constructed using \mathbf{r}, \mathbf{m} with coefficients in \mathbb{Z}_p , we will only be able to prove knowledge of vectors $\bar{\mathbf{r}}, \bar{\mathbf{m}}$ with larger norms. This extracted commitment is still binding as long as the parameters are set so that the norm of the vector $\begin{bmatrix} \bar{\mathbf{r}} - \bar{\mathbf{r}}' \\ \bar{\mathbf{m}} - \bar{\mathbf{m}}' \end{bmatrix}$ is smaller than the bound in (7).

We say the proof system is *public coin* if the verifier's challenges are chosen uniformly at random independently of the prover's messages. A proof system is special honest-verifier zero-knowledge if it is possible to simulate the proof without knowing the witness whenever the verifier's challenges are known in advance.

Definition 2.6. *A public-coin argument of knowledge $(\mathcal{K}, \mathcal{P}, \mathcal{V})$ is said to be statistical special honest-verifier zero-knowledge (SHVZK) if there exists a PPT simulator \mathcal{S} such that for all interactive and stateful adversaries \mathcal{A}*

$$\Pr \left[\begin{array}{l} pp \leftarrow \mathcal{K}(1^\lambda); (u, w, \varrho) \leftarrow \mathcal{A}(pp); (tr, b) \leftarrow \langle \mathcal{P}(pp, u, w), \mathcal{V}(\sigma, u; \varrho) \rangle \\ (pp, u, w) \in R \text{ and } \mathcal{A}(tr) = 1 \end{array} \right] \\ \approx \Pr \left[\begin{array}{l} pp \leftarrow \mathcal{K}(1^\lambda); (u, w, \varrho) \leftarrow \mathcal{A}(pp); (tr, b) \leftarrow \mathcal{S}(pp, u, \varrho) \\ (pp, u, w) \in R \text{ and } \mathcal{A}(tr) = 1 \end{array} \right],$$

where ϱ is the randomness used by the verifier.




Fig. 2. Amortized proof for ℓ equations.

2.3 Amortized Proofs of Knowledge

Baum et al. [5] give an amortized proof of knowledge for preimages of SIS commitments (see Fig. 2). The prover \mathcal{P} wants to prove knowledge of the secret matrix \mathbf{S} such that $\mathbf{AS} \equiv \mathbf{T} \pmod{q}$, where \mathbf{A}, \mathbf{T} are known to the verifier \mathcal{V} .

The protocol begins with \mathcal{P} selecting a “masking” value \mathbf{Y} with small coefficients and sending $\mathbf{W} = \mathbf{AY} \pmod{q}$. Then \mathcal{V} picks a random challenge matrix $\mathbf{C} \in \{0, 1\}^{\ell \times n}$, and sends it to \mathcal{P} . Then, \mathcal{P} computes $\mathbf{Z} = \mathbf{SC} + \mathbf{Y}$ and performs a rejection-sampling step (Fig. 3) to make the distribution of \mathbf{Z} independent of \mathbf{S} , and if it passes, sends \mathbf{Z} to \mathcal{V} . Finally, \mathcal{V} checks that all columns of \mathbf{Z} have small norms and that $\mathbf{AZ} \equiv \mathbf{TC} + \mathbf{W} \pmod{q}$.




Fig. 3. Rejection Sampling [31,32].

This protocol can be proved zero-knowledge using exactly the same techniques as in [31,32], i.e. Lemma 2.7. One proves knowledge-soundness using a standard heavy-row argument (Lemma A.1).

Lemma 2.7 ([32]). *Let $\mathbf{B} \in \mathbb{Z}^{r \times n}$ be any matrix. Consider a procedure that samples $\mathbf{Y} \leftarrow D_\sigma^{r \times n}$ and then returns the output of $\text{Rej}(\mathbf{Z} := \mathbf{Y} + \mathbf{B}, \mathbf{B}, \sigma, \rho)$*

where $\sigma \geq \frac{12}{\ln \rho} \cdot \|\mathbf{B}\|$. The probability that this procedure outputs 1 is within 2^{-100} of $1/\rho$. The distribution of \mathbf{Z} , conditioned on the output being 1, is within statistical distance of 2^{-100} of $D_\sigma^{r \times n}$.

By choosing appropriate parameters (r, v, n, ℓ) , Baum et al. obtain a $\tilde{O}(\sqrt{N})$ proof size for the standard SIS commitment scheme where $N = v\ell$ is the number of entries in the matrix \mathbf{S} .

3 Levelled Commitments

In this section, we define levelled lattice commitments and show how to obtain proofs of knowledge with proof size $\tilde{O}(N^{1/c})$ where N is the number of secrets and c is a constant. Recall that Baum et al. [5] give an amortized proof of knowledge for statements of the form $\mathbf{T} = \mathbf{A}\mathbf{S} \bmod q$. We call this a *level-one commitment*. Roughly speaking, the main idea is to apply lattice commitments $c - 1$ times to the secret \mathbf{S} in a structured way.

In the full version of this paper, we extend this result and sketch out the details of an arithmetic circuit satisfiability argument which uses the proof of knowledge based on levelled commitments as a key component.

From now on, we assume that the secret matrix \mathbf{S} already includes the randomness. This not only significantly improves the readability of our protocol, but also ensures that the standard SIS commitment defined in Section 2.1 is both binding and hiding.

3.1 Overview

We define our levelled commitment scheme with d levels for constant d . Let $n, m_0, m_1, \dots, m_d, m_{d+1} \in \mathbb{N}$ such that $m_0 = 1$ and $N = m_1 \cdot \dots \cdot m_{d+1}$. We denote $M_{i,j} = m_i \cdot m_{i+1} \cdot \dots \cdot m_j$ and for simplicity, we write $M_i = M_{0,i}$. Consider d distinct moduli $q_1 > q_2 > \dots > q_d$. Let $\mathbf{A}_d, \dots, \mathbf{A}_1$ be matrices such that $\mathbf{A}_d \in \mathbb{Z}_{q_d}^{n \times m_d}$ and $\mathbf{A}_i \in \mathbb{Z}_{q_i}^{n \times m_i}$ for $i \in [d-1]$. Then, the levelled commitment is a function F defined as follows:

$$F_{i,j}(\mathbf{S}) := \begin{cases} \mathbf{A}_i \mathbf{S} \bmod q_i, & \text{if } i = j \\ F_{i,j-1}((\mathbf{I}_{M_{i,j-1}} \otimes \mathbf{A}_j) \mathbf{S} \bmod q_j) & \text{if } i < j. \end{cases} \quad (8)$$

For example, when $d = 2$, the explicit formula for F is

$$F_{1,2}(\mathbf{S}) = \mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \mathbf{S} \bmod q_2) \bmod q_1. \quad (9)$$

When $d = 3$, the explicit formula for F is

$$F_{1,3}(\mathbf{S}) = \mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) ((\mathbf{I}_{m_1 \cdot m_2} \otimes \mathbf{A}_3) \cdot \mathbf{S} \bmod q_3) \bmod q_2) \bmod q_1. \quad (10)$$

Observe that explicit formulae for F written without tensor notation bear some similarity to Merkle trees of SIS commitments. For instance, if $d = 3$ then

$$\mathbf{T} = F_{1,3} \left(\begin{bmatrix} \mathbf{S}_1 \\ \vdots \\ \mathbf{S}_{m_1 m_2} \end{bmatrix} \right) = \mathbf{A}_1 \cdot \begin{bmatrix} \mathbf{A}_2 \cdot \begin{bmatrix} \mathbf{A}_3 \mathbf{S}_1 \\ \vdots \\ \mathbf{A}_3 \mathbf{S}_{m_2} \end{bmatrix} \\ \vdots \\ \mathbf{A}_2 \cdot \begin{bmatrix} \mathbf{A}_3 \mathbf{S}_{m_1(m_2-1)+1} \\ \vdots \\ \mathbf{A}_3 \mathbf{S}_{m_1 m_2} \end{bmatrix} \end{bmatrix}.$$

Here, \mathbf{T} represents a commitment to the whole tree. In our protocol, the statement will be $F_{1,d}(\mathbf{S}) \equiv \mathbf{T} \pmod{q_1}$, where \mathbf{S} is a matrix consisting of small elements.

For readability, let us introduce commitments for intermediate vertices in this tree. We start from the leaves and denote them as $\mathbf{S}_{[i_1, \dots, i_{d-1}]}$ where each $i_k \in [m_k]$. More concretely, write

$$\mathbf{S} = \begin{bmatrix} \mathbf{S}_{[1, \dots, 1, 1]} \\ \mathbf{S}_{[1, \dots, 1, 2]} \\ \vdots \\ \mathbf{S}_{[1, \dots, 1, m_{d-1}]} \\ \mathbf{S}_{[1, \dots, 2, 1]} \\ \vdots \\ \mathbf{S}_{[m_1, \dots, m_{d-1}]} \end{bmatrix}, \text{ where } \mathbf{S}_{[i_1, \dots, i_{d-1}]} \in \mathbb{Z}^{m_d \times m_{d+1}}. \quad (11)$$

Now we can define commitments for the intermediate vertices in the commitment tree. Fix $k \in [d-2]$ and recursively define

$$\mathbf{S}_{[i_1, \dots, i_k]} := (\mathbf{I}_{m_{k+1}} \otimes \mathbf{A}_{k+2}) \begin{bmatrix} \mathbf{S}_{[i_1, \dots, i_k, 1]} \\ \mathbf{S}_{[i_1, \dots, i_k, 2]} \\ \vdots \\ \mathbf{S}_{[i_1, \dots, i_k, m_{k+1}]} \end{bmatrix} \pmod{q_{k+2}} \in \mathbb{Z}^{n_{m_{k+1}} \times m_{d+1}}. \quad (12)$$

Let us also set $\mathbf{S}_{[]} := (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \begin{bmatrix} \mathbf{S}_{[1]} \\ \vdots \\ \mathbf{S}_{[m_1]} \end{bmatrix} \pmod{q_2} \in \mathbb{Z}^{n_{m_1} \times m_{d+1}}$. Then, we have

$$\mathbf{A}_1 \mathbf{S}_{[]} \equiv \mathbf{T} \pmod{q_1}.$$

Relaxed Opening. Recall that our protocol aims to prove knowledge of a small matrix \mathbf{S} such that $F_{1,d}(\mathbf{S}) \equiv \mathbf{T} \pmod{q_1}$. However, our extraction algorithm finds a slightly larger (but still small) matrix \mathbf{S}' and additional matrices $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$ such that $\tilde{F}_{1,d}(\mathbf{S}'; \mathbf{R}_1, \dots, \mathbf{R}_{d-1}) \equiv \mathbf{T} \pmod{q_1}$ where \tilde{F} is defined by

$$\tilde{F}_{i,j}(\mathbf{S}'; \mathbf{R}_i, \dots, \mathbf{R}_{j-1}) := \tilde{F}_{i,j-1}(\mathbf{X} \pmod{q_j} + q_j \mathbf{R}_{j-1}; \mathbf{R}_i, \dots, \mathbf{R}_{j-2}) \quad (13)$$

and $\mathbf{X} := (\mathbf{I}_{M_{i-1,j-1}} \otimes \mathbf{A}_j) \mathbf{S}'$ for $i < j$ and $\tilde{F}_{i,i}(\mathbf{S}') := (\mathbf{I}_{m_{i-1}} \otimes \mathbf{A}_i) \mathbf{S}' \bmod q_i$. For example, if $d = 2$ then $\tilde{F}_{1,d}$ is defined to be

$$\tilde{F}_{1,d}(\mathbf{S}'; \mathbf{R}_1) = \mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \mathbf{S}' \bmod q_2 + \mathbf{R}_1 \cdot q_2) \bmod q_1 \quad (14)$$

similarly to (2). Clearly, if $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$ are all zero matrices then $\tilde{F}_{1,d}(\mathbf{S}'; \mathbf{R}_1, \dots, \mathbf{R}_{d-1}) = F_{1,d}(\mathbf{S}')$.

We observe that this is enough for practical applications as long as $\mathbf{A}_1, \dots, \mathbf{A}_d$ are binding. Indeed, one can show, using similar methods to Section 1.2, that $F_{i,j}$ is binding based on the hardness of SIS for appropriate parameter choice q_1, \dots, q_d (see Section 3.4).

Formally, given matrices $\mathbf{A}_d, \dots, \mathbf{A}_1$ such that $\mathbf{A}_d \in \mathbb{Z}_q^{n \times m_d}$ and $\mathbf{A}_i \in \mathbb{Z}_q^{n \times m_i}$ for $i \in [d-1]$, the relation we give a zero-knowledge proof of knowledge for the relation

$$R = \left\{ \begin{array}{l} (pp, u, w) = ((\mathbf{q}, \mathbf{m}, n, B, B_R, \mathbf{A}_1, \dots, \mathbf{A}_d), \mathbf{T}, (\mathbf{S}', \bar{\mathbf{R}})) \mid \\ \left[\|\mathbf{s}_i\| \leq 2^d B \right]_{i \in [m_{d+1}]} \wedge \left[\|\mathbf{R}_i\|_\infty \leq B_R \right]_{i \in [d-1]} \wedge \left[\mathbf{R}_i \in \mathbb{Z}^{n M_i \times m_{d+1}} \right]_{i \in [d-1]} \\ \wedge (\mathbf{S}', \mathbf{T}) \in \mathbb{Z}^{M_d \times m_{d+1}} \times \mathbb{Z}_{q_1}^{n \times m_{d+1}} \wedge \tilde{F}_{1,d}(\mathbf{S}'; \bar{\mathbf{R}}) \equiv \mathbf{T} \pmod{q_1} \end{array} \right\}$$

where we denote $\bar{\mathbf{R}} := (\mathbf{R}_1, \dots, \mathbf{R}_{d-1})$, $\mathbf{S}' := [\mathbf{s}_1 \dots \mathbf{s}_{m_{d+1}}]$, $\mathbf{q} := (q_1, \dots, q_d)$ and $\mathbf{m} := (m_0, \dots, m_{d+1})$.

3.2 The Main Protocol

We present our zero-knowledge proof of knowledge in Fig. 4. First, we describe supporting algorithms that we will use in the protocol. Firstly, BT_i takes a matrix \mathbf{Z} which has a number of rows divisible by m_i and outputs its *block transpose*:

$$\text{BT}_i(\mathbf{Z}) := [\mathbf{Z}_1 \cdots \mathbf{Z}_{m_i}], \text{ where } \mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 \\ \vdots \\ \mathbf{Z}_{m_i} \end{bmatrix}.$$

On the other hand, Fold_i is a recursive algorithm which takes as input M_i matrices $\mathbf{U}_1, \dots, \mathbf{U}_{M_i}$ and $i+1$ challenge matrices $\mathbf{C}_1, \dots, \mathbf{C}_{i+1}$. If $i = 0$ then it simply outputs $\mathbf{U}_1 \mathbf{C}_1$. Otherwise, it splits the vector $(\mathbf{U}_1, \dots, \mathbf{U}_{M_i})$ into m_i shorter ones, i.e. $\bar{\mathbf{U}}_j = (\mathbf{U}_{(j-1)M_i+1}, \dots, \mathbf{U}_{jM_i})$ and runs $\mathbf{U}'_j = \text{Fold}_{i-1}(\bar{\mathbf{U}}_j; \mathbf{C}_1, \dots, \mathbf{C}_i)$ for each $j \in [m_i]$. Eventually, it outputs $[\mathbf{U}'_1 \cdots \mathbf{U}'_{m_i}] \mathbf{C}_{i+1}$. We show more properties of this algorithm in the correctness section.

The statement is $F_{1,d}(\mathbf{S}) \equiv \mathbf{T} \pmod{q_1}$. The protocol begins with the prover \mathcal{P} selecting a masking value for \mathbf{Y} with small coefficients and sending $\mathbf{W} = \mathbf{A}_d \mathbf{Y} \bmod q_d$. In the i -th round, the verifier \mathcal{V} picks a random challenge \mathbf{C}_i and sends it to \mathcal{P} . The prover applies Fold to the intermediate commitments

$$\mathbf{V}_i = (\mathbf{S}_{[1, \dots, 1]}, \mathbf{S}_{[2, 1, \dots, 1]}, \dots, \mathbf{S}_{[m_1, 1, \dots, 1]}, \mathbf{S}_{[1, 2, \dots, 1]}, \dots, \mathbf{S}_{[m_1, \dots, m_{i-1}]})$$

as well as all the previous challenges $\mathbf{C}_1, \dots, \mathbf{C}_i$ sent by \mathcal{V} . If $i = d$ then \mathcal{P} also adds \mathbf{Y} and runs rejection sampling. Next, it returns

$$\mathbf{Z}_i = \text{Fold}_{i-1}(\mathbf{V}_i; \mathbf{C}_1, \dots, \mathbf{C}_i) + \llbracket i = d \rrbracket \mathbf{Y}.$$

Finally, the verifier checks that all the \mathbf{Z}_i are small and for all $i \in [d-1]$:

$$\mathbf{A}_{i+1} \mathbf{Z}_{i+1} \equiv \text{BT}_i(\mathbf{Z}_i) \mathbf{C}_{i+1} + \llbracket i = d-1 \rrbracket \mathbf{W} \pmod{q_{i+1}}.$$

We assume that $(\mathbf{I}_{M_{d-1}} \otimes \mathbf{A}_d) \mathbf{S} \pmod{q_i}$ is public, although this information is not used by the verifier. Consequently, for all $0 \leq k < d-1$ and any $i_1, \dots, i_k \in [m_1] \times \dots \times [m_k]$, $\mathbf{S}_{[i_1, \dots, i_k]}$ is known as well.

<p>$\mathcal{P}_0(pp, \mathbf{S})$</p> <p>01 $\mathbf{Y} \leftarrow D_{\sigma}^{m_d \times \lambda}$</p> <p>02 $\mathbf{W} := \mathbf{A}_d \mathbf{Y} \pmod{q_d}$</p> <p>03 return $(\mathbf{W}, St = \mathbf{Y})$</p> <p>$\mathcal{P}_i(pp, (\mathbf{C}_1, \dots, \mathbf{C}_i), \mathbf{S}, St = \mathbf{Y})$ for $i \in [d]$</p> <p>04 if $i = 1$ then</p> <p>05 $\mathbf{V}_i := \mathbf{S}_{\square}$</p> <p>06 else $\mathbf{V}_i := (\mathbf{S}_{[1, \dots, 1]}, \mathbf{S}_{[2, 1, \dots, 1]}, \dots, \mathbf{S}_{[m_1, \dots, m_{i-1}]})$</p> <p>07 $\mathbf{V}'_i := \text{Fold}_{i-1}(\mathbf{V}_i; \mathbf{C}_1, \dots, \mathbf{C}_i)$</p> <p>08 if $i = d$ then</p> <p>09 $\mathbf{Z}_d := \mathbf{V}'_d + \mathbf{Y}$</p> <p>10 Abort if $\text{Rej}(\mathbf{Z}_d, \mathbf{V}'_d, \sigma, \rho) = 1$</p> <p>11 else $\mathbf{Z}_i := \mathbf{V}'_i$</p> <p>12 return $(\mathbf{Z}_i, St = \mathbf{Y})$</p>	<p>$\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1})$</p> <p>13 if $i = 0$</p> <p>14 then return $\mathbf{U}_1 \mathbf{C}_1$</p> <p>15 $\mathbf{U}'_j := \text{Fold}_{i-1}(\mathbf{U}_{(j-1)M_{i-1}+1}, \dots, \mathbf{U}_{jM_{i-1}}; \mathbf{C}_1, \dots, \mathbf{C}_i)$</p> <p>16 return $[\mathbf{U}'_1 \dots \mathbf{U}'_{m_i}] \mathbf{C}_{i+1}$</p> <p>$\text{BT}_i(\mathbf{Z})$</p> <p>17 Write $\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 \\ \vdots \\ \mathbf{Z}_{m_i} \end{bmatrix}$</p> <p>18 return $[\mathbf{Z}_1 \dots \mathbf{Z}_{m_i}]$</p> <p>$\mathcal{V}(pp, \mathbf{T}, (\mathbf{C}_i, \mathbf{Z}_i)$ for $i \in [d]$</p> <p>19 $[\mathbf{z}_1, \dots, \mathbf{z}_{m_{d+1}}] := \mathbf{Z}_d$</p> <p>20 Check:</p> <p>21 1. $\forall j \in [m_{d+1}], \ \mathbf{z}_j\ \leq B$</p> <p>22 2. $\forall i \in [d-1], \ \mathbf{Z}_i\ _{\infty} \leq M_{i-1} m_{d+1} \lambda^{i-1} \frac{q_{i+1}-1}{2}$</p> <p>23 3. $\mathbf{A}_1 \mathbf{Z}_1 \equiv \mathbf{T} \mathbf{C}_1 \pmod{q_1}$</p> <p>24 4. $\forall i \in [d-2], \mathbf{A}_{i+1} \mathbf{Z}_{i+1} \equiv \text{BT}_i(\mathbf{Z}_i) \mathbf{C}_{i+1} \pmod{q_{i+1}}$</p> <p>25 5. $\mathbf{A}_d \mathbf{Z}_d \equiv \text{BT}_{d-1}(\mathbf{Z}_{d-1}) \mathbf{C}_d + \mathbf{W} \pmod{q_d}$</p>
---	--

Fig. 4. Levelled Lattice Commitment Protocol.

3.3 Security Analysis

We start by proving certain properties of the Fold algorithm defined in Fig. 4. They will be crucial when proving correctness of our protocol.

Lemma 3.1. *Let $i \in [d]$ and $k, \ell, m \in \mathbb{N}$. Take arbitrary $\mathbf{U}_1, \dots, \mathbf{U}_{M_i} \in \mathbb{Z}^{k \times m_{d+1}}$ and $\mathbf{C}_1, \dots, \mathbf{C}_{i+1}$ such that $\mathbf{C}_1 \in \{0, 1\}^{m_{d+1} \times \lambda}$ and for $j > 1$, $\mathbf{C}_j \in \{0, 1\}^{m_{j-1} \lambda \times \lambda}$. Then, the following hold.*

- (i) *There exist matrices $\mathbf{D}_1, \dots, \mathbf{D}_{M_i} \in \mathbb{Z}^{m_{d+1} \times \lambda}$ such that $\|\mathbf{D}_i\|_{\infty} \leq \lambda^i$ and $\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \sum_{t=1}^{M_i} \mathbf{U}_t \mathbf{D}_t$.*
- (ii) *For all $\mathbf{A} \in \mathbb{Z}^{m \times k}$,*

$$\mathbf{A} \cdot \text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \text{Fold}_i(\mathbf{A} \mathbf{U}_1, \dots, \mathbf{A} \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}).$$

(iii) Suppose that each \mathbf{U}_j can be written as $\mathbf{U}_j = \begin{bmatrix} \mathbf{U}_{j,1} \\ \vdots \\ \mathbf{U}_{j,\ell} \end{bmatrix}$, where all matrices $\mathbf{U}_{j,j'}$ have the same dimensions. Then:

$$\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \begin{bmatrix} \text{Fold}_i(\mathbf{U}_{1,1}, \dots, \mathbf{U}_{M_i,1}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) \\ \vdots \\ \text{Fold}_i(\mathbf{U}_{1,\ell}, \dots, \mathbf{U}_{M_i,\ell}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) \end{bmatrix}.$$

Proof. Each part of Lemma 3.1 is proved by induction on i . A detailed proof can be found in the full version of this paper.

We are now ready to prove security properties of our protocol.

Theorem 3.2. *Let $s \geq \max_{i_1, \dots, i_{d-1}} s_1(\mathbf{S}_{[i_1, \dots, i_{d-1}]})$, $\rho > 1$ be a constant, $\sigma \in \mathbb{R}$ be such that $\sigma \geq \frac{12}{\ln \rho} M_{d-1} s \lambda^{d-1} \sqrt{m_{d+1} \lambda}$, and $B = \sqrt{2m_d} \sigma$. Then the protocol described in Fig. 4 is a zero-knowledge proof of knowledge for R .*

Proof. We prove correctness and zero-knowledge, and prove knowledge soundness separately in Theorem 3.3.

Correctness. If \mathcal{P} and \mathcal{V} are honest then the probability of abort is exponentially close to $1 - 1/\rho$ (see Lemma 2.7). Indeed, note that by Lemma 3.1 (i) and the triangle inequality we know that $\|V'_d\|$ is bounded above by $M_{d-1} s \lambda^{d-1} \sqrt{m_{d+1} \lambda}$. In a similar manner, one can show that the second verification condition is satisfied. Now, we show that the equations verified by \mathcal{V} are true.

Firstly, note that $\mathbf{A}_1 \mathbf{Z}_1 = \mathbf{A}_1 \text{Fold}(\mathbf{S}_{\square}; \mathbf{C}_1) = \mathbf{A}_1 \mathbf{S}_{\square} \mathbf{C}_1 \equiv \mathbf{TC}_1 \pmod{q_1}$. Now, fix $i \in [d-1]$. We know that $\mathbf{Z}_i = \text{Fold}_{i-1}(\mathbf{V}_i; \mathbf{C}_1, \dots, \mathbf{C}_i)$ (line 7) where

$$\mathbf{V}_i = (\mathbf{S}_{[1, \dots, 1]}, \mathbf{S}_{[2, 1, \dots, 1]}, \dots, \mathbf{S}_{[m_1, 1, \dots, 1]}, \mathbf{S}_{[1, 2, \dots, 1]}, \dots, \mathbf{S}_{[m_1, \dots, m_{i-1}]}).$$

By definition, each $\mathbf{S}_{[j_1, \dots, j_{i-1}]}$ is equal to

$$\begin{bmatrix} \mathbf{A}_{i+1} \mathbf{S}_{[j_1, \dots, j_{i-1}, 1]} \\ \vdots \\ \mathbf{A}_{i+1} \mathbf{S}_{[j_1, \dots, j_{i-1}, m_i]} \end{bmatrix}.$$

By Lemma 3.1 (ii) and (iii), we have

$$\mathbf{Z}_i = \begin{bmatrix} \mathbf{A}_{i+1} \text{Fold}_{i-1}(\mathbf{V}_{i,1}; \mathbf{C}_1, \dots, \mathbf{C}_i) \\ \vdots \\ \mathbf{A}_{i+1} \text{Fold}_{i-1}(\mathbf{V}_{i,m_i}; \mathbf{C}_1, \dots, \mathbf{C}_i) \end{bmatrix},$$

where

$$\mathbf{V}_{i,j} = (\mathbf{S}_{[1, \dots, 1, j]}, \mathbf{S}_{[2, 1, \dots, 1, j]}, \dots, \mathbf{S}_{[m_1, 1, \dots, 1, j]}, \mathbf{S}_{[1, 2, \dots, 1, j]}, \dots, \mathbf{S}_{[m_1, \dots, m_{i-1}, j]}).$$

Observe that \mathbf{V}_{i+1} is indeed equal to the concatenation of vectors $\mathbf{V}_{i,1}, \dots, \mathbf{V}_{i,m_i}$. Then, by applying the BT function to \mathbf{Z}_i and by definition of Fold, we obtain:

$$\begin{aligned}
\text{BT}_i(\mathbf{Z}_i)\mathbf{C}_{i+1} &= [\mathbf{A}_{i+1}\bar{\mathbf{V}}_1 \cdots \mathbf{A}_{i+1}\bar{\mathbf{V}}_{m_i}] \mathbf{C}_{i+1} \\
&= \mathbf{A}_{i+1} [\bar{\mathbf{V}}_1 \cdots \bar{\mathbf{V}}_{m_i}] \mathbf{C}_{i+1} \\
&= \mathbf{A}_{i+1} \text{Fold}_i(\mathbf{V}_{i+1}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) \\
&= \mathbf{A}_{i+1} \mathbf{Z}_{i+1},
\end{aligned} \tag{15}$$

where $\bar{\mathbf{V}}_j := \text{Fold}_{i-1}(\mathbf{V}_{i,j}; \mathbf{C}_1, \dots, \mathbf{C}_i)$. The last verification equation is also satisfied using the same argument as before and noting that $\mathbf{A}_d \mathbf{Y} = \mathbf{W}$.

Eventually, since each coefficient of \mathbf{Z} is statistically close to D_σ , then according to (6) we have $\|\mathbf{z}_i\| \leq \sqrt{2m_d}\sigma$ with overwhelming probability.

Honest-Verifier Zero-Knowledge. We will now prove that our protocol is honest-verifier zero-knowledge. More concretely, we show that it is zero-knowledge when the prover does not abort prior to sending \mathbf{Z}_d . We recall that for all $0 \leq k < d-1$, $\mathbf{S}_{[i_1, \dots, i_k]}$ is known to adversaries.

Define a simulator \mathcal{S} as follows. It first selects $\mathbf{C}_1 \xleftarrow{\$} \{0, 1\}^{m_{d+1} \times \lambda}$ and $\mathbf{C}_j \xleftarrow{\$} \{0, 1\}^{m_{j-1} \lambda \times \lambda}$ for $j = 2, \dots, d$. Next, \mathcal{S} samples $\mathbf{Z}_d \leftarrow D_\sigma^{M_{d-1} \times \lambda}$. Then, for $i \in [d-1]$, the simulator sets $\mathbf{Z}_i := \text{Fold}_{i-1}(\mathbf{V}_i; \mathbf{C}_1, \dots, \mathbf{C}_i)$ where

$$\mathbf{V}_i = (\mathbf{S}_{[1, \dots, 1]}, \mathbf{S}_{[2, 1, \dots, 1]}, \dots, \mathbf{S}_{[m_1, 1, \dots, 1]}, \mathbf{S}_{[1, 2, \dots, 1]}, \dots, \mathbf{S}_{[m_1, \dots, m_{i-1}]}).$$

Finally, \mathcal{S} sets $\mathbf{W} := \mathbf{A}_d \mathbf{Z}_d - \text{BT}_{d-1}(\mathbf{Z}_{d-1})\mathbf{C}_d$ and outputs $(\mathbf{W}, \mathbf{C}_1, \mathbf{Z}_1, \dots, \mathbf{C}_d, \mathbf{Z}_d)$.

It is clear that \mathcal{V} verifies with overwhelming probability. We already argued in the section on correctness that in the real protocol when no abort occurs the distribution of \mathbf{Z}_d is within statistical distance 2^{-100} of $D_\sigma^{M_{d-1} \times \lambda}$. Since \mathbf{W} is completely determined by $\mathbf{A}_d, \mathbf{Z}_{d-1}, \mathbf{Z}_d, \mathbf{C}_d$ and additionally, the distribution of \mathbf{Z}_i output by \mathcal{S} is identical to the one in the real protocol for $i \in [d-1]$, the distribution of $(\mathbf{W}, \mathbf{C}_1, \mathbf{Z}_1, \dots, \mathbf{C}_d, \mathbf{Z}_d)$ output by \mathcal{S} is within 2^{-100} of the distribution of these variables in the actual non-aborting run of the protocol. \square

Knowledge Soundness. We describe a knowledge extractor \mathcal{E} which finds small matrices \mathbf{S}' and $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$ such that $\mathbf{T} = \tilde{F}_{1,d}(\mathbf{S}'; \mathbf{R}_1, \dots, \mathbf{R}_{d-1})$.

Theorem 3.3. *For any prover \mathcal{P}^* who succeeds with probability $\varepsilon > 2^{-\lambda+1} \cdot (4dN)^{2d}$ over its random tape $\chi \in \{0, 1\}^x$ and the challenge choice $\mathbf{C}_1, \dots, \mathbf{C}_d$, such that $\mathbf{C}_1 \xleftarrow{\$} \{0, 1\}^{m_{d+1} \times \lambda}$ and $\mathbf{C}_j \xleftarrow{\$} \{0, 1\}^{m_{j-1} \lambda \times \lambda}$ for $j > 1$, there exists a knowledge extractor \mathcal{E} running in expected time $\text{poly}(\lambda)/\varepsilon$ who can extract \mathbf{S}' and $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$ such that $\tilde{F}_{1,d}(\mathbf{S}'; \mathbf{R}_1, \dots, \mathbf{R}_{d-1}) = \mathbf{T}$. Moreover, each column of \mathbf{S}' has norm at most $2^d B$ and $\forall k \in [d-1]$, we have $\|\mathbf{R}_k\|_\infty \leq 2^k (M_{k-1} m_{d+1} \lambda^{k-1} + 2)$.*

Proof. We provide a sketch of the proof here and include more detail in the full version of this paper. First, the extractor \mathcal{E} constructs a tree \mathcal{T} of partial transcripts similar to [14,12] where each vertex of \mathcal{T} (apart from the root) is created using extraction techniques from [5] based on the heavy-rows argument. The

tree-construction algorithm `TreeConstruct` is given in Fig. 7 in Section A. Next, \mathcal{E} computes relaxed openings of the levelled commitments, using the algorithm in Fig. 8 in Section A.

We sketch some of the steps of the extraction algorithm. First, we can fix $\alpha \in [m_{d+1}]$ and define an extractor \mathcal{E} which finds small vectors $\mathbf{s}', \mathbf{r}_1, \dots, \mathbf{r}_{i-1}$ such that $F_{1,d}(\mathbf{s}'; \mathbf{r}_1, \dots, \mathbf{r}_{d-1}) = \mathbf{t}_\alpha$, where \mathbf{t}_α is the α -th column vector of \mathbf{T} ¹¹. Then, using the extraction strategy from [5], we can find $\mathbf{Z}'_1, \mathbf{Z}''_1$ such that $\mathbf{A}_1(\mathbf{z}'_{1,u} - \mathbf{z}''_{1,u}) \equiv \mathbf{t}_\alpha \pmod{q_1}$ for some u , where $\mathbf{z}'_{1,u}$ (resp. $\mathbf{z}''_{1,u}$) is the u -th column of \mathbf{Z}'_1 (resp. \mathbf{Z}''_1). Hence, \mathcal{E} must find a preimage of $\mathbf{z}'_{1,u}$ and $\mathbf{z}''_{1,u}$. We focus on the former. By symmetry, the latter can be obtained analogously.

Suppose that we continue running the prover \mathcal{P}^* given the first response \mathbf{Z}'_1 . We want to get a preimage of the u -th column of \mathbf{Z}'_1 . Note that when applying BT_1 to \mathbf{Z}'_1 , the u -th column vector gets split into the u -th, $u + \lambda$ -th, ..., $u + (m_1 - 1)\lambda$ -th columns. Take arbitrary $j \in \{u + i\lambda : 0 \leq i < m_1\}$. Then, again by rewinding \mathcal{P}^* , we can get $\mathbf{Z}'_2, \mathbf{Z}''_2$ such that

$$\mathbf{A}_2 \hat{\mathbf{z}}_{2,j} = \mathbf{A}_2(\mathbf{z}'_{2,v} - \mathbf{z}''_{2,v}) \equiv \text{BT}_1(\mathbf{Z}'_1)_j \pmod{q_2}$$

for some v , where $\hat{\mathbf{z}}_{2,j} := \mathbf{z}'_{2,v} - \mathbf{z}''_{2,v}$ and $\text{BT}_1(\mathbf{Z}'_1)_j$ denotes the j -th column of $\text{BT}_1(\mathbf{Z}'_1)$. By repeating this argument for all possible j , we obtain:

$$(\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \begin{bmatrix} \hat{\mathbf{z}}_{2,u} \\ \vdots \\ \hat{\mathbf{z}}_{2,u+(m_1-1)\lambda} \end{bmatrix} = \begin{bmatrix} \text{BT}_1(\mathbf{Z}'_1)_u \\ \vdots \\ \text{BT}_1(\mathbf{Z}'_1)_{u+(m_1-1)\lambda} \end{bmatrix} = \mathbf{z}'_{1,u} \pmod{q_2}.$$

Observe how the tree structure appears in the argument. We first find \mathbf{Z}'_1 and \mathbf{Z}''_1 which correspond to the two children of the root. Then, for each such vertex V , we repeat the same argument m_1 times and add new children $W_1, W'_1, \dots, W_{m_1}, W'_{m_1}$ of V . In general, the tree \mathcal{T} has exactly $2^i M_{i-1}$ vertices on each level $i > 0$.

Eventually, the extracted solution consists of responses which correspond to the leaves of \mathbf{T} . We also get additional terms \mathbf{R}_i since each verification equation holds for different moduli. Hence, in order to make any implications from them, we need to first “lift” the previous verification equation and then we can apply it to the next one. The \mathbf{R}_i terms are the result of such lifting. \square

3.4 Asymptotic Parameter Choice

In this section, we set parameters for our protocol which minimise the total communication size (see Table 5). More concretely, we pick q_1, \dots, q_d and m_1, \dots, m_{d+1} (conditioned on the fact that $N = \prod_{i=1}^{d+1} m_i$ is fixed and $N = \mathcal{O}(\lambda^r)$ for some constant, integer r). For readability, we consider asymptotic parameter choice, neglecting constant terms and focussing on the leading terms using “big- \mathcal{O} ” notation.

¹¹ By collecting extracted solutions for all α , we can merge them and thus obtain the overall solution.

Parameter	Size	Description
λ		Security parameter
p	$\text{poly}(\lambda)$	The largest value of the secrets, i.e. $\ \mathbf{S}\ _\infty$
s		Operator norm of \mathbf{S}
N	$m_1 \cdot \dots \cdot m_{d+1} = \text{poly}(\lambda)$	Number of secrets
n	$d \cdot \mathcal{O}(\log N)$	Number of rows in $\mathbf{A}_1, \dots, \mathbf{A}_d$
q_i	$\mathcal{O}(N^{d-i+2}(2\lambda)^d p^2)$	Modulus corresponding to the commitment \mathbf{A}_i
m_i	$\left(\mathcal{O}\left(\frac{d \cdot \log N}{d^2 \lambda \cdot \log^2 N + \lambda^2}\right) \cdot N\right)^{\frac{1}{d+1}}$	i -th dimension of \mathbf{S} for $i \in [d-1]$
m_d	$\mathcal{O}\left(\frac{d^2 \cdot \log^2 N + \lambda}{\lambda d \cdot \log N}\right)^{\frac{d}{d+1}} \cdot \left(\frac{N}{\lambda}\right)^{\frac{1}{d+1}}$	d -th dimension of \mathbf{S}
m_{d+1}	$\left(\mathcal{O}\left(\frac{\lambda^d d \cdot \log N}{d^2 \cdot \log^2 N + \lambda}\right) \cdot N\right)^{\frac{1}{d+1}}$	$(d+1)$ -th dimension of \mathbf{S}
σ	$\frac{12}{\ln \rho} M_{d-1} s \lambda^{d-1} \sqrt{m_{d+1} \lambda}$	Standard deviation for rejection sampling
B	$\sqrt{2m_d} \cdot \sigma$	Soundness slack from proof of knowledge
B_R	$(2\lambda)^d N \cdot \sigma$	Infinity norm of extracted matrices $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$

Fig. 5. Parameter choice for our protocol.

To begin with, we compute simple upper bounds for the norms of the prover's responses. First, let us assume that secret elements in \mathbf{S} have size at most $p < N$, i.e. $\|\mathbf{S}\|_\infty \leq p$. Using the Cauchy-Schwarz inequality and the definition of an operator norm, we get a bound $s \leq Np^2$. Now, we provide a simple bound on B which is defined in Theorem 3.2:

$$B = \sqrt{2m_d} \sigma = \sqrt{2m_d} \cdot \frac{12}{\ln \rho} M_{d-1} s \lambda^{d-1} \sqrt{m_{d+1} \lambda} = \mathcal{O}(\lambda^d N^2 p^2).$$

We note this bound can be substantially improved. Concretely, $s \leq m_d m_{d+1} p^2$ since we only consider the operator norm of $m_d \times m_{d+1}$ matrices in \mathbb{Z}_p . By picking the parameters set below, we get $s = \mathcal{O}(\lambda^2 N^{2/d+1})$. However, for readability, we demonstrate a simpler bound.

We know from Theorem 3.3 that for $k \in [d-1]$ we have

$$\|\mathbf{R}_k\|_\infty \leq 2^k (M_{k-1} m_{d+1} \lambda^{k-1} + 2) \leq (2\lambda)^d N =: B_R.$$

We are ready to set q_d . In order to make \mathbf{A}_d binding and satisfy (7), one needs to pick $q_d > 2\|\mathbf{s}'_d\|$ where \mathbf{s}'_i is the i -th column of the extracted matrix \mathbf{S}' in Theorem 3.2. We know that $\|\mathbf{s}'_i\| \leq 2^d B$ and therefore choose $q_d = \mathcal{O}((2\lambda)^d N^2 p^2)$.

Next, let us fix $i \in [d-1]$ and consider the explicit formula for $F_{i,j}$ in (13) without tensor notation. One observes that each copy of the matrix \mathbf{A}_i is multiplied from the right-hand side by a matrix of the form $\mathbf{U} = (\mathbf{V} \bmod q_{i+1}) + q_{i+1} \mathbf{R}$ and we know that $\|\mathbf{R}\|_\infty \leq (2\lambda)^d N$. Thus, we just need to choose q_i which satisfies $q_i > 2N\|\mathbf{U}\|_\infty \geq N \cdot (q_{i+1} + 2 \cdot (2\lambda)^d N) = Nq_{i+1} + 2 \cdot (2\lambda)^d N^2$. We solve this recursive formula for q_i and obtain

$$\begin{aligned} q_i &= \mathcal{O}\left(N^{d-i} \left((2\lambda)^d N^2 p^2 + \frac{2 \cdot (2\lambda^4) N^2}{N-1} \right) - \frac{2 \cdot (2\lambda^4) N^2}{N-1}\right) \\ &= \mathcal{O}(N^{d+2-i} (2\lambda)^d p^2). \end{aligned} \quad (16)$$

Hence, we have $\log q_i \leq \log q_1 = d \cdot \mathcal{O}(\log N)$ for $i \in [d]$. Finally, in order to make all commitments $\mathbf{A}_1, \dots, \mathbf{A}_d$ satisfy (7), we pick $n = d \cdot \mathcal{O}(\log N)$.

Now, let us set m_1, \dots, m_{d+1} which minimise the total communication cost of our protocol, including the statement \mathbf{T} . First, note that the verifier \mathcal{V} sends $\lambda m_{d+1} + \lambda^2 \cdot (m_1 + \dots + m_{d-1})$ bits as challenges. Next, consider the communication cost from the prover's side. At the beginning, \mathcal{P} sends \mathbf{W} which has $n\lambda \log q_d = \mathcal{O}(d^2 \lambda \log^2 N)$ bits. Since it does not contain any m_1, \dots, m_{d+1} , we ignore this term for now. Next, we note that from the second verification equation, each \mathbf{Z}_i sent by \mathcal{P} satisfies:

$$\log(2\|\mathbf{Z}_i\|_\infty) \leq \log(M_{i-1} m_{d+1} \lambda^{i-1} (q_{i+1} - 1)) \leq \log(N \lambda^d q_{i+1}) = d \cdot \mathcal{O}(\log N)$$

for $i \in [d-1]$. On the other hand, with overwhelming probability we have $\|\mathbf{Z}_d\|_\infty \leq 6\sigma = \mathcal{O}(\lambda^d N^2 p^2)$ and thus

$$\log(2\|\mathbf{Z}_d\|_\infty) = \mathcal{O}(d \log \lambda + \log N) = d \cdot \mathcal{O}(\log N).$$

Therefore, \mathcal{P} sends in total (excluding \mathbf{W})

$$\begin{aligned} nm_{d+1} \log q_1 + \sum_{i=1}^{d-1} nm_i \lambda \log(2\|\mathbf{Z}_i\|) + m_d \lambda \log(2\|\mathbf{Z}_d\|_\infty) \\ \leq \left(nm_{d+1} + \sum_{i=1}^{d-1} nm_i \lambda + m_d \lambda \right) d \cdot \mathcal{O}(\log N) \end{aligned} \quad (17)$$

bits. Eventually, this can be upper-bounded by:

$$\sum_{i=1}^{d-1} (n\lambda d \cdot \mathcal{O}(\log N) + \lambda^2) \cdot m_i + \lambda d \cdot \mathcal{O}(\log N) \cdot m_d + (nd \cdot \mathcal{O}(\log N) + \lambda) \cdot m_{d+1}.$$

In order to minimise this expression, we want to set m_1, \dots, m_{d+1} in such a way that all these $d+1$ terms are (almost) equal. Fix m_{d+1} . Then,

$$m_d = \frac{nd \cdot \mathcal{O}(\log N) + \lambda}{\lambda d \cdot \mathcal{O}(\log N)} \cdot m_{d+1} \text{ and } m_i = \frac{m_{d+1}}{\lambda} \text{ for } i \in [d-1].$$

We compute an exact expression for m_{d+1} as follows:

$$N = \prod_{i=1}^{d+1} m_i = \frac{nd \cdot \mathcal{O}(\log N) + \lambda}{\lambda^d d \cdot \mathcal{O}(\log N)} (m_{d+1})^{d+1}$$

and hence we can set

$$m_{d+1} = \left(\frac{\lambda^d d \cdot \mathcal{O}(\log N)}{nd \cdot \mathcal{O}(\log N) + \lambda} \cdot N \right)^{\frac{1}{d+1}} < (\lambda^{d+1} N)^{\frac{1}{d+1}} = \lambda \cdot N^{\frac{1}{d+1}}.$$

Then, the total communication cost (now including \mathbf{W}) is bounded above by:

$$\begin{aligned} \mathcal{O}(d^2 \lambda \log^2 N) + (d+1)(nd \cdot \mathcal{O}(\log N) + \lambda) \cdot m_{d+1} \\ = \mathcal{O} \left(d^2 \lambda \log^2 N + (d+1)(d^2 \cdot \log^2 N + \lambda) \lambda N^{\frac{1}{d+1}} \right) \quad (18) \\ = \mathcal{O} \left(N^{\frac{1}{d+1}} \cdot (d^3 \lambda \log^2 N + d\lambda^2) \right). \end{aligned}$$

To obtain logarithmic proof size, set $d+1 = \log N$, giving communication cost $\lambda \cdot \mathcal{O}(\log^5 N + \lambda \log N)$.

4 Bulletproofs Folding Protocol

In the discrete logarithm setting, one can apply recursive arguments as in [14,12] and thus obtain logarithmic proof sizes. We show how these techniques can also be used in the lattice setting. Concretely, suppose the statement is as usual $\mathbf{A}\mathbf{s} = \mathbf{t}$ where $\mathbf{A} \in R^{1 \times k}$, $\mathbf{s} \in R^k$ with $\|\mathbf{s}\|_\infty \leq p$ and $R = \mathbb{Z}[X]/(X^n + 1)$. Then the number of secrets N is equal to kn . We highlight that the only variables which are defined the same in this section and the previous one are λ (security parameter), N (number of secrets) and p (the largest coefficient of the secrets).

We fold the initial statement as follows. Let us write $\mathbf{A} = [\mathbf{A}_1 \ \mathbf{A}_2]$ and $\mathbf{s} = \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \end{bmatrix}$ where $\mathbf{s}_1, \mathbf{s}_2 \in R^{k/2}$. Hence, if we define $\mathbf{l} = \mathbf{A}_1\mathbf{s}_2 \in R$ and $\mathbf{r} = \mathbf{A}_2\mathbf{s}_1 \in R$ then for all $c \in R$, $(c\mathbf{A}_1 + \mathbf{A}_2)(\mathbf{s}_1 + c\mathbf{s}_2) = c^2\mathbf{l} + c\mathbf{t} + \mathbf{r}$. This gives the following proof of knowledge of \mathbf{s} .

$$\begin{array}{ccc}
 \mathcal{P} & & \mathcal{V} \\
 \mathbf{l} = \mathbf{A}_1\mathbf{s}_2, \mathbf{r} = \mathbf{A}_2\mathbf{s}_1 & \xrightarrow{\mathbf{l}, \mathbf{r}} & \\
 & \xleftarrow{c} & c \stackrel{\$}{\leftarrow} \{X^i : i \in \mathbb{Z}_{2n}\} \subset R \\
 \mathbf{z} = \mathbf{s}_1 + c\mathbf{s}_2 & \xrightarrow{\mathbf{z}} & (c\mathbf{A}_1 + \mathbf{A}_2)\mathbf{z} \stackrel{?}{=} c^2\mathbf{l} + c\mathbf{t} + \mathbf{r} \\
 & & \|\mathbf{z}\|_\infty \stackrel{?}{\leq} 2p
 \end{array}$$

The vector \mathbf{z} has length $k/2$, so this protocol has half the communication cost of simply sending \mathbf{s} . We can repeat this protocol for the new statement $\mathbf{B}\mathbf{z} = \mathbf{t}'$ where $\mathbf{B} = c\mathbf{A}_1 + \mathbf{A}_2$ and $\mathbf{t}' = c^2\mathbf{l} + c\mathbf{t} + \mathbf{r}$.

Iterating the folding trick down to vectors of length 1 yields a protocol with communication cost $O(\log k)$. Extraction works in principle as follows. First, let us focus on extracting in the one-round protocol presented above. By rewinding we can get three equations

$$(c_i\mathbf{A}_1 + \mathbf{A}_2)\mathbf{z}_i = c_i^2\mathbf{l} + c_i\mathbf{t} + \mathbf{r}, \quad i = 1, 2, 3$$

for three different challenges c_i and answers \mathbf{z}_i . Combine these to obtain

$$\mathbf{A}_1 \left(\sum_{i=1}^3 \lambda_i c_i \mathbf{z}_i \right) + \mathbf{A}_2 \left(\sum_{i=1}^3 \lambda_i \mathbf{z}_i \right) = \sum_{i=1}^3 \lambda_i c_i^2 \mathbf{l} + \sum_{i=1}^3 \lambda_i c_i \mathbf{t} + \sum_{i=1}^3 \lambda_i \mathbf{r}. \quad (19)$$

If $\lambda = (\lambda_1, \lambda_2, \lambda_3)^T$ is a solution of the system

$$\begin{pmatrix} c_1^2 & c_2^2 & c_3^2 \\ c_1 & c_2 & c_3 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix},$$

then Equation (19) implies

$$\mathbf{A}_1 \left(\sum_{i=1}^3 \lambda_i c_i \mathbf{z}_i \right) + \mathbf{A}_2 \left(\sum_{i=1}^3 \lambda_i \mathbf{z}_i \right) = \mathbf{A} \sum_{i=1}^3 \lambda_i \begin{bmatrix} c_i \mathbf{z}_i \\ \mathbf{z}_i \end{bmatrix} = \mathbf{t}.$$

Hence, we get a preimage of \mathbf{t} but the problem is that in general it will not be short since λ_i can be large. In order to estimate the size of λ_i , we use the fact that for $i \neq j$, polynomials of the form $2/(X^i - X^j) \in R$ have coefficients in $\{-1, 0, 1\}$ ([9]). Also, we know by the properties of Vandermonde matrices that λ_i are of the form $\pm f \cdot (X^u - X^v)^{-1} \cdot (X^v - X^w)^{-1} \cdot (X^w - X^u)^{-1}$ for some pairwise distinct $u, v, w \in \mathbb{Z}_{2n}$ and $\|f\|_1 \leq 2$. Therefore, we have $\|8\lambda_i\|_\infty \leq 2n^2$. Hence, we have extracted a solution $\bar{\mathbf{z}}$ which satisfies $\mathbf{A}\bar{\mathbf{z}} = 8\mathbf{t}$ and

$$\|\bar{\mathbf{z}}\|_\infty = \left\| \sum_{i=1}^3 8\lambda_i \begin{bmatrix} c_i \mathbf{z}_i \\ \mathbf{z}_i \end{bmatrix} \right\|_\infty \leq \sum_{i=1}^3 \left\| 8\lambda_i \begin{bmatrix} c_i \mathbf{z}_i \\ \mathbf{z}_i \end{bmatrix} \right\|_\infty \leq \sum_{i=1}^3 2n^2 \cdot 2np = 12n^3p.$$

The extractor for the full protocol constructs a tree of partial transcripts similar to [14,12] and applies the strategy we described above at every level. Due to the small soundness error of order $1/n$, the protocol has to be repeated sufficiently many times to achieve negligible soundness error.

Proof size and slack. Let us consider the protocol with $d \leq \log k$ rounds. Then, using the same extraction strategy as above recursively, we obtain a relaxed opening $\bar{\mathbf{z}}$ to the modified equation: $\mathbf{A}\bar{\mathbf{z}} = 8^d \mathbf{t}$ such that $\|\bar{\mathbf{z}}\|_\infty = ((6n^3)^d \cdot 2^d \cdot p) = \mathcal{O}(n^{3d} \cdot 12^d \cdot p)$. Therefore, we set $q = \mathcal{O}(n^{3d} \cdot 12^d \cdot p)$. The proof size is then equal to $N \log(2^d p)/2^d + 2dn \log q$ which is $\mathcal{O}(N \log(2^d p)/2^d + d^2 n \log n)$.

Since this gives a soundness error of $O(1/n)$, we repeat the protocol $\lambda/\log n$ times in order to get soundness error $2^{-\lambda}$. This gives a total proof size of $\mathcal{O}\left(\frac{\lambda N \log(2^d p)}{2^d \log n} + \lambda d^2 n\right)$.

Suppose that we follow this protocol all the way down to vectors of length 1, i.e. $d = \log k$. Then, we have a “slack”¹² of $\|\bar{\mathbf{z}}\|_\infty = \mathcal{O}(n^{3 \log N} N^4 p)$ since $k = N/n < N$. The proof size is bounded by $\mathcal{O}(\lambda n \log N + \lambda n \log^2 N)$.

Comparison. We compare the Bulletproofs approach with levelled commitments introduced in Section 3 in terms of both proof sizes and slack. The latter one is not clearly defined in context of levelled commitments since one extracts some secret matrix \mathbf{S}' along with additional terms $\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_{d-1}$ (where d is a number of levels). Therefore, we only focus on the size of \mathbf{S}' and ignore the other terms. We provide a comparison of sizes for both techniques in Fig. 6. Firstly, we observe that none of these methods provide a way to extract an exact solution to the original equation. Indeed, with lattice commitments we only manage to extract \mathbf{S}' along with extra terms $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$ which satisfy (13). On the other hand, with Bulletproofs we extract a relaxed opening $\bar{\mathbf{z}}$ such

¹² Slack here means the Euclidean norm of an extracted solution.

	Bulletproofs	Levelled Commitments
Logarithmic proof size	$\mathcal{O}(\lambda n \log N + \lambda n \log^2 N)$	$\mathcal{O}(\lambda \log^5 N + \lambda^2 \log N)$
Corresponding slack	$\mathcal{O}(n^{3 \log N} N^4 \sqrt{N} p)$	$\mathcal{O}(\lambda^{\log N} N^3 p^2)$
$\text{poly}(\lambda, N^{1/c})$ proof size	$\mathcal{O}(\lambda N^{1/c} \log N + \lambda n \log^2 N)$	$\mathcal{O}(N^{1/c} \cdot (c^3 \lambda \log^2 N + c \lambda^2))$
Corresponding slack	$\mathcal{O}(n^{3(c-1) \log N/c} \cdot N^{4(c-1)/c} \sqrt{N} \cdot p)$	$\mathcal{O}((2\lambda)^c N^2 p^2)$

Fig. 6. Comparison of lattice Bulletproofs and levelled commitments.

that $\mathbf{A}\bar{\mathbf{z}} = 8^d \mathbf{t}$. In practice, this implies that the slack we have for $\bar{\mathbf{z}}$ gets also multiplied by the relaxation factor 8^d in front of \mathbf{t} . For $d = \log k$, this factor becomes $k^3 = N^3/n^3$.

From Fig. 6 we deduce that Bulletproofs folding offers smaller proof size at the cost of larger slack. Indeed, if one is not limited with any particular amount of slack then one can achieve quadratic-logarithmic proof size as shown on the top-left part of the table. Now, suppose that we can only tolerate $B = N^\alpha$ of slack for some α . The question would be which method achieves smaller proof size given this condition. Note that if $\alpha = 7.5$ then by the argument above, one would simply use Bulletproofs (by setting $n = 2$). Hence, assume that $3 \leq \alpha \leq 7$. For readability, from now on we do not write the “big- \mathcal{O} ” for each expression. Nevertheless, we still consider asymptotic parameters.

Let us first focus on levelled commitments – we find c such that $(2\lambda)^c N^2 p^2 = B$. Then

$$c = \frac{\log(B/N^2 p^2)}{\log(2\lambda)} \approx (\alpha - 2) \cdot \frac{\log N}{\log(2\lambda)} \approx (\alpha - 2)r$$

where $N = \lambda^r$ for some constant r ¹³. Then, the levelled commitments achieve $\tilde{\mathcal{O}}(N^{1/(\alpha-2)r})$ proof size. Now consider the Bulletproofs solution. To begin with, we would like to find d such that $n^{3d} \cdot 12^d \cdot \sqrt{N} p = B$. By solving this equation we have

$$d \approx \frac{\log(B/\sqrt{N})}{3 \log n + 4} = \frac{(\alpha - 1/2) \log N}{3 \log n + 4} = \gamma \log N$$

where $\gamma = (\alpha - 1/2)/(3 \log n + 4)$. Then, the Bulletproofs protocol has $\tilde{\mathcal{O}}(N^{1-\gamma})$ proof size. Therefore, we just need to compare $1 - \gamma$ with $1/(\alpha - 2)r$. The main observation is that for $r \geq 3$, the quadratic function $f_r(x) := (15 - 2x)(x - 2)r - 14$ is positive when $3 \leq x \leq 7$. Hence

$$\frac{1}{(\alpha - 2)r} < \frac{15 - 2\alpha}{14} \leq 1 - \frac{\alpha - 1/2}{3 \log n + 4} = 1 - \gamma.$$

This shows that if one is given only a limited (and relatively small) slack, one should consider using the levelled commitments approach to obtain small sub-linear proof sizes.

¹³ We neglect the $\log p$ term.

References

1. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 99–108, 1996.
2. Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017*, pages 2087–2104, 2017.
3. Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. *IACR Cryptol. ePrint Arch.*, 2020:517, 2020.
4. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, Dec 1993.
5. Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In *Proceedings of the 38th Annual International Cryptology Conference, CRYPTO'18*, pages 669–699, 2018.
6. Eli Ben-Sasson, Iddo Bentov, Alessandro Chiesa, Ariel Gabizon, Daniel Genkin, Matan Hamilis, Evgenya Pergament, Michael Riabzev, Mark Silberstein, Eran Tromer, and Madars Virza. Computational integrity with a public random string from quasi-linear pcps. In *Advances in Cryptology - EUROCRYPT 2017. Proceedings, Part III*, pages 551–579, 2017.
7. Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In *Advances in Cryptology - CRYPTO 2019. Proceedings, Part III*, pages 701–732, 2019.
8. Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. In *Advances in Cryptology - EUROCRYPT 2019. Proceedings, Part I*, pages 103–128, 2019.
9. 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. Proceedings, Part I*, pages 551–572, 2014.
10. Daniel J. Bernstein, Andreas Hülsing, Stefan Kölbl, Ruben Niederhagen, Joost Rijneveld, and Peter Schwabe. The sphincs⁺ signature framework. In *CCS*, pages 2129–2146. ACM, 2019.
11. Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. Quasi-optimal snargs via linear multi-prover interactive proofs. In *Proceedings of the 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, EUROCRYPT'18*, pages 222–255, 2018.
12. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *Advances in Cryptology - EUROCRYPT 2016, Proceedings*, pages 327–357, 2016.
13. Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In *Advances in Cryptology - CRYPTO 2019. Proceedings, Part I*, pages 176–202, 2019.
14. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *Proceedings of the 39th IEEE Symposium on Security and Privacy, S&P '18*, pages 315–334, 2018.

15. Jelle Don, Serge Fehr, and Christian Majenz. The measure-and-reprogram technique 2.0: Multi-round Fiat-Shamir and more. *CoRR*, abs/2003.05207, 2020.
16. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In *CRYPTO (2)*, volume 11693 of *Lecture Notes in Computer Science*, pages 356–383. Springer, 2019.
17. Léo Ducas, Eike Kiltz, Tancreède 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.
18. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In *ASIACRYPT (2)*, volume 8874 of *Lecture Notes in Computer Science*, pages 22–41. Springer, 2014.
19. Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. *IACR Cryptol. ePrint Arch.*, 2020:518, 2020.
20. Muhammed F. Esgin, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In *Advances in Cryptology - CRYPTO 2019. Proceedings, Part I*, pages 115–146, 2019.
21. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In *Advances in Cryptology - CRYPTO 2018. Proceedings, Part II*, pages 33–62, 2018.
22. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In *Advances in Cryptology - EUROCRYPT 2008. Proceedings*, pages 31–51, 2008.
23. Rosario Gennaro, Michele Minelli, Anca Nitulescu, and Michele Orrù. Lattice-based zk-snarks from square span programs. In *Proceedings of the 25th ACM Conference on Computer and Communications Security, CCS’18*, pages 556–573, 2018.
24. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
25. Jens Groth. Efficient zero-knowledge arguments from two-tiered homomorphic commitments. In *Advances in Cryptology - ASIACRYPT 2011. Proceedings*, pages 431–448, 2011.
26. Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology - EUROCRYPT 2016. Proceedings, Part II*, pages 305–326, 2016.
27. Jens Groth, Markulf Kohlweiss, Mary Maller, Sarah Meiklejohn, and Ian Miers. Updatable and universal common reference strings with applications to zk-snarks. In *Proceedings of the 38th Annual International Cryptology Conference, CRYPTO’18*, pages 698–728, 2018.
28. Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. In *EUROCRYPT (3)*, volume 10822 of *Lecture Notes in Computer Science*, pages 552–586. Springer, 2018.
29. Leslie Lamport. Constructing digital signatures from a one-way function, 1979.
30. Qipeng Liu and Mark Zhandry. Revisiting post-quantum fiat-shamir. In *CRYPTO (2)*, volume 11693 of *Lecture Notes in Computer Science*, pages 326–355. Springer, 2019.
31. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *Advances in Cryptology - ASIACRYPT 2009. Proceedings*, pages 598–616, 2009.

32. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *Advances in Cryptology - EUROCRYPT 2012. Proceedings*, pages 738–755, 2012.
33. Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, volume 4948 of *Lecture Notes in Computer Science*, pages 37–54. Springer, 2008.
34. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.
35. Ralph C. Merkle. A certified digital signature. In *CRYPTO*, volume 435 of *Lecture Notes in Computer Science*, pages 218–238. Springer, 1989.
36. Anca Nitulescu. Lattice-based zero-knowledge snargs for arithmetic circuits. In *Proceedings of the 6th International Conference on Cryptology and Information Security in Latin America, LATINCRYPT'19*, pages 217–236, 2019.
37. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2017. <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>.
38. Victor Shoup. Lower bounds for discrete logarithms and related problems. In *Advances in Cryptology - EUROCRYPT '97. Proceedings*, pages 256–266, 1997.

A Knowledge Soundness

In this section, we state the heavy-rows lemma and describe the extraction algorithms used in the proof of Theorem 3.3. A detailed analysis of the extraction algorithms is provided in the full version of this paper.

Lemma A.1. *Let $K > 1$ and $\mathbf{H} \in \{0, 1\}^{\ell \times n}$ for some $n, \ell > 1$, such that a fraction ε of the inputs of \mathbf{H} are 1. We say that a row of \mathbf{H} is “heavy” if it contains a fraction at least ε/K of ones. Then less than $1/K$ of the ones in \mathbf{H} are located in heavy rows.*

```

TreeConstructi([j1, ..., ji], Ci, Zi, t)
01  $V = \mathcal{T}[\mathbf{j}_1, \dots, \mathbf{j}_i]$ 
02  $\text{chal}(V) = \mathbf{C}_i, \text{resp}(V) = \mathbf{Z}_i, \text{index}(V) = t$ 
03  $(\text{root}, V_1, \dots, V_i) := \text{vertices on the path from root to } V \text{ where } V = V_i$ 
04 if  $i = d$  then
05   return
06  $\mathbf{z}_t := t\text{-th column vector of } \mathbf{Z}_i$ 
07 Write  $\mathbf{z}_t = \begin{bmatrix} \mathbf{z}_{t,0} \\ \vdots \\ \mathbf{z}_{t,m_i-1} \end{bmatrix}$ 
08  $t_j := t + (j-1)\lambda$  for  $j \in [m_i]$ 
09 for  $j \in [m_i]$  :
10   Select random  $\mathbf{C}'_{i+1}$  and then  $\mathbf{C}''_{i+1}$  such that  $\forall u \neq t_j, \mathbf{c}'_{i+1,u} = \mathbf{c}''_{i+1,u}$  and  $\mathbf{c}''_{i+1,t_j}$ 
is freshly sampled
11   Run  $\mathcal{P}^*$  on the  $i+1$ -th random challenge  $\mathbf{C}'_{i+1}$  until it outputs  $\mathbf{Z}'_{i+1}$ 
12   Rewind  $\mathcal{P}^*$  and re-run it on the  $i+1$ -th challenge  $\mathbf{C}''_{i+1}$  until it outputs  $\mathbf{Z}''_{i+1}$ 
13    $T' := (\mathbf{W}, \text{chal}(V_1), \text{resp}(V_1), \dots, \mathbf{C}'_{i+1}, \mathbf{Z}'_{i+1})$ 
14    $T'' := (\mathbf{W}, \text{chal}(V_1), \text{resp}(V_1), \dots, \mathbf{C}''_{i+1}, \mathbf{Z}''_{i+1})$ 
15   count = 0
16   while  $i = d-1$  and count  $< \lambda(4dN)^{2d}/\varepsilon$  and  $T'$  is not a valid transcript:
17     Rewind  $\mathcal{P}^*$  and run  $\mathcal{P}^*$  on the new  $\mathbf{C}'_d$  until it outputs  $\mathbf{Z}'_d$ 
18      $T' = (\mathbf{W}, \text{chal}(V_1), \text{resp}(V_1), \dots, \mathbf{C}'_d, \mathbf{Z}'_d)$ 
19     count = count + 1
20   if count  $\geq \lambda(4dN)^{2d}/\varepsilon$  then abort
21   count = 0
22   while  $i = d-1$  and count  $< 2\lambda(4dN)^{2d}/\varepsilon$  and  $T''$  is not a valid transcript:
23     Rewind  $\mathcal{P}^*$  and run  $\mathcal{P}^*$  on  $\mathbf{C}''_{i+1}$  such that  $\forall u \neq t_j, \mathbf{c}'_{i+1,u} = \mathbf{c}''_{i+1,u}$  and  $\mathbf{c}''_{i+1,t_j}$ 
is freshly sampled
24     Get response  $\mathbf{Z}''_d$ 
25      $T'' = (\mathbf{W}, \text{chal}(V_1), \text{resp}(V_1), \dots, \mathbf{C}''_{i+1}, \mathbf{Z}''_d)$ 
26     count = count + 1
27   if count  $\geq 2\lambda(4dN)^{2d}/\varepsilon$  then abort
28   Let  $\ell$  be an index where  $\mathbf{c}'_{i+1,u}[\ell] \neq \mathbf{c}''_{i+1,u}[\ell]$  (w.l.o.g.  $\mathbf{c}'_{i+1,u}[\ell] - \mathbf{c}''_{i+1,u}[\ell] = 1$ ,
otherwise swap)
29   TreeConstructi+1([j1, ..., ji, (j, 0)],  $\mathbf{C}'_{i+1}, \mathbf{Z}'_{i+1}, \ell$ )
30   TreeConstructi+1([j1, ..., ji, (j, 1)],  $\mathbf{C}''_{i+1}, \mathbf{Z}''_{i+1}, \ell$ )

```

Fig. 7. Construction of a tree \mathcal{T} of partial transcripts for \mathcal{P}^* . We denote $\mathbf{c}'_{i+1,j}$ (resp. $\mathbf{c}''_{i+1,j}$) to be the j -th row of \mathbf{C}'_{i+1} (resp. \mathbf{C}''_{i+1}).

```

Extracti(V)
01  $\mathcal{T}[\mathbf{j}_1, \dots, \mathbf{j}_i] = V$ 
02 if  $i = d$  then return  $\text{resp}_{\text{index}}(V)$ 
03 for  $j \in [m_i]$ :
04    $(\mathbf{s}'_{j,0}; \mathbf{r}_{j,0,i+2}, \dots, \mathbf{r}_{j,0,d-1}) \leftarrow \text{Extract}_{i+1}(\mathcal{T}[\mathbf{j}_1, \dots, \mathbf{j}_i, (j, 0)])$ 
05    $(\mathbf{s}'_{j,1}; \mathbf{r}_{j,1,i+2}, \dots, \mathbf{r}_{j,1,d-1}) \leftarrow \text{Extract}_{i+1}(\mathcal{T}[\mathbf{j}_1, \dots, \mathbf{j}_i, (j, 1)])$ 
06   if  $i < d - 1$  then
07      $\mathbf{w} = F_{i+2,d}(\mathbf{s}'_{j,b}; \mathbf{r}_{j,b,i+2}, \dots, \mathbf{r}_{j,b,d-1})$ 
08      $\mathbf{r}_{j,b,i+1} := (\text{resp}_{\text{index}}(\mathcal{T}[\mathbf{j}_1, \dots, \mathbf{j}_i, (j, b)]) - \mathbf{w}) / q_{i+2}$  for  $b \in \{0, 1\}$ 
09      $\mathbf{v}_b := (\mathbf{I}_{M_{i+1,d-1}} \otimes \mathbf{A}_d) \mathbf{s}'_{j,b} \bmod q_d$  for  $b \in \{0, 1\}$ 
10     for  $k = d - 1, d - 2, \dots, i + 1$ :
11        $\mathbf{u} := (\mathbf{v}_0 - \mathbf{v}_1 - (\mathbf{v}_0 - \mathbf{v}_1 \bmod q_k)) / q_k$ 
12        $\hat{\mathbf{r}}_{j,k} := \mathbf{r}_{j,0,k} - \mathbf{r}_{j,1,k} + \mathbf{u}$ 
13       if  $k > i + 1$  then  $\mathbf{v}_b = (\mathbf{I}_{M_{i+1,k-1}} \otimes \mathbf{A}_k)(\mathbf{v}_b + \mathbf{r}_{j,b,k}) \bmod q_k$  for  $b \in \{0, 1\}$ 
14 return  $\hat{\mathbf{s}} = \begin{bmatrix} \mathbf{s}'_{1,0} - \mathbf{s}'_{1,1} \\ \vdots \\ \mathbf{s}'_{m_i,0} - \mathbf{s}'_{m_i,1} \end{bmatrix}$  and  $\hat{\mathbf{r}}_k = \begin{bmatrix} \hat{\mathbf{r}}_{1,k} \\ \vdots \\ \hat{\mathbf{r}}_{m_i,k} \end{bmatrix}$  for  $k = i + 1, \dots, d - 1$ 

```

Fig. 8. Extracting relaxed openings of levelled commitments, or more concretely, preimages of $F_{i+1,d}$.