

Sub-Linear Lattice-Based Zero-Knowledge Arguments for Arithmetic Circuits^{*}

Carsten Baum¹, Jonathan Bootle², Andrea Cerulli²,
Rafael del Pino³, Jens Groth², and Vadim Lyubashevsky³

¹ Bar-Ilan University
carsten.baum@biu.ac.il

² University College London, London, UK
{jonathan.bootle.14, andrea.cerulli.13, j.groth}@ucl.ac.uk

³ IBM Research - Zurich

Abstract. We propose the first zero-knowledge argument with sub-linear communication complexity for arithmetic circuit satisfiability over a prime p whose security is based on the hardness of the short integer solution (SIS) problem. For a circuit with N gates, the communication complexity of our protocol is $O\left(\sqrt{N\lambda\log^3 N}\right)$, where λ is the security parameter. A key component of our construction is a surprisingly simple zero-knowledge proof for pre-images of linear relations whose amortized communication complexity depends only logarithmically on the number of relations being proved. This latter protocol is a substantial improvement, both theoretically and in practice, over the previous results in this line of research of Damgård et al. (CRYPTO 2012), Baum et al. (CRYPTO 2016), Cramer et al. (EUROCRYPT 2017) and del Pino and Lyubashevsky (CRYPTO 2017), and we believe it to be of independent interest.

Keywords: Sigma-protocol, zero-knowledge argument, arithmetic circuit, SIS assumption.

1 Introduction

Zero-knowledge proofs and arguments are used throughout cryptography as a key ingredient to ensure security in complex protocols. They form an important part of applications such as authentication protocols, electronic voting systems, encryption primitives, multi-party computation schemes, and verifiable computation protocols. Therefore, designing zero-knowledge protocols with strong security and high efficiency is of the utmost importance.

^{*} Jonathan Bootle, Andrea Cerulli and Jens Groth were supported by funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307937. Rafael del Pino and Vadim Lyubashevsky were supported in part by the SNSF ERC Transfer Starting Grant CRETP2-166734-FELICITY. Carsten Baum acknowledges support by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Ministers Office.

A zero-knowledge argument allows a prover to convince a verifier that a particular statement is true, without the prover revealing any other information that she knows about the statement. Statements are of the form $u \in \mathcal{L}$, where \mathcal{L} is a language in NP. We call w a witness for statement u if $(u, w) \in R$, where R is a polynomial time decidable binary relation associated with \mathcal{L} . Zero-knowledge arguments must be complete, sound and zero-knowledge.

Completeness: A prover with witness w for $u \in \mathcal{L}$ can convince the verifier.

Soundness: A prover cannot convince the verifier when $u \notin \mathcal{L}$.

Zero-knowledge: The interaction should not reveal anything to the verifier except that $u \in \mathcal{L}$. In particular, it should not reveal the prover’s witness w .

We wish to design a zero-knowledge argument based on the short integer solution (SIS) assumption. Lattice problems appear to resist quantum attacks, and possess attractive worst-case to average-case reductions, in stark contrast with number theoretic assumptions such as the hardness of factoring or computing discrete logarithms. Moreover, using SIS (and the even more efficient Ring-SIS) yields better computational efficiency, which is a significant bottleneck in many zero-knowledge arguments.

1.1 Our Contributions

We provide an honest verifier zero-knowledge argument for arithmetic circuit satisfiability over \mathbb{Z}_p , for an arbitrary prime p . Our argument is based on the SIS assumption [Ajt96, MR04], which is conjectured to be secure even against a quantum adversary. Our argument has an expected constant number of moves and sub-linear communication complexity, as shown in Table 1. Moreover, it achieves small soundness error in a single protocol execution. Moreover, both the prover and verifier have quasi-linear computational complexity in the amount of computation it would require to evaluate the arithmetic circuit directly. The argument therefore improves on the state-of-the-art in communication complexity for lattice proof systems and is efficient on all performance parameters.

Expected # Moves	Communication (bits)	Prover Complexity (bit ops)	Verifier Complexity (bit ops)
$O(1)$	$O(\sqrt{N\lambda} \log^3 N)$	$O(N \log N (\log \lambda)^2)$	$O(N (\log \lambda)^3)$

Table 1: Performance of our zero-knowledge argument for arithmetic circuit satisfiability. Here N is the number of gates in the arithmetic circuit, and λ is the security parameter.

Techniques. We draw inspiration from the discrete logarithm based arithmetic circuit satisfiability argument of Bootle et al. [BCC⁺16], which requires 5 moves and has square root communication complexity in the number of multiplication gates. In their argument the prover commits to all the wires using homomorphic

commitments, and embeds the wire values into a polynomial that verifies products and linear relations simultaneously, avoiding the cost for addition gates.

Almost all parts of the original arguments adapt seamlessly to the SIS setting, except for two important issues:

- To achieve sub-linear communication, we need a technique for proving knowledge of commitment openings in sub-linear space.
- Due to the new algebraic setting, we require new techniques for achieving negligible soundness in a single run of the protocol.

The first of these issues has been an open problem in a fairly active area of research, and we sketch our solution below.

Proof of Knowledge. Suppose that we have a linear relation

$$\mathbf{A}\mathbf{s} = \mathbf{t} \bmod q, \tag{1}$$

where $\mathbf{A} \in \mathbb{Z}_q^{r \times v}$, $\mathbf{t} \in \mathbb{Z}_q^r$ are public and $\mathbf{s} \in \mathbb{Z}_q^v$ is a vector with small coefficients, and we want to give a zero-knowledge proof of knowledge of an $\bar{\mathbf{s}}$ with small coefficients (the coefficients of $\bar{\mathbf{s}}$ may be larger than those of \mathbf{s}) that satisfies

$$\mathbf{A}\bar{\mathbf{s}} = \mathbf{t} \bmod q. \tag{2}$$

We do not currently know of any an efficient linear-communication protocol for proving knowledge of a single relation of the above form in a direct way. There are protocols, however, that allow for proofs of many such relations for the same \mathbf{A} but different \mathbf{s}_i (and thus different \mathbf{t}_i) in linear amortized complexity. We will mention these previous works in more detail in Section 1.2.

In this work, we give a protocol for proving (1) where the proof length is a factor $\frac{\lambda}{\ell} \cdot O(\log v\ell\lambda)$ larger than the total bit-length of ℓ pre-images $\mathbf{s}_1, \dots, \mathbf{s}_\ell$ of the relations, where λ is the security parameter. More specifically, to prove knowledge of ℓ pre-images $\mathbf{s}_1, \dots, \mathbf{s}_\ell$ whose coefficients have $\log s$ bits each, the prover needs to send λ vectors in \mathbb{Z}_q^v whose coefficients require $O(\log v\ell\lambda s)$ bits to represent. Ignoring logarithmic terms, our proof essentially requires a fixed-size proof regardless of the number of relations being proved. The previously best results had proofs that were at least linear in the total size of the pre-images.

Surprisingly, the proof of knowledge protocol turns out to just be a parallel repetition of λ copies of the ZKPoK implicit in the signing protocol from [Lyu12]. In particular, if we write the ℓ relations as $\mathbf{A}\mathbf{S} = \mathbf{T} \bmod q$, where $\mathbf{S} \in \mathbb{Z}_q^{v \times \ell}$, then the protocol begins with the prover selecting a “masking” value \mathbf{Y} with small coefficients and sending $\mathbf{W} = \mathbf{A}\mathbf{Y} \bmod q$. The verifier then picks a random challenge matrix $\mathbf{C} \in \{0, 1\}^{\ell \times (\lambda+2)}$, and sends it to the prover. The prover computes $\mathbf{Z} = \mathbf{C}\mathbf{S} + \mathbf{Y}$ and performs a rejection sampling step in order to make the distribution of \mathbf{Z} independent from \mathbf{S} , and if it passes, sends \mathbf{Z} to the verifier. The verifier checks that all columns comprising \mathbf{Z} have small norms and that $\mathbf{A}\mathbf{Z} = \mathbf{C}\mathbf{T} + \mathbf{W} \bmod q$. This protocol can be shown to be zero-knowledge using exactly the same techniques as in [Lyu09, Lyu12].

To show that the protocol is a proof of knowledge, we make the following observation: if the prover succeeds with probability $\epsilon > 2^{-\lambda}$, and she succeeded for a random \mathbf{C} , then there is a probability of $\epsilon - 2^{-\lambda-2}$ that she would successfully answer another challenge $\mathbf{C}' \neq \mathbf{C}$ that is constructed such that all rows except the i^{th} are the same as that of \mathbf{C} , and the i^{th} row is picked uniformly at random. This property follows from an averaging (or “heavy row”) type argument. The implication is that if the prover succeeds in time t with probability ϵ , then the extractor can extract responses to two such commitments \mathbf{C}, \mathbf{C}' in expected time $O(t/\epsilon)$. Obtaining two responses \mathbf{Z}, \mathbf{Z}' for two such challenges allows the extractor to compute $\mathbf{A}(\mathbf{Z} - \mathbf{Z}') = \mathbf{T}(\mathbf{C} - \mathbf{C}')$ where $\mathbf{C} - \mathbf{C}'$ is 0 everywhere except in row i . Since $\mathbf{C} \neq \mathbf{C}'$, this implies that some position in row i is ± 1 . If \mathbf{t}_i is the i^{th} column of \mathbf{T} and \mathbf{z}_i is the i^{th} column of $\mathbf{Z} - \mathbf{Z}'$, then we have a solution $\mathbf{A}\mathbf{z}_i = \pm\mathbf{t}_i$. Repeating this extraction ℓ times, each time rewinding by fixing all rows in the challenge except for the i^{th} , results in an algorithm that runs in expected time $O(\ell \cdot t/\epsilon)$, which is only a factor of ℓ larger than the expected running time of a successful prover.

In the case that we are proving (1) over the polynomial ring $\mathbb{Z}[X]/(X^d + 1)$, the proof can be even shorter, as we can reduce the number of columns in \mathbf{C} to $\approx \lambda/\log 2d$ because we can use challenges of the form $\pm X^i$ and prove the knowledge of $\bar{\mathbf{s}}$ such that $\mathbf{A}\bar{\mathbf{s}} = 2\mathbf{t}$ using the observation from [BCK⁺14].

Commitment Scheme. Central to the main proof of proving circuit satisfiability is being able to commit to N values in \mathbb{Z}_p and giving a ZKPoK for the values such that the total size of the commitments and the proofs is sub-linear in N . For this, it is necessary to use a compressing commitment scheme – i.e. one in which we can commit to n elements of \mathbb{Z}_p in space less than n elements. The scheme that we will use is the “classic” statistically-hiding commitment scheme based on the hardness of SIS that was already implicit in the original work of Ajtai [Ajt96]. The public randomness consists of two matrices $\mathbf{A} \in \mathbb{Z}_q^{r \times 2r \log_p q}$, $\mathbf{B} \in \mathbb{Z}_q^{r \times n}$, and committing to a message string $\mathbf{s} \in \mathbb{Z}_p^n$ where $p < q$ involves picking a random vector $\mathbf{r} \in \mathbb{Z}_p^{2r \log_p q}$ and outputting the commitment $\mathbf{t} = \mathbf{A}\mathbf{r} + \mathbf{B}\mathbf{s} \bmod q$. Thus the commitment of n elements of \mathbb{Z}_p requires $r \log q$ bits. One can set the parameters such that $n = \text{poly}(r)$ and the commitment scheme will still be computationally binding based on the worst-case hardness of approximating SIVP for all lattices of dimension r .

We now explain the intuition for putting together this commitment scheme with the zero-knowledge proof system we described above to produce a commitment to N values in \mathbb{Z}_p such that the total size of the commitments and the ZKPoK of the committed values is $O(\sqrt{N\lambda \log N})$. The idea is to create N/n commitments (for some choice of n which will be optimized later), with each one committing to n values. Our motivation is that an arithmetic circuit over \mathbb{Z}_p with N gates has $3N$ wire values in \mathbb{Z}_p . Now, we can arrange all of the wire values in the circuit into, for example, a $3N/n \times n$ matrix over \mathbb{Z}_p , and make one homomorphic commitment to all of the elements in each row of the matrix. Then, we can employ techniques from [Gro09a, BCC⁺16], where checking

arithmetic circuit satisfiability is reduced to checking linear-algebraic statements over committed matrices, using a homomorphic commitment scheme.

The total space requirement for these commitments is therefore $\frac{N}{n} \cdot r \log_p q$. We now have a linear equation of the form $[\mathbf{A} \ \mathbf{B}] \begin{bmatrix} \mathbf{R} \\ \mathbf{S} \end{bmatrix} = \mathbf{T} \bmod q$. Using our new zero-knowledge proof, the communication complexity of proving the knowledge of a short $\begin{bmatrix} \mathbf{R} \\ \mathbf{S} \end{bmatrix} \in \mathbb{Z}_q^{(r \log_p q + n) \times N/n}$ such that $[\mathbf{A} \ \mathbf{B}] \begin{bmatrix} \mathbf{R} \\ \mathbf{S} \end{bmatrix} = \mathbf{T} \bmod q$ requires sending λ vectors of length $r \log_p q + n$ with coefficients requiring $O(\log N \lambda p)$ bits, for a total bit-length of $n \cdot \lambda \cdot O(\log N \lambda p)$. Combining the proof size with the commitment size results in a total bit-size of

$$\frac{N}{n} \cdot r \log q + n \cdot \lambda \cdot O(\log N \lambda p).$$

We minimize the above by setting $n = \sqrt{\frac{Nr \log_p q}{\lambda \log N \lambda p}}$, which makes the size

$$O\left(\sqrt{Nr \lambda (\log_p q) (\log N \lambda p)}\right).$$

Based on the complexity of the best known algorithm against the SIS problem, one can set $\log q, r = O(\log N)$, thus making the proof size of order $O(\sqrt{N \lambda \log^3 N})$.

1.2 Related Work

Zero-knowledge proofs were invented by Goldwasser et al. [GMR85]. It is useful to distinguish between zero-knowledge *proofs*, with statistical soundness, and zero-knowledge *arguments* with computational soundness. The most efficient proofs have communication proportional to the size of the witness [IKOS07, KR08, GGI⁺15] and proofs cannot in general have communication that is smaller than the witness size unless surprising results about the complexity of solving SAT instances hold [GH98, GVW02]. Kilian [Kil92] showed that in contrast to proofs, zero-knowledge arguments can have very low communication complexity. His construction relied on the PCP theorem, and thus incurred a large computational cost.

Group theoretic zero-knowledge arguments. Schnorr [Sch91] and Guillou and Quisquater [GQ88] gave early examples of practical zero-knowledge arguments for concrete number theoretic problems. Extending Schnorr's protocols, there have been many constructions of zero-knowledge arguments based on the discrete logarithm assumption, for instance [CD97, Gro09b]. The most efficient discrete logarithm based zero-knowledge arguments for arithmetic circuits are by Bootle et al. [BCC⁺16] and later optimised in [BBB⁺17], which have logarithmic communication complexity and require a linear number of exponentiations.

An exciting line of research [Gro10a, Lip12, BCCT12, GGPR13, BCCT13, PHGR13, Gro16] on succinct non-interactive arguments (SNARGs) has yielded pairing-based constructions where the arguments consist of a constant number

of group elements. However, it can be shown that all SNARKs must rely on non-falsifiable knowledge extractor assumptions [GW11]. In contrast, since our argument is interactive, we do not need to rely on these strong assumptions.

Lattice-based zero-knowledge arguments. The first zero-knowledge proofs from lattice-based assumptions were aimed at lattice problems themselves. Goldreich and Goldwasser [GG98] presented constant round interactive zero knowledge proofs for the complements of the approximate Shortest Vector Problem (SVP) and the approximate Closest Vector Problem (CVP). Micciancio and Vadhan [MV03] later constructed statistical zero knowledge proofs for these problems which had efficient provers.

Stern’s protocol [Ste94] was one of the first zero-knowledge identification protocols to be based on a post-quantum assumption, namely, on the hardness of syndrome decoding for a random linear code, which is essentially proving (1) where $q = 2$ and $\|\mathbf{s}\| \ll \sqrt{v}$. The protocol achieves constant soundness error, and thus requires many parallel repetitions. Stern’s work prompted many variants and similar protocols. For example, [LNSW13] adapts the protocol for larger q , which implies proving knowledge of SIS solutions.

Another technique for creating zero-knowledge proofs is the “Fiat-Shamir with Aborts” approach [Lyu09, Gro10b, Lyu12]. When working over polynomial rings R , it gives a proof of knowledge of a vector $\bar{\mathbf{s}}$ with small coefficients (though larger than those in \mathbf{s}) and a ring element \bar{c} with very small coefficients satisfying $A\bar{\mathbf{s}} = \bar{c}\mathbf{t}$. As long as the ring R has many elements with small coefficients, such proofs are very efficient, producing soundness of $1 - 2^{-128}$ with just one iteration. While these proofs are good enough for constructing practical digital signatures (e.g. [GLP12, DDLL13, BG14]), commitment schemes with proofs of knowledge [BKLP15, BDOP16], and certain variants of verifiable encryption schemes [LN17], they prove less than what the honest prover knows. In many applications where zero-knowledge proofs are used, in particular those that need to take advantage of additive homomorphisms, the presence of the element \bar{c} makes these kinds of “approximate” proofs too weak to be useful. As of today, we do not have any truly practical zero-knowledge proof systems that give a proof of (1).

The situation is more promising when one considers *amortized* proofs. The work of [BD10] uses MPC-in-the-head to prove knowledge of plaintexts for multiple Regev [Reg05] ciphertexts. Damgård and López-Alt [DL12] extend the [BD10] results to prove knowledge of plaintext in \mathbb{Z}_p , rather than bits, and provide a proof for the correctness of multiplications. Combining these together gives a zero-knowledge proof for the satisfiability of arithmetic circuits with linear communication in the circuit size.

Another idea for proving the relation in (1) is to use the above-mentioned “Fiat-Shamir with Aborts” protocol, but with challenges that come from the set $\{0, 1\}$. The works of [BDLN16, CDXY17, dPL17] gave a series of improved protocols that were able to employ this technique in the amortized setting. Their proofs had a small polynomial “slack” (i.e. the ratio between the original committed \mathbf{s} and the extracted $\bar{\mathbf{s}}$) and were of approximate linear size when the

number of commitments was a couple of thousand. The schemes are considerably less efficient when one is proving fewer relations.

The amortized zero-knowledge proof in the current work improves on the above series of papers in two important ways. First, the number of relations necessary before the size of our proof is linear only in λ . But more importantly, if we have more than λ relations, the communication complexity does not increase except for small logarithmic factors (i.e. the proof size becomes sub-linear).

Hash-based zero-knowledge arguments. Recently Bootle et al [BCG⁺17] used error-correcting codes and linear-time collision-resistant hash functions to give proof systems for the satisfiability of an arithmetic circuit where the prover uses a linear number of field multiplications. Verification is even more efficient, requiring only a linear number of additions. While their proofs and arguments are asymptotically very efficient, they are not quite practical as their choices of error-correcting codes and hash functions involves very large constants.

Another effective way to construct efficient zero-knowledge proofs is to follow the so-called MPC-in-the-head paradigm of [IKOS07]. This approach proved itself to give very efficient constructions both theoretically and practically. Most notably, ZKBOO [GMO16] and subsequent optimisation ZKB++ [CDG⁺17] use hash functions to construct zero-knowledge arguments for the satisfiability of boolean circuits. Their communication complexity is linear in the circuit size, but the use of symmetric primitives gives good performances in practice. Ligerio [AHIV17] provides another implementation of the MPC-in-the-head paradigm and used techniques similar to [BCG⁺17] to construct sublinear arguments for arithmetic circuits.

2 Preliminaries

Algorithms in our schemes receive a security parameter λ as input (sometimes implicitly) written in unary. The intuition is that the higher the security parameter, the lower the risk of the scheme being broken. 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 any integer N , $[N]$ denotes the set $\{0, 1, \dots, N - 1\}$ of integers.

2.1 Notation

Throughout this paper we will consider a ring \mathcal{R} , which will be either \mathbb{Z} or the polynomial ring $\mathbb{Z}[X]/(X^d + 1)$ for d some power of 2. We will denote elements of \mathcal{R} by lowercase letters, (column) vectors over \mathcal{R} in bold lowercase and matrices over \mathcal{R} in bold uppercase. e.g. $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_k] \in \mathcal{R}^{l \times k}$ with $\mathbf{a}_i = (a_{i1}, \dots, a_{im})^T \in \mathcal{R}^l$. We will consider the norm of elements in \mathcal{R} to be $\|a\|_2 = |a|$ if $a \in \mathbb{Z}$, and $\|a\|_2 = \sqrt{\sum a_i^2}$ if $a = \sum a_i X^i \in \mathbb{Z}[X]/(X^d + 1)$. We extend the notation to vectors and matrices $\|\mathbf{a}\|_2 = \sqrt{\sum \|a_i\|_2^2}$, $\|\mathbf{A}\|_2 = \sqrt{\sum \|\mathbf{a}_i\|_2^2}$. We

will also consider the quotient ring $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$ for odd q . In the quotient ring, the norm of an element \mathcal{R}_q will be the norm of its unique representative \mathcal{R} with coefficients in $[-\frac{q-1}{2}, \frac{q-1}{2}]$.

We will also consider the operator norm of matrices over \mathcal{R} defined as $s_1(\mathbf{A}) = \max_{\|x\|_2 \neq 0} \left(\frac{\|\mathbf{A}x\|_2}{\|x\|_2} \right)$.

Probability Distributions. Let \mathcal{D} denote a distribution over some set. 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

$$D_\sigma(x) = \frac{\rho(x)}{\rho(\mathbb{Z})} \text{ where } \rho(\mathbb{Z}) = \sum_{v \in \mathbb{Z}} \rho(v).$$

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

Using the tail bounds for the 0-centered discrete Gaussian distribution (cf. [Ban93]), we can show that for any $\sigma > 0$ the norm of $x \leftarrow D_\sigma$ can be upper-bounded using σ . Namely, for any $k > 0$ it holds that

$$\Pr_{x \leftarrow D_\sigma} [|x| > k\sigma] \leq 2e^{-k^2/2}, \quad (3)$$

and when \mathbf{x} is drawn from D_σ^r , we have

$$\Pr_{\mathbf{x} \leftarrow D_\sigma^r} [\|\mathbf{x}\|_2 > \sqrt{2r} \cdot \sigma] < 2^{-r/4}. \quad (4)$$

We will abuse the notation $x \leftarrow D_\sigma$ when $x \in \mathbb{Z}[X]/(X^d + 1)$ to denote the distribution in which each coefficient of x is taken from D_σ . It is clear that in this case $\|x\|_2$ can be bounded using equation 4 with d instead of r .

2.2 Lattice-based Commitment Schemes

A commitment scheme allows a sender to create commitments to secret values, which she might then decide to reveal later. The main properties of commitment schemes are hiding and binding. Hiding guarantees that commitments do not leak information about the committed values, while binding guarantees that the sender cannot change her mind and open commitments to different values.

Formally, a non-interactive commitment scheme is a pair of probabilistic polynomial-time algorithms (Gen, Com). The setup algorithm $ck \leftarrow \text{Gen}(1^\lambda)$ generates a commitment key ck , which specifies message, randomness and commitment spaces $\mathbb{M}_{ck}, \mathbb{R}_{ck}, \mathbb{C}_{ck}$. It also specifies an efficiently sampleable probability

distribution $D_{R_{ck}}$ over R_{ck} and a binding set $B_{ck} \subset M_{ck} \times R_{ck}$. The commitment key also specifies a deterministic polynomial-time commitment function $\text{Com}_{ck} : M_{ck} \times R_{ck} \rightarrow C_{ck}$. We define $\text{Com}_{ck}(\mathbf{m})$ to be the probabilistic algorithm that given $\mathbf{m} \in M_{ck}$ samples $\mathbf{r} \leftarrow D_{R_{ck}}$ 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 $\mathbf{m}_0, \mathbf{m}_1 \in M_{ck}$ and for all $\mathbf{r}_0, \mathbf{r}_1 \in R_{ck}$

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

Definition 1 (Hiding). *The commitment scheme is computationally hiding if a commitment does not reveal the committed value. Formally, we say the commitment scheme is hiding if for all probabilistic polynomial time stateful interactive adversaries \mathcal{A}*

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

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

Definition 2 (Binding). *The commitment scheme is computationally binding if a commitment can only be opened to one value within the binding set B_{ck} . For all probabilistic polynomial time adversaries \mathcal{A}*

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

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

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

Ajtai's One-Way Function. The standard one-way function used in lattice cryptography maps a vector \mathcal{R}^n to \mathcal{R}^r via the function

$$f_{\mathbf{A}}(\mathbf{s}) = \mathbf{A}\mathbf{s},$$

where \mathbf{A} is a fixed, randomly-chosen matrix in $\mathcal{R}^{r \times n}$. Ajtai's seminal result [Ajt96] stated that when $\mathcal{R} = \mathbb{Z}_q$, it is as hard to find elements \mathbf{s} with some bounded norm $\|\mathbf{s}\| \leq B$ such that $f_{\mathbf{A}}(\mathbf{s}) = 0$ for random \mathbf{A} , as it is to find short vectors in any lattice of dimension r . This is called the short integer solution (SIS) problem and its hardness increases as r, q increase and B decreases; but somewhat surprisingly, the hardness of SIS is essentially unaffected by n as soon as n is large enough. The independence of the hardness from n holds both theoretically and in practice.

When solving SIS, one can ignore, if one wishes, any columns of \mathbf{A} by setting the corresponding coefficient of \mathbf{s} to 0, and solving SIS over the remaining columns. It was computed in [MR08] that if n is very large, then one should solve SIS for a submatrix where the number of columns is $n' = \sqrt{r \log q / \log \delta}$ for some

constant δ .⁴ With such a setting of n' , one expects to find a vector of length approximately

$$\min\{q, 2^{\sqrt{r \log q \log \delta}}\}. \quad (5)$$

Compressing Commitments Based on SIS. The fact that a larger n (after a certain point) does not decrease the security of the scheme allows one to construct simple compressing commitment schemes where the messages are elements in \mathbb{Z}_p for $p < q$. The commitment scheme, which was already implicit in the aforementioned work of Ajtai [Ajt96], 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. 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 $\mathbf{v} = \mathbf{A}_1 \mathbf{r} + \mathbf{A}_2 \mathbf{m} \bmod q$. By the leftover hash lemma, $(\mathbf{A}_1, \mathbf{A}_1 \mathbf{r} \bmod q)$ is statistically close to uniform, and so the commitment scheme is statistically hiding.⁵

To prove binding, note that if there are two different $(\mathbf{r}, \mathbf{m}) \neq (\mathbf{r}', \mathbf{m}')$ such that

$$\mathbf{v} = \mathbf{A}_1 \mathbf{r} + \mathbf{A}_2 \mathbf{m} = \mathbf{A}_1 \mathbf{r}' + \mathbf{A}_2 \mathbf{m}' \bmod q,$$

then

$$\mathbf{A}_1(\mathbf{r} - \mathbf{r}') + \mathbf{A}_2(\mathbf{m} - \mathbf{m}') = \mathbf{0} \bmod q,$$

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 $\mathbf{A} = [\mathbf{A}_1 \ \mathbf{A}_2]$. As long as the parameters are set such that $\|\mathbf{s}\|$ is smaller than the value in (5), the binding property of the commitment is based on an intractable version of the SIS problem.

The commitment scheme we will be working with in this paper works as follows:

$\text{Gen}(1^\lambda) \rightarrow ck$: Select a ring \mathcal{R} (either \mathbb{Z} or $\mathbb{Z}[X]/(X^d + 1)$), and parameter p, q, r, v, N, B, σ according to Table 2, and let $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$.

Pick uniformly at random matrices $\mathbf{A}_1 \leftarrow \mathcal{R}_q^{r \times r \log_p q}$ and $\mathbf{A}_2 \leftarrow \mathcal{R}_q^{r \times n}$.

Return $ck = (p, q, r, v, \ell, N, B, \mathcal{R}_q, \mathbf{A}_1, \mathbf{A}_2)$.

The commitment key defines message, randomness, commitment and binding spaces and distribution

$$\mathbf{M}_{ck} = \mathcal{R}_q^n \quad \mathbf{R}_{ck} = \mathcal{R}_q^{2r \log_p q} \quad \mathbf{C}_{ck} = \mathcal{R}_q^r$$

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

⁴ This constant δ is related to the optimal block-size in BKZ reduction [GN08], which is the currently best way of solving the SIS problem. Presently, the optimal lattice reductions set $\delta \approx 1.005$.

⁵ 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.

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

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 $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_\ell]$. It 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 $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_\ell]$.

2.3 Arguments of Knowledge

We aim to give efficient lattice-based proofs for arithmetic circuit satisfiability over \mathbb{Z}_p . The strategy we will employ is to commit to the values of a satisfying assignment to the wires, execute a range proof to demonstrate the committed values are within a suitable range, and to prove the committed values satisfy the constraints imposed by the arithmetic circuit. We will now formally define arguments of knowledge.

Let R be a polynomial time decidable ternary relation. The first input will contain some public parameters (aka common reference string) pp . We define the corresponding language L_{pp} indexed by the public parameters that consists of elements 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 PGen, 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 . Our convention is $b = 0$ means reject and $b = 1$ means accept.

Definition 3 (Argument of knowledge). *The proof system $(\text{PGen}, \mathcal{P}, \mathcal{V})$ is called an argument of knowledge for the relation R if it is complete and knowledge sound as defined below.*

Definition 4 (Statistical completeness). *$(\text{PGen}, \mathcal{P}, \mathcal{V})$ has statistical completeness with completeness error $\rho : \mathbb{N} \rightarrow [0; 1]$ if for all adversaries \mathcal{A}*

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

Definition 5 (Computational knowledge soundness). *$(\mathcal{K}, \mathcal{P}, \mathcal{V})$ is knowledge sound with knowledge soundness error $\epsilon : \mathbb{N} \rightarrow [0; 1]$ if for all deterministic polynomial time \mathcal{P}^* there exists an expected polynomial time extractor \mathcal{E} such that for all PPT adversaries \mathcal{A}*

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

It is sometimes useful to relax the definition of knowledge soundness to hold only for a larger relation \bar{R} such that $R \subset \bar{R}$. In this work, our zero-knowledge proofs

of pre-images will for instance 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 such that the vector $\bar{\mathbf{s}} = \begin{bmatrix} \bar{\mathbf{r}} - \bar{\mathbf{r}}' \\ \bar{\mathbf{m}} - \bar{\mathbf{m}}' \end{bmatrix}$ has norm smaller than the bound in (5).⁶

Concretely, if we would like to make a commitment to N values in \mathbb{Z}_p , then to satisfy (5) we need to make sure that $q > \|\bar{\mathbf{s}}\|$ and $\sqrt{r \log q \log \delta} > \log \|\bar{\mathbf{s}}\|$. In the protocols in our paper, we will have $\|\bar{\mathbf{s}}\| < N^2 p^2$ and $p < N$, which implies that $r = O(\log N)$.

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 6 (Special honest-verifier zero-knowledge). *A public-coin argument of knowledge (PGen, \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}*

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

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

where ϱ is the randomness used by the verifier.

Full Zero-Knowledge. In real life applications special honest verifier zero-knowledge may not suffice since a malicious verifier may give non-random challenges. However, it is easy to convert an SHVZK argument into a full zero-knowledge argument secure against *arbitrary* verifiers in the common reference string model using standard techniques, and when using the Fiat-Shamir heuristic to make the argument non-interactive SHVZK suffices to get zero-knowledge in the random oracle model.

3 Amortized Proofs of Knowledge

We will consider amortized proofs of knowledge for preimages of the Ajtai one-way function. Formally, given a matrix $\mathbf{A} \in \mathcal{R}_q^{r \times v}$ the relation we want to give a zero-knowledge proof of knowledge for is

⁶ Commitments over other rings, such as $\mathbb{Z}_q[X]/(X^d + 1)$ can be done in the same manner as above based on the hardness of the Ring-SIS problem [PR06, LM06] for which the bound in (5) still appears to hold in practice.

$$R = \left\{ \begin{array}{l} (pp, u, w) = ((q, \ell, \beta, \mathcal{R}, \mathbf{A}, c), \mathbf{T}, \mathbf{S}) \mid \\ (\mathbf{A}, \mathbf{S}, \mathbf{T}) \in \mathcal{R}_q^{r \times v} \times \mathcal{R}^{v \times \ell} \times \mathcal{R}_q^{r \times \ell} \wedge \mathbf{A}\mathbf{S} = c \cdot \mathbf{T} \wedge [\|\mathbf{s}_i\|_2 \leq \beta]_{i \in [\ell]} \end{array} \right\}$$

with $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_\ell]$ where \mathcal{R} is implicitly fixed in advance. The multiplier c depends on the instantiation of the proof: for $\mathcal{R} = \mathbb{Z}$ our proof achieves $c = 1$ and is exact, while for $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ it only guarantees that $c = 2$.




Fig. 1: Amortized proof for ℓ equations. The ring \mathcal{R} can be either \mathbb{Z} or $\mathbb{Z}[X]/(X^d + 1)$, the challenge set \mathcal{C} will be respectively $\{0, 1\}$ or $\{0\} \cup \{\pm X^j\}_{j < d}$

We consider a generalization of Σ -Protocols in which honest instances only complete with some constant probability $1/\rho$, this is to accommodate the fact that the rejection sampling step described in Lemma 1 only outputs 1 with probability $1/\rho$. In practice such a restriction is not too inconvenient: though the interactive protocol has to be repeated an average of ρ times to terminate, what we are interested in is usually the non-interactive protocol obtained by using the Fiat-Shamir transform, in which case the prover only has to output a proof when she obtains a challenge which passes the rejection step.

In our zero-knowledge proof, the prover will want to output a matrix \mathbf{Z} whose distribution should be independent of the secret matrix \mathbf{S} . During the protocol, the prover obtains $\mathbf{Z}' = \mathbf{B} + \mathbf{Y}$ where \mathbf{B} depends on the secret \mathbf{S} and \mathbf{Y} is a “masking” matrix each of whose coefficients is a discrete Gaussian with standard deviation σ . To remove the dependency of \mathbf{Z}' on \mathbf{B} , we use the rejection sampling procedure from [Lyu12] in Algorithm 1, which has the properties described in Lemma 1.

Algorithm 1 $\text{Rej}(\mathbf{Z}, \mathbf{B}, \sigma, \rho)$

```

 $u \leftarrow [0, 1)$ 
if  $u > \frac{1}{\rho} \cdot \exp\left(\frac{-2\langle \mathbf{Z}, \mathbf{B} \rangle + \|\mathbf{B}\|^2}{2\sigma^2}\right)$  then
  return 0
else
  return 1
end if

```

Lemma 1 ([Lyu12]). *Let $\mathbf{B} \in \mathcal{R}^{r \times n}$ be any matrix. Consider a procedure that samples a $\mathbf{Y} \leftarrow D_\sigma^{r \times n}$ 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 $D_\sigma^{r \times n}$.*

We give a useful lemma for knowledge extraction. In essence this lemma will be used to show that a prover who can output a verifying output for a challenge $\mathbf{c}_1, \dots, \mathbf{c}_\ell$ has a high probability of also being able to answer a challenge $\mathbf{c}'_1, \mathbf{c}'_2, \dots, \mathbf{c}'_\ell$ in which only $\mathbf{c}'_1 \neq \mathbf{c}_1$.

Lemma 2 ([Dam10]). *Let $\mathbf{H} \in \{0, 1\}^{\ell \times n}$ 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 $\varepsilon/2$ of ones. Then more than half of the ones in \mathbf{H} are located in heavy rows.*

We describe our proof system in Figure 1. Our first instantiation is with $\mathcal{R} = \mathbb{Z}$ in which case the one-way function will rely on the SIS problem and the challenge set will be $\mathcal{C}^{\ell \times n}$ for $\mathcal{C} = \{0, 1\}$, this solution allows the extractor of the protocol to obtain exact preimages of the \mathbf{t}_i and requires $n \geq \lambda + 2$. This ensures that communication only grows linearly in λ regardless of the size of ℓ (since $\mathbf{Z} \in \mathbb{Z}_q^{v \times n}$).

Theorem 1. *Let $\mathcal{R} = \mathbb{Z}$, $\mathcal{C} = \{0, 1\}$, $v, r = \text{poly}(\lambda)$, and $n \geq \lambda + 2$. Let $s > 0$ be an upper bound on $s_1(\mathbf{S})$, $\rho > 1$ be a constant, $\sigma \in \mathbb{R}$ be such that $\sigma \geq \frac{12}{\ln \rho} s \sqrt{\ell n}$, and $B = \sqrt{2v}\sigma$. Then the protocol described in Figure 1 is a zero-knowledge proof of knowledge for R .*

Proof. We will prove correctness and zero-knowledge here as the proofs are straightforward and very similar to prior works. We will however defer the proof of soundness to Lemma 3.

Correctness: If \mathcal{P} and \mathcal{V} are honest then the probability of abort is exponentially close to $1 - 1/\rho$ since $\|\mathbf{S}\mathbf{C}\|_2 \leq s_1(\mathbf{S})\|\mathbf{C}\|_2 \leq s\sqrt{\ell n}$. The equation verified by \mathcal{V} is true by construction of \mathbf{Z} . Since each coefficient of \mathbf{Z} is statistically close to D_σ , then according to (4) we have $\|\mathbf{z}_i\|_2 \leq \sqrt{2v}\sigma$ with overwhelming probability.

Honest-Verifier Zero-Knowledge: We will now prove that our protocol is honest-verifier zero-knowledge. More concretely, we show that the protocol is zero-knowledge when the prover does not abort prior to sending \mathbf{Z} . The reason that this is enough for practical purposes is that HVZK Σ -protocols can be turned into non-interactive proofs via the Fiat-Shamir transform. The non-interactive protocol generates the challenge \mathbf{C} as the hash of \mathbf{W} and \mathbf{T} , and otherwise repeats the prover's part of the protocol until a non-abort occurs, whereupon the prover outputs the transcript $(\mathbf{W}, \mathbf{C}, \mathbf{Z})$. Only the non-aborting transcripts will be seen by \mathcal{V} , and thus only they need to be simulated. Further below we will also sketch how to modify our protocol to obtain an interactive zero-knowledge proof.

Let $\mathcal{S}(\mathbf{A}, \mathbf{T})$ be the following PPT algorithm:

1. Sample $\mathbf{C} \leftarrow \{0, 1\}^{\ell \times n}$
2. Sample $\mathbf{Z} \leftarrow D_{\sigma}^{v \times n}$
3. Set $\mathbf{W} = \mathbf{AZ} - \mathbf{TC}$
4. Output $(\mathbf{W}, \mathbf{C}, \mathbf{Z})$

It is clear that \mathbf{Z} verifies with overwhelming probability. We already showed in the section on correctness that in the real protocol when no abort occurs the distribution of \mathbf{Z} is within statistical distance 2^{-100} of $D_{\sigma}^{v \times n}$. Since \mathbf{W} is completely determined by $\mathbf{A}, \mathbf{T}, \mathbf{Z}$ and \mathbf{C} , the distribution of $(\mathbf{W}, \mathbf{C}, \mathbf{Z})$ output by \mathcal{S} is within 2^{-100} of the distribution of these variables in the actual non-aborting run of the protocol.

To turn our proof into a full interactive HVZK proof, one can use the above simulator together with a standard transformation: in the first message of the protocol, \mathcal{P} will send a statistically hiding commitment of \mathbf{W} to the verifier. Later in the third round, she will then send both the opening and the message \mathbf{Z} , given that the protocol would not abort. The above simulator $\mathcal{S}(\mathbf{A}, \mathbf{T})$ can then, in the beginning of the protocol, flip a coin to determine if the simulation is aborting. If so, then it can just commit to a uniformly random value, and otherwise to the correct value \mathbf{W} . In order to make the protocol secure against arbitrary verifiers one can run an interactive coin-flipping protocol to generate \mathbf{C} .

Lemma 3 (Knowledge Soundness). *For any prover \mathcal{P}^* who succeeds with probability $\varepsilon > 2^{-\lambda}$ (i.e. $\geq 2^{-n+2}$) over her random tape $\chi \in \{0, 1\}^x$ and the challenge choice $\mathbf{C} \stackrel{\$}{\leftarrow} \mathcal{C}^{\ell \times n}$, there exists a knowledge extractor \mathcal{E} running in expected time $\text{poly}(\lambda)/\varepsilon$ who can extract a witness $\mathbf{S}' := (\mathbf{s}'_1, \dots, \mathbf{s}'_{\ell}) \in \mathcal{R}^{v \times \ell}$, such that $\mathbf{AS}' = \mathbf{T}$, and $\forall i \in [\ell] \|\mathbf{s}'_i\|_2 \leq 2B$.*

Proof. For $i \in [\ell]$, let $\mathbf{t}_i \in \mathcal{R}^n$ be the i th column of \mathbf{T} , and $\mathbf{c}_i^T \in \mathcal{R}^{1 \times n}$ be the i th row of \mathbf{C} (note that \mathbf{c}_i^T are not the transpose of the columns of \mathbf{C} but really its rows). Note that $\mathbf{t}_i \mathbf{c}_i^T \in \mathcal{R}^{r \times n}$ and $\mathbf{TC} = \sum_{i=1}^{\ell} \mathbf{t}_i \mathbf{c}_i^T$. For any fixed i , we describe an extractor \mathcal{E}_i who can extract a preimage of \mathbf{t}_i of norm less than $2B$ in expected $O(1/\varepsilon)$ executions, and the full result follows by running each extractor (of which there are $\ell = \text{poly}(\lambda)$).

Consider a matrix $\mathbf{H}_i \in \{0, 1\}^{2^{n(\ell-1)+x} \times 2^n}$ whose rows are indexed by the value of

$$(\chi, \mathbf{c}_1^T, \dots, \mathbf{c}_{i-1}^T, \mathbf{c}_{i+1}^T, \dots, \mathbf{c}_\ell^T)$$

and whose columns are indexed by the value of \mathbf{c}_i^T . An entry of \mathbf{H}_i will be 1 if \mathcal{P}^* succeeds for the corresponding challenge (i.e. produces an accepting \mathbf{Z}). We will say that a row of \mathbf{H}_i is "heavy" if it contains a fraction of at least $\varepsilon/2$ ones, i.e. if it contains more than $2^k * \varepsilon/2 > 2$ ones. The extractor \mathcal{E}_i will proceed as follow:

1. Run \mathcal{P}^* on random challenges \mathbf{C}' until it succeeds, and obtains \mathbf{Z}' that verifies. This takes expected time $1/\varepsilon$.
2. Run \mathcal{P}^* on random challenges \mathbf{C}'' where $\forall j \neq i, \mathbf{c}_j'^T = \mathbf{c}_j''^T$ and $\mathbf{c}_i''^T$ is freshly sampled. If after λ/ε attempts \mathcal{P}^* has not output a valid response \mathbf{Z}'' , abort.

The extractor \mathcal{E}_i runs in expected time $\text{poly}(\lambda)/\varepsilon$, and aborts with probability less than $1/2 + 2^{-\lambda}$. The running time is clear from the definition of \mathcal{E}_i . To compute the abort probability note that in step 2 all the challenges \mathbf{C}'' considered are in the same row of \mathbf{H}_i as \mathbf{C}' , if we call **Abort** the event where \mathcal{E}_i aborts and **Heavy** the event that \mathbf{C}' is in a row of \mathbf{H}_i , we have:

$$\Pr[\text{Abort}] = \Pr[\text{Abort}|\text{Heavy}] \Pr[\text{Heavy}] + \Pr[\text{Abort}|\neg\text{Heavy}] \Pr[\neg\text{Heavy}]$$

According to Lemma 2, $\Pr[\neg\text{Heavy}] < 1/2$. On the other hand if the row is heavy then for a random sample in this row \mathcal{P}^* has probability at least $\varepsilon/2 - 2^{-n} > \varepsilon/4$ of outputting a valid answer (the probability is $\varepsilon/2 - 2^{-n}$ and not $\varepsilon/2$ because we want a reply for a challenge different from \mathbf{C}'). Thus the probability that \mathcal{P}^* does not succeed on any of the λ/ε challenges \mathbf{C}'' is

$$\Pr[\text{Abort}|\text{Heavy}] < (1 - \varepsilon/4)^{\lambda\varepsilon} < e^{-4\lambda} < 2^{-\lambda}$$

and therefore $\Pr[\text{Abort}] < 1/2 + 2^{-\lambda}$. By running \mathcal{E}_i $O(\lambda)$ times we obtain an extractor that runs in expected time $\text{poly}(\lambda)/\varepsilon$ and outputs two valid pairs \mathbf{C}', \mathbf{Z}' and $\mathbf{C}'', \mathbf{Z}''$ such that $\forall j \neq i, \mathbf{c}_j'^T = \mathbf{c}_j''^T$, and $\mathbf{c}_i'^T \neq \mathbf{c}_i''^T$. Since both transcripts verify we know:

$$\mathbf{AZ}' = \mathbf{TC}' + \mathbf{W} = \sum_{j=1}^r \mathbf{t}_j \mathbf{c}_j'^T + \mathbf{W}$$

$$\mathbf{AZ}'' = \mathbf{TC}'' + \mathbf{W} = \sum_{j=1}^r \mathbf{t}_j \mathbf{c}_j''^T + \mathbf{W}$$

Which implies:

$$\mathbf{A}(\mathbf{Z}' - \mathbf{Z}'') = \sum_{j=1}^r \mathbf{t}_j (\mathbf{c}_j'^T - \mathbf{c}_j''^T) = \mathbf{t}_i (\mathbf{c}_i'^T - \mathbf{c}_i''^T)$$

If we consider an index $l \in [\ell]$ such that $\mathbf{c}'_i{}^T[l] \neq \mathbf{c}''_i{}^T[l]$, and assume w.l.o.g that $\mathbf{c}'_i{}^T[l] - \mathbf{c}''_i{}^T[l] = 1$, then by only considering the l^{th} column of the previous equation we obtain:

$$\mathbf{A}(\mathbf{z}'_i - \mathbf{z}''_i) = \mathbf{t}_i$$

where $\|\mathbf{z}'_i - \mathbf{z}''_i\|_2 \leq 2B$.

Our second instantiation uses $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ and $\mathcal{C} = \{0\} \cup \{\pm X^j\}_{j < d}$. This protocol only proves R with $c = 2$, i.e. the extractor will only obtain preimages of $2\mathbf{t}_i$ but the number of columns in the response matrix \mathbf{Z} can be reduced by a factor of $\log(2d + 1)$ as the soundness now only requires that $n \log(2d + 1) \geq \lambda + 2$. It is worth noting that in this protocol the values of r and v would typically be chosen to be around d times smaller than in the instantiation with $\mathcal{R} = \mathbb{Z}$, because \mathbf{A} will be a matrix of polynomials of degree d . We first give a lemma about the difference on monomials in $\mathbb{Z}[X]/(X^d + 1)$ which will be useful in the extraction.

Lemma 4 ([BCK⁺14] Lemma 3.2). *Let d be a power of 2, let $a, b \in \{\pm X^i : i \geq 0\} \cup \{0\}$. Then $2(a - b)^{-1} \bmod X^d + 1$ only has coefficients in $\{-1, 0, 1\}$. In particular $\|2(a - b)^{-1}\|_2 \leq \sqrt{d}$.*

Theorem 2. *Let $\mathcal{R} = \mathbb{Z}[X]/X^d + 1$, $\mathcal{C} = \{0\} \cup \{\pm X^j\}$, $v, r = \text{poly}(\lambda)$, and $n \geq (\lambda + 2)/\log(2d + 1)$. Let $s \in \mathbb{R}$ be an upper bound on $s_1(\mathbf{S})$, $\rho > 1$ be a constant, $\sigma \in \mathbb{R}$ be such that $\sigma \geq \frac{12}{\ln \rho} s \sqrt{\ell n}$, and $B = \sqrt{2md\sigma}$. Then the protocol described in Figure 1 is a SHVZK proof of knowledge.*

Proof. The proofs for correctness and zero-knowledge are nearly identical to the ones of Theorem 1. We will prove soundness in Lemma 5.

Lemma 5 (Knowledge Soundness). *For any prover \mathcal{P}^* who succeeds with probability $\varepsilon > 2^{-\lambda} (\geq 2^{-n \log(2d+1)+2})$ over his random tape $\chi \in \{0, 1\}^x$ and the challenge choice $\mathbf{C} \leftarrow \mathcal{C}^{\ell \times n}$ there exists a knowledge extractor \mathcal{E} who can extract a witness $\mathbf{S}' := (s'_1, \dots, s'_\ell) \in \mathcal{R}^{v \times \ell}$, such that $\mathbf{A}\mathbf{S}' = 2\mathbf{T}$, and $\forall i \in [\ell] \|\mathbf{s}'_i\|_2 \leq 2\sqrt{d}B$, in expected time $\text{poly}(\lambda)/\varepsilon$.*

Proof. The first part of the proof (obtaining \mathbf{C}', \mathbf{Z}' and $\mathbf{C}'', \mathbf{Z}''$) is identical to the one of Lemma 3 except for the fact that the matrix \mathbf{H}_i has different dimensions. Let $\delta = \log(2d + 1)$. Since for each $j \in [\ell]$, \mathbf{c}_j^T is sampled from a set of size $2^{n\delta}$, we have $\mathbf{H}_i \in \{0, 1\}^{2^{n\delta(\ell-1)+x} \times 2^{n\delta}}$. The heavy rows of \mathbf{H}_i will contain $2^{n\delta}\varepsilon/2 > 2$ ones, and the extractor can proceed as in the proof of Lemma 3. Assume that \mathcal{E}_i has extracted $\mathbf{C}', \mathbf{Z}''$ and $\mathbf{C}'', \mathbf{Z}'$ such that $\forall j \neq i, \mathbf{c}'_j{}^T = \mathbf{c}''_j{}^T$, and $\mathbf{c}'_i{}^T \neq \mathbf{c}''_i{}^T$. As previously we have:

$$\mathbf{A}(\mathbf{Z}' - \mathbf{Z}'') = \sum_{j=1}^{\ell} \mathbf{t}_j(\mathbf{c}'_j{}^T - \mathbf{c}''_j{}^T) = \mathbf{t}_i(\mathbf{c}'_i{}^T - \mathbf{c}''_i{}^T)$$

If we consider an index $l \in [\ell]$ such that $\mathbf{c}_i'^T[l] \neq \mathbf{c}_i''^T[l]$, since $\mathcal{C} = \{0\} \cup \{\pm X^j\}_{0 \leq j \leq d-1}$, we have according to Lemma 4 that there exists a $\mathbf{g} \in \mathcal{R}$ such that $2^{-1}(\mathbf{c}_i'^T[l] - \mathbf{c}_i''^T[l])\mathbf{g} = 1$ and $\|\mathbf{g}\|_2 \leq \sqrt{d}$. Hence:

$$\mathbf{A}(\mathbf{z}'_i - \mathbf{z}''_i)\mathbf{g} = 2\mathbf{t}_i \cdot 2^{-1}(\mathbf{c}_i'^T[l] - \mathbf{c}_i''^T[l])\mathbf{g} = 2\mathbf{t}_i$$

with $\|(\mathbf{z}'_i - \mathbf{z}''_i)\mathbf{g}\|_2 \leq 2\sqrt{d}B$.

4 Argument for the Satisfiability of an Arithmetic Circuit

In this section, we show how to construct arguments for the satisfiability of an arithmetic circuit based on the SIS assumption. We take inspiration from the arguments of [Gro09a, BCC⁺16] which rely on homomorphic commitments based on the hardness of discrete logarithm and translate them into the lattice settings. We obtain sublinear communication arguments with improved computational efficiency with respect to [Gro09a, BCC⁺16].

At a high level, [BCC⁺16] reduces the satisfiability of an arithmetic circuit to the verification of two sets of constraints: multiplication constraints, arising from multiplication gates; linear constraints, arising from additions and multiplication by constant gates. Then, it shows how to embed each of these sets of constraints into a polynomial equation over \mathbb{Z}_p . An argument for the satisfiability of an arithmetic circuit can then be constructed by giving arguments for the satisfiability of such polynomial equations. These arguments typically involve testing polynomial identities which can be done by evaluating the polynomials at random challenge points in \mathbb{Z}_p . In order to achieve negligible soundness one can invoke the Schwarz-Zippel lemma by setting p to be super-polynomial in the security parameter.

We give arithmetic circuit arguments over \mathbb{Z}_p for much smaller p (e.g. $p = \text{poly}(\lambda)$). Therefore, a straightforward translation of the above approach yield arguments which only have inverse polynomial soundness error, as $O(1/p)$ is inverse-polynomial in the security parameter in this setting. The soundness error could be reduced by repeated the protocol multiple times in parallel, resulting into a significant computational and communication overhead. For example, k parallel repetitions gives a soundness error of $O(1/p^k)$, but increases the communication cost and computational overhead by a factor k .

Therefore, we devise a more complex embedding technique in order to apply the Schwarz-Zippel lemma over larger fields. Cramer, Damgård and Keller give in [CDK14] an amortised proof of knowledge of k commitments over \mathbb{Z}_p are embedded into $GF(p^k)$, with soundness error $O(1/p^k)$. We follow a similar approach and embed the constraints for the satisfiability of the circuit into polynomial equations over an extension field. While [CDK14] only give a proof of knowledge, we also construct a product argument for the openings of k commitments over \mathbb{Z}_p embedded into an extension field of degree $2k$ with soundness $O(1/p^{2k})$.

We start by recalling how [BCC⁺16] embedded the satisfiability of an arithmetic circuit into a polynomial equations over \mathbb{Z}_p and then extend it to $GF(p^{2k})$.

Reduction of Circuit Satisfiability to a Hadamard Matrix Product and Linear Constraints over \mathbb{Z}_p . We consider arithmetic circuits with fan-in 2 addition and multiplication gates. Multiplication gates are directly represented as equations of the form $a \cdot b = c$, and we refer to a, b, c as the left, right and output wires, respectively.

We consider an arithmetic circuit containing $N = nm$ multiplication gates over a field \mathbb{Z}_p . Without loss of generality, we assume that the circuit has been pre-processed (see [BCC⁺16]), so that the input and the output wires feed into and go out from multiplication gates only. We number the multiplication gates from 1 to N and we arrange the inputs and outputs of these gates into three $m \times n$ matrices A, B and C such that the (i, j) entries of the matrices correspond to the left input, right input and output of the same multiplication gate. This process will lead to duplicate values, when a wire is the output of one multiplication gate and the input of another, or when it is used as input multiple times. We keep track of this by using a series of linear constraints. For example, if the output of the first multiplication gate is the right input of the second, we would write $c_1 - b_2 = 0$.

In addition to the constraints arising from wire duplication, we also add linear constraints representing the addition gates and multiplication by constant gates of the circuit. We then rewrite those equations so that the only wires that are referenced in the equations are those linked to (non-constant) multiplication gates, namely the wires included in A, B, C . We refer to [BCC⁺16] for a more detailed explanation of this process.

The satisfiability of an arithmetic circuit can be then described as a system of equations in the entries of A, B, C . The multiplication gates define a set of N equations

$$A \circ B = C \tag{6}$$

where \circ is the Hadamard (entry-wise) product.

The circuit description also contains constraints on the wires between multiplication gates. Denoting the rows of the matrices A, B, C as

$$\mathbf{a}_i = (a_{i,1}, \dots, a_{i,n}) \quad \mathbf{b}_i = (b_{i,1}, \dots, b_{i,n}) \quad \mathbf{c}_i = (c_{i,1}, \dots, c_{i,n}) \quad \text{for } i \in \{1, \dots, m\}$$

these constraints can be expressed as $U < 2N$ linear equations of inputs and outputs of multiplication gates of the form

$$\sum_{i=1}^m \mathbf{a}_i \cdot \mathbf{w}_{u,a,i} + \sum_{i=1}^m \mathbf{b}_i \cdot \mathbf{w}_{u,b,i} + \sum_{i=1}^m \mathbf{c}_i \cdot \mathbf{w}_{u,c,i} = K_u \quad \text{for } u \in \{1, \dots, U\} \tag{7}$$

for constant vectors $\mathbf{w}_{u,a,i}, \mathbf{w}_{u,b,i}, \mathbf{w}_{u,c,i}$ and scalars K_u .

For example, suppose that the circuit contains a single addition gate, with $a_{1,1}$ and $a_{1,2}$ as inputs, and $b_{1,1}$ as output. In this case, $U = 1$ and we would set $\mathbf{w}_{1,a,1} = (1, 1, 0, \dots, 0)$, $\mathbf{w}_{1,b,1} = (-1, 0, \dots, 0)$, and all other \mathbf{w} vectors would be set to $\mathbf{0}$. Then (7) would simply read

$$a_{1,1} + a_{1,2} - b_{1,1} = 0$$

to capture the constraint imposed by the addition gate.

In total, to capture all multiplications and linear constraints, we have $N + U$ equations that the wires must satisfy in order for the circuit to be satisfiable.

Reduction to Two Polynomial Equations. Let Y be a formal indeterminate. We will reduce the $N + U$ equations above to a two polynomial equations in Y by embedding distinct equations into distinct powers of Y . In our argument we will then require the prover to prove that these two equations hold when replacing Y by a random challenge received from the verifier.

Let \mathbf{Y} denote (Y, Y^2, \dots, Y^m) . Then, we can multiply (6) by \mathbf{Y} from the left to obtain $\mathbf{Y}(A \circ B) = \mathbf{Y}C$, or equivalently

$$\sum_{i=1}^m Y^i (\mathbf{a}_i \circ \mathbf{b}_i) = \sum_{i=1}^m Y^i \mathbf{c}_i$$

Since $(\mathbf{a} \circ \mathbf{b}) \cdot \mathbf{Y}' = (\mathbf{a} \circ \mathbf{Y}') \cdot \mathbf{b}$, we obtain the following expression

$$\sum_{i=1}^m \mathbf{a}_i \cdot (\mathbf{b}_i \circ \mathbf{Y}') Y^i = \left(\sum_{i=1}^m \mathbf{c}_i Y^i \cdot \mathbf{Y}' \right)$$

This is easily seen to be equivalent to (6), because $\mathbf{a}_i \circ \mathbf{b}_i = \mathbf{c}_i$ appears in the coefficient of Y^i .

Moreover, the U linear constraints on the wires in Eq. 7 are satisfied if and only if

$$\sum_{u=1}^U \left(\sum_{i=1}^m \mathbf{a}_i \cdot \mathbf{w}_{u,a,i} + \sum_{i=1}^m \mathbf{b}_i \cdot \mathbf{w}_{u,b,i} + \sum_{i=1}^m \mathbf{c}_i \cdot \mathbf{w}_{u,c,i} \right) Y^u = \sum_{u=1}^U K_u Y^u$$

since the u th constraint arises from comparing the coefficients of Y^u . Let us define

$$\begin{aligned} \mathbf{w}_{a,i}(Y) &= \sum_{u=1}^U \mathbf{w}_{u,a,i} Y^{N+1+u} & \mathbf{w}_{b,i}(Y) &= \sum_{u=1}^U \mathbf{w}_{u,b,i} Y^{N+1+u} \\ \mathbf{w}_{c,i}(Y) &= \sum_{u=1}^U \mathbf{w}_{u,c,i} Y^{N+1+u} & K(Y) &= \sum_{u=1}^U K_u Y^{N+1+u} \end{aligned}$$

Then the circuit is satisfied if and only if

$$\sum_{i=1}^m \mathbf{a}_i \cdot \mathbf{w}_{a,i}(Y) + \sum_{i=1}^m \mathbf{b}_i \cdot \mathbf{w}_{b,i}(Y) + \sum_{i=1}^m \mathbf{c}_i \cdot \mathbf{w}_{c,i}(Y) - K(Y) = 0 \quad (8)$$

$$\sum_{i=1}^m \mathbf{a}_i \circ \mathbf{b}_i Y^i = \sum_{i=1}^m \mathbf{c}_i Y^i \quad (9)$$

Sublinear Communication Product Argument. To give an argument for the satisfiability of an arithmetic circuit it is sufficient to give arguments showing that (8) and (9) are satisfied. For the purpose of constructing sublinear communication arguments, we craft polynomials which will have particular terms equal to zero if and only if (8) and (9) are satisfied. This can then be proved by having the prover reveal evaluations of the polynomials at random points to the verifier, who can check that the evaluations are correct using the homomorphic property of the commitment scheme. We define

$$\begin{aligned} \mathbf{a}(X) &:= \mathbf{a}_0 + \sum_{i=1}^m \mathbf{a}_i y^i X^i \\ \mathbf{b}(X) &:= \mathbf{b}_{m+1} + \sum_{i=1}^m \mathbf{b}_i X^{m+1-i} \\ \mathbf{c} &:= \sum_{i=1}^m \mathbf{c}_i y^i \\ \mathbf{a}(X) \circ \mathbf{b}(X) &= \mathbf{c} + \sum_{\substack{i=0 \\ i \neq m+1}}^{2m} \mathbf{h}_i X^i \end{aligned} \tag{10}$$

We have designed these polynomials such that the X^{m+1} term of $\mathbf{a}(X) \circ \mathbf{b}(X)$ is equal to $\sum_{i=1}^m \mathbf{c}_i y^i$. We conclude that the X^{m+1} term of $\mathbf{a}(X) \circ \mathbf{b}(X)$ is exactly \mathbf{c} if and only if (9) is satisfied. A similar approach can be followed to embed the satisfiability of (8) into the constant term of polynomial which is tested at random challenge evaluation points.

4.1 Amortisation Over Field Extensions

We now show how to extend the previous approach to work over field extensions. This will allow us to give an efficient amortised argument for the product of openings of commitments. This will be used to give efficient arguments for the satisfiability of an arithmetic circuit achieving sublinear communication and $O(1/p^{2k})$ soundness error.

Let $GF(p^{2k}) \simeq \mathbb{Z}_p[\phi]/\langle f(\phi) \rangle$, where f is a polynomial of degree $2k$ that is irreducible over \mathbb{Z}_p . Our goal is to embed k elements of \mathbb{Z}_p into the extension field in a way so that we can multiply two $GF(p^{2k})$ elements in a way that does not interfere with the products of the original \mathbb{Z}_p elements. Let e_1, \dots, e_k be distinct interpolation points in \mathbb{Z}_p (note that in particular, this forces $p > k$). Let $l_1(X), \dots, l_k(X)$ be the Lagrange polynomials associated with the points e_i , which have degree $k-1$. Let $l_0(X) = \prod_{j=1}^k (X - e_j)$, which has degree k .

Now, suppose that we have $a_1, \dots, a_k, b_1, \dots, b_k$ and c_1, \dots, c_k in \mathbb{Z}_p such that $a_j \cdot b_j = c_j \pmod p$ for each j . By evaluating the expression at each interpolation

point, we see that the following statement about polynomials holds over \mathbb{Z}_p .

$$\left(\sum_{j=1}^k a_j l_j(X) \right) \cdot \left(\sum_{j=1}^k b_j l_j(X) \right) \equiv \left(\sum_{j=1}^k c_j l_j(X) \right) \pmod{l_0(X)}$$

Therefore, there are $c'_0, \dots, c'_{k-2} \in \mathbb{Z}_p$ such that

$$\left(\sum_{j=1}^k a_j l_j(X) \right) \cdot \left(\sum_{j=1}^k b_j l_j(X) \right) = \left(\sum_{j=1}^k c_j l_j(X) \right) + l_0(X) \sum_{j=0}^{k-2} c'_j X^j$$

The degree of f is $2k$, so if we choose the basis

$$\mathcal{B} = l_1(\phi), \dots, l_k(\phi), l_0(\phi), \phi l_0(\phi), \dots, \phi^{k-1} l_0(\phi)$$

for $GF(p^{2k})$, we can perform multiplications of extension field elements without any overflow modulo f interfering with the individual product relations $a_i b_i = c_i$ in \mathbb{Z}_p . We can therefore port the above equality into $GF(p^{2k})$ as the equality

$$\left(\sum_{j=1}^k a_j l_j(\phi) \right) \cdot \left(\sum_{j=1}^k b_j l_j(\phi) \right) = \left(\sum_{j=1}^k c_j l_j(\phi) \right) + l_0(\phi) \sum_{j=0}^{k-2} c'_j \phi^j$$

This allows one multiplication of committed values to be performed without any overflow modulo f . As we shall see in the next subsection, this is sufficient for verifying multiplication triples for arithmetic circuit satisfiability.

We also need to be able to view single commitments to elements of \mathbb{Z}_p as elements of the extension field in a way that helps to verify linear consistency relations between the elements.

Now, suppose that we have $a_1, \dots, a_k, b_1, \dots, b_k$ and c_1, \dots, c_k in \mathbb{Z}_p , and coefficients $w_{a,1}, \dots, w_{a,k}, w_{b,1}, \dots, w_{b,k}$ and $w_{c,1}, \dots, w_{c,k}$ in \mathbb{Z}_p such that $\sum_{j=1}^k a_j w_{a,j} + \sum_{j=1}^k b_j w_{b,j} + \sum_{j=1}^k c_j w_{c,j} = K \pmod{p}$. By comparing coefficients, we see that the following statement about polynomials holds over \mathbb{Z}_p .

$$\begin{aligned} & \left(\sum_{j=1}^k a_j X^{j-1} \right) \cdot \left(\sum_{j=1}^k w_{a,j} X^{k-j} \right) + \left(\sum_{j=1}^k b_j X^{j-1} \right) \cdot \left(\sum_{j=1}^k w_{b,j} X^{k-j} \right) \\ & + \left(\sum_{j=1}^k c_j X^{j-1} \right) \cdot \left(\sum_{j=1}^k w_{c,j} X^{k-j} \right) = K X^{k-1} + \sum_{j=0, j \neq k-1}^{2k-2} K_j X^j \end{aligned}$$

where the K_j are extra terms determined from the a, b, c and w values.

If we choose the basis $\mathcal{B}' = 1, \phi, \phi^2, \dots, \phi^{2k-1}$ for $GF(p^{2k})$, we can perform multiplications of extension field elements in a way that always yields a useful linear relation in the ϕ^{k-1} term without any overflow modulo f .

By viewing multiplication in $GF(p^{2k})$ as a linear map over \mathbb{Z}_p^{2k} , we can simulate arithmetic in the extension field using arithmetic in \mathbb{Z}_p^{2k} .

Let $A_1, \dots, A_{2k} \in \mathcal{C}^{2k}$ be homomorphic commitments to single elements, $a_1, \dots, a_k \in \mathbb{Z}_p$. We can consider the tuple $\mathbf{A} = (A_1, \dots, A_k)$ to be a commitment to an element $\mathbf{a} = (a_1, \dots, a_{2k})$ of $GF(p^{2k})$. Now, if we consider $\mathbf{x} \in \mathbb{Z}_p^{2k}$ as an element of $GF(p^{2k})$, then there is a matrix $M_{\mathbf{x}}$ which simulates multiplication by \mathbf{x} in \mathbb{Z}_p^{2k} when we multiply on the left by $M_{\mathbf{x}}$. Since the A_i are homomorphic commitments, we can obtain a commitment to $\mathbf{x} * \mathbf{a}$ by computing $M_{\mathbf{x}}\mathbf{A}$, where $*$ represents multiplication in $GF(p^{2k})$.

Reduction of Circuit Satisfiability to a Hadamard Matrix Product and Linear Constraints over $GF(p^{2k})$. Let $N = mnk$ be the number of multiplication gates in the arithmetic circuit. To reduce circuit satisfiability to constraints over $GF(p^{2k})$, we can consider the same polynomial equations as before, written over $GF(p^{2k})$ rather than \mathbb{Z}_p . We consider the rows of matrices A , B , and C as before, but this time, we label the row vectors of the matrices $\mathbf{a}_{i,j}$, $\mathbf{b}_{i,j}$ and $\mathbf{c}_{i,j} \in \mathbb{Z}_p^n$, for $1 \leq i \leq m$ and $1 \leq j \leq k$. Now, we consider the row vectors $\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,k}$, which are elements of \mathbb{Z}_p^n , as an element in $GF(p^{2k})^n$.

Let $a_i = \begin{pmatrix} \mathbf{a}_{i,1} \\ \mathbf{a}_{i,2} \\ \vdots \\ \mathbf{a}_{i,k} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix}$ represent this element in $GF(p^{2k})^n$. Each column of the

matrix represents a separate element of $GF(p^{2k})$.

Satisfiability conditions over \mathbb{Z}_p were embedded using scalar products, denoted by \cdot , and element-wise products, denoted by \circ . If a and b in $\mathbb{Z}_p^{2k \times n}$ represent elements of $GF(p^{2k})^n$, then each column represents an element of $GF(p^{2k})$, and the scalar products and element-wise products of a and b are computed using the columns. We denote the element-wise product by $a \circ b$ and the scalar product by $a \odot b$ to avoid confusion with any other matrix products on a and b .

$$a = \begin{pmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \dots & \mathbf{v}_n \end{pmatrix} \quad b = \begin{pmatrix} \mathbf{w}_1 & \mathbf{w}_2 & \dots & \mathbf{w}_n \end{pmatrix}$$

$$a \circ b = \begin{pmatrix} M_{\mathbf{v}_1} \mathbf{w}_1 & M_{\mathbf{v}_2} \mathbf{w}_2 & \dots & M_{\mathbf{v}_n} \mathbf{w}_n \end{pmatrix} \quad a \odot b = M_{\mathbf{v}_1} \mathbf{w}_1 + M_{\mathbf{v}_2} \mathbf{w}_2 + \dots + M_{\mathbf{v}_n} \mathbf{w}_n$$

Note that in the verification equations, although the verifier computes high powers of random challenges \mathbf{x} and \mathbf{y} , the verifier only computes quadratic polynomials of values such as a and b which have been sent by the prover. This

is important, because when we expand a and b in terms of their coefficients a_i and b_i , we see that the verifier only computes expressions which have degree 2 in the prover's secret committed wire values, embedded as elements of $GF(p^{2k})$. Therefore, considering a field extension of degree $2k$ with the basis \mathcal{B} is sufficient for our purposes: we only need to ensure that a single multiplication in $GF(p^{2k})$ preserves the individual product relations embedded in the $GF(p)$ elements.

When embedding satisfiability conditions into a polynomial over \mathbb{Z}_p , using random challenges $x, y \in \mathbb{Z}_p$, the prover could send linear combinations of vectors $\mathbf{a}_i \in \mathbb{Z}_p^n$ such as

$$a(x) = a_0 + \sum_{i=1}^m \mathbf{a}_i y^i x^i$$

to the verifier.

However, when embedding satisfiability conditions into a polynomial over $GF(p^{2k})$, using random challenges $\mathbf{x}, \mathbf{y} \in GF(p^{2k})$, the prover sends linear combinations of vectors $\mathbf{a}_i \in GF(p^{2k})^n$ such as

$$a(x) = a_0 + \sum_{i=1}^m (M_{\mathbf{y}})^i (M_{\mathbf{x}})^i \mathbf{a}_i$$

Committing and Performing Calculations in a Lattice Setting. Commitment schemes based on lattice assumptions often require messages to be ‘small’ elements of the base ring in which the commitment is computed. Therefore, we consider the wire values in the arithmetic circuit to be integers in $[p]$ inside a larger ambient ring \mathbb{Z}_q where the commitments are computed.

We can still simulate the action of $GF(p^{2k})$ over the integers by applying the same multiplication matrices over the integers rather than working modulo p . Whenever the prover and verifier multiply by powers of random challenges $\mathbf{x} \in [p]^{2k}$, they reduce powers of matrices such as $M_{\mathbf{x}}$ and $M_{\mathbf{y}}$ modulo p before applying these matrices to commitments or openings. For example, the prover will send openings a and b to the verifier.

$$a = \sum_{i=0}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \bmod p) a_i \quad b = \sum_{i=0}^m (M_{\mathbf{x}}^{(m+1-i)} \bmod p) b_i$$

For this reason, the verification equations will compare quantities that are congruent modulo p , but not equal over the integers, or in \mathbb{Z}_q , as the prover and verifier will have computed and reduced various terms modulo p , but performed this reduction at different times during the computation. Therefore, the prover will send an additional commitment \mathbf{D} containing a message which is a multiple of p and corrects the discrepancy.

5 Parameter Selection

In this section we introduce notation for the parameters in our arithmetic circuit argument, and specify the choice of values in our arguments to ensure asymptotic security. Due to the large number of different variables used in the arithmetic circuit argument, and the fact that the arithmetic circuit argument and earlier proof of knowledge are quite independent of one another, we redefine certain variable names which were used earlier on for use in the arithmetic circuit argument. Parameter λ is dictated by the desired security level, and p and N come from the arithmetic circuit whose satisfiability is to be proven. All other parameters are derived from the table below, can be written in terms of λ, p and N , and are chosen in order to ensure that the commitment scheme is binding on a large enough message space for security.

Parameters and Asymptotic Sizes. In order to satisfy the constraints above, we choose the parameters in Table 2. Let λ be the security parameter, and suppose that we wish to verify an arithmetic circuit with N gates, over \mathbb{Z}_p .

Parameter	Size	Description
λ		Security parameter for our arguments
p	$\text{poly}(\lambda)$	Underlying field for the arithmetic circuit
N	$kmn = \text{poly}(\lambda)$	Number of multiplication gates in the arithmetic circuit
P	$O(nk^2m^2p^2)$	Maximum size of elements committed by honest prover
B	$O(PN)$	Soundness slack from proof of knowledge
P'	$P' = BP$	Commitment scheme must be binding up to elements in $[P']$.
n	$n \approx \sqrt{\frac{Nr \log q}{\lambda \log N \lambda p}}$	Controls length of vectors in the SAT argument
k	$k \approx \lambda / \log_2 p$	Controls soundness error of the SAT argument
m	$m = N / kn$	Number of commitments in SAT argument is $O(mk)$
q	$q \approx P' \sqrt{r}$	Modulus for SIS instances.
r	$r = O(\log n)$	Commitments lie in \mathbb{Z}_q^d .

Table 2: Parameter choices for our arithmetic circuit argument.

6 Product Argument

The following protocol allows the prover to prove that they know $N = nmk$ triples satisfying multiplicative relations.

We give parameters for our protocol in Section 5.

Consider the commitment scheme $\text{Com}_{ck} : \mathbb{Z}_q^n \times \mathbb{Z}_q^{n'} \mapsto \mathcal{C}$ introduced earlier in section 2.2, where ck consists of the public matrices used to generate a SIS instance. Let $A \in \mathbb{Z}_q^{2k \times n}$ and $R \in \mathbb{Z}_q^{2k \times n'}$. Define the extended commitment

scheme Com_{ck}^* as

$$\text{Com}_{ck}^*(A; R) := \begin{pmatrix} \text{Com}_{ck}(\mathbf{a}_1; \mathbf{r}_1) \\ \text{Com}_{ck}(\mathbf{a}_2; \mathbf{r}_2) \\ \vdots \\ \text{Com}_{ck}(\mathbf{a}_{2k}; \mathbf{r}_{2k}) \end{pmatrix}$$

where $\mathbf{a}_i \in \mathbb{Z}_p^n$, $\mathbf{r}_i \in \mathbb{Z}_p^{n'}$ are the row vectors of A and R .

Common Reference String: Commitment key ck . The basis \mathcal{B} for the extension field $GF(p^{2k})$, which specifies how elements should be multiplied.

Statement: Description of a set of $N = kmn$ multiplication relations over \mathbb{Z}_p .

Prover's Witness: Values $A_i, B_i, C_i \in \mathbb{Z}_p^{k \times n}$, $1 \leq i \leq m$, such that $\forall i, A_i \circ B_i \equiv C_i \pmod{p}$.

Argument:

\mathcal{P} Since $\forall i, A_i \circ B_i \equiv C_i \pmod{p}$, then for $1 \leq i \leq m$, we can write

$$\begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \circ \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} = \begin{bmatrix} C_i \\ C'_i \end{bmatrix} \pmod{p}$$

for some $C'_i \in [p]^{k \times n}$, $1 \leq i \leq m$, by our choice of basis \mathcal{B} .

The prover randomly selects $A_0, B_{m+1} \leftarrow D_{\sigma_1}^{2k \times n}$.

The prover selects α_i and β_i uniformly at random from $[p]^{k \times n'}$ and γ_i uniformly at random from $[p]^{2k \times n'}$ for $1 \leq i \leq m$, and selects $\alpha_0, \beta_{m+1} \leftarrow D_{\sigma_1}^{2k \times n'}$.

For $1 \leq i \leq m$, the prover computes

$$\begin{aligned} \mathbf{A}_i &= \text{Com}_{ck}^* \left(\begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix}; \begin{bmatrix} \alpha_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \right) & \mathbf{C}_i &= \text{Com}_{ck}^* \left(\begin{bmatrix} C_i \\ C'_i \end{bmatrix}; \gamma_i \right) \\ \mathbf{B}_i &= \text{Com}_{ck}^* \left(\begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix}; \begin{bmatrix} \beta_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \right) \end{aligned}$$

Note that by definition, \mathbf{A}_i and $\mathbf{B}_i \in \mathcal{C}^{2k}$ consist of k commitments and k trivial commitments in the k final components. The prover also computes

$$\mathbf{A}_0 = \text{Com}_{ck}^*(A_0; \alpha_0), \quad \mathbf{B}_{m+1} = \text{Com}_{ck}^*(B_{m+1}; \beta_{m+1})$$

The prover sends $\{\mathbf{A}_i\}_{i=0}^m, \{\mathbf{B}_i\}_{i=1}^{m+1}, \{\mathbf{C}_i\}_{i=1}^m$ to the verifier.

\mathcal{V} The verifier picks $\mathbf{y} \leftarrow [p]^{2k}$, and sends \mathbf{y} to the prover.

\mathcal{P} The prover computes polynomials $A(\mathbf{X}), B(\mathbf{X})$, which have matrix coefficients, in the indeterminate $\mathbf{X} \in \mathbb{Z}_q^{2k}$, and also computes C .

$$\begin{aligned} A(\mathbf{X}) &= A_0 + \sum_{i=1}^m M_{\mathbf{X}}^i (M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \\ B(\mathbf{X}) &= B_{m+1} + \sum_{i=1}^m M_{\mathbf{X}}^{m+1-i} \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \\ C &= \sum_{i=1}^m M_{\mathbf{y}}^i \begin{bmatrix} C_i \\ C'_i \end{bmatrix} \pmod p \end{aligned}$$

The prover computes $A(\mathbf{X}) \circ B(\mathbf{X}) \pmod p$.

$$A(\mathbf{X}) \circ B(\mathbf{X}) \pmod p = M_{\mathbf{X}}^{m+1} C + \sum_{l=0, l \neq m+1}^{2m} M_{\mathbf{X}}^l H_l \pmod p$$

where $H_l \in [p]^{2k \times n}$.

For $0 \leq l \leq 2m, l \neq 0$, the prover selects η_l uniformly at random from $[p]^{2k \times n'}$, and computes $\mathbf{H}_l = \text{Com}_{ck}^*(H_l; \eta_l)$.

The prover sends $\{\mathbf{H}_l\}_{l=0, l \neq m}$ to the verifier.

\mathcal{V} The verifier picks $\mathbf{x} \leftarrow [p]^{2k}$, and sends \mathbf{x} to the prover.

\mathcal{P} The prover computes the following values modulo p .

$$\begin{aligned} A &= A_0 + \sum_{i=1}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \\ \alpha &= \alpha_0 + \sum_{i=1}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} \alpha_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \\ B &= B_{m+1} + \sum_{i=1}^m (M_{\mathbf{x}}^{m+1-i} \pmod p) \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \\ \beta &= \beta_{m+1} + \sum_{i=1}^m (M_{\mathbf{x}}^{m+1-i} \pmod p) \begin{bmatrix} \beta_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \end{aligned}$$

Note that $A \equiv A(\mathbf{x}) \pmod p$ and $B \equiv B(\mathbf{x}) \pmod p$.

The prover computes

$$D = (A \circ B \pmod p) - \sum_{i=1}^m (M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} C_i \\ C'_i \end{bmatrix} - \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \pmod p) H_l$$

The prover randomly selects $\delta \leftarrow D_{\sigma_2}^{2k \times n'}$ and computes $\mathbf{D} = \text{Com}_{ck}^*(D; \delta)$.

The prover randomly selects $E \leftarrow p \cdot D_{\sigma_3}^{2k \times n}$, $\epsilon \leftarrow D_{\sigma_4}^{2k \times n'}$ and computes $\mathbf{E} = \text{Com}_{ck}^*(E; \epsilon)$. Note that E is 0 modulo p .

The prover sends \mathbf{D} and \mathbf{E} to the verifier.

- \mathcal{V} The verifier picks $\mathbf{z} \leftarrow [p]^{2k}$, and sends \mathbf{z} to the prover.
 \mathcal{P} The prover runs $\text{Rej}((A\|\alpha\|B\|\beta), (A\|\alpha\|B\|\beta) - (A_0\|\alpha_0\|B_{m+1}\|\beta_{m+1}), \sigma_1, e)$, and aborts according to the result.
 The prover computes the following

$$\rho = \sum_{i=1}^m (M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^i \bmod p) \gamma_i + \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \bmod p) \eta_l + \delta$$

- The prover runs $\text{Rej}(\rho, \rho - \delta, \sigma_2, e)$.
 The prover computes $\bar{D} = (M_{\mathbf{z}} \bmod p)D + E$ and $\bar{\delta} = (M_{\mathbf{z}} \bmod p)\delta + \epsilon$.
 The prover runs $\text{Rej}(\bar{D}/p, D/p, \sigma_3, e)$.
 The prover runs $\text{Rej}(\bar{\delta}, \delta, \sigma_4, e)$.
 The prover sends $A, \alpha, B, \beta, \rho, \bar{D}, \bar{\delta}$ to the verifier.
 \mathcal{V} The prover and the verifier engage in a proof-of-knowledge, as shown in Figure 1, including every commitment sent from the prover to the verifier.
 The verifier accepts if and only if

$$\begin{aligned} \text{Com}_{ck}^*(A; \alpha) &= \sum_{i=0}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \bmod p) \mathbf{A}_i \\ \text{Com}_{ck}^*(B; \beta) &= \sum_{i=1}^{m+1} (M_{\mathbf{x}}^{m+1-i} \bmod p) \mathbf{B}_i \\ \text{Com}_{ck}^*(A \circ B \bmod p; \rho) &= \sum_{i=1}^m (M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^i \bmod p) \mathbf{C}_i + \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \bmod p) \mathbf{H}_l + \mathbf{D} \\ \text{Com}_{ck}^*(\bar{D}; \bar{\delta}) &= (M_{\mathbf{z}} \bmod p) \mathbf{D} + \mathbf{E} \\ \bar{D} = 0 \bmod p \quad \|\bar{D}\|_2 &\leq 2\sqrt{kn}\sigma_3 p \\ \|(A\|\alpha\|B\|\beta)\|_2 \leq 4\sqrt{kn}\sigma_1 \quad \|\rho\|_2 \leq 2\sqrt{kn}\sigma_2 \quad \|\bar{\delta}\|_2 &\leq 2\sqrt{kn}\sigma_4 \end{aligned}$$

and the proof-of-knowledge is accepting.

Sizes of Standard Deviations.

$$\begin{aligned} \sigma_1 &= 48\sqrt{knkmp^2}, & \sigma_2 &= 72\sqrt{2knkmp}, \\ \sigma_3 &= 24\sqrt{2knkp}(1 + 6kmp), & \sigma_4 &= 24\sqrt{2k^2pn}\sigma_2 \end{aligned}$$

Security Analysis.

Theorem 3. *Given the statistically hiding, computationally binding commitment scheme based on SIS, the argument for multiplication triples has statistical completeness, statistical special honest verifier zero-knowledge and computational knowledge-soundness.*

The proof of Theorem 3 can be found in Appendix B

Efficiency. The above argument uses 7 moves of interaction and results in an overall 9 move argument when combined with the proof-of-knowledge sub-protocols. For the product argument, the prover must send $8mk+6k$ commitments to the verifier, and $14nk$ integers as commitments openings, plus the communication for the proof-of-knowledge. Sub-linear communication is achieved by setting parameters as in Table 2. This gives communication of approximately $O(\sqrt{N \log N})$ elements of \mathbb{Z}_q .

The prover’s computational costs are dominated by three computations:

- The cost of multiplying polynomials over $GF(p^{2k})$ to compute $A(\mathbf{X}) \circ B(\mathbf{X})$. This is $O(N \log m \log p (\log \lambda)^2)$ bit operations.
- The cost of computing commitments. Note that \mathbf{A}_i and \mathbf{B}_i consist partly of trivial commitments for $1 \leq i \leq m$. The dominant cost is bounded by $O(N \log N)$ multiplications in \mathbb{Z}_q .
- The cost of computing polynomials such as $A(\mathbf{X})$ and matrices such as A , α , and ρ . This is $O(N \log q (\log k)^2)$ bit operations.

For $q = \text{poly}(\lambda)$, this gives computational costs of $O(N \log N (\log \lambda)^2)$ bit operations for the prover.

The verifier’s computational costs are dominated by computing the same types of linear combinations as the prover, giving computational costs of $O(N (\log \lambda)^3)$ bit operations.

7 Linear Constraint Argument Description.

In Section A of the Supplementary material, we give a protocol which allows the prover to prove that $N = nmk$ committed values satisfy the linear consistency relations

$$\sum_{i=1, j=1}^{m, k} \mathbf{a}_{i, j} \cdot \mathbf{w}_{u, a, i, j} + \sum_{i=1, j=1}^{m, k} \mathbf{b}_{i, j} \cdot \mathbf{w}_{u, b, i, j} + \sum_{i=1, j=1}^{m, k} \mathbf{c}_{i, j} \cdot \mathbf{w}_{u, c, i, j} = K_u \quad \text{for } u \in \{1, \dots, U\} \quad (11)$$

Without loss of generality, we pad the linear consistency relations so that U is divisible by k .

The protocol, security proof, and complexity analysis are very similar to that of the argument for proving multiplication triples in the previous section.

We select parameters for our protocol in Section 5.

Security Analysis.

Theorem 4. *Given the statistically hiding, computationally binding commitment scheme based in SIS, the argument for linear consistency constraints has statistical completeness, statistical special honest verifier zero-knowledge and computational knowledge-soundness.*

The proof of Theorem 4 can be found in Appendix C

Efficiency. The above argument uses 7 moves of interaction and results in an overall 9 move argument when combined with the proof-of-knowledge sub-protocols. For the product argument, the prover must send $7km + 9k - 1$ commitments to the verifier, and $10nk + 2k$ integers as commitment openings, plus the communication for the proof-of-knowledge. The asymptotic costs of the protocol are the same as for the argument for multiplication triples in the previous section. Combined with the proof of knowledge, this gives an arithmetic circuit argument with the stated efficiency.

8 Arithmetic Circuit Argument

The product protocol given in 6 and the linear consistency protocol given in 7 imply an arithmetic circuit protocol with the same asymptotic efficiency as the two subprotocols, in which the prover forms $O(mk)$ commitments, each to n wire values in p , and runs both subprotocols in order to prove that they satisfy the arithmetic circuit, reusing the same commitments A_i , *to the wires in both subprotocols*.

This yields a zero-knowledge argument for arithmetic circuit satisfiability with communication costs $O(\sqrt{N \log N})$ elements of \mathbb{Z}_q , computational costs of $O(N \log N)$ for the prover, and approximately $O(N)$ for the verifier.

References

- AHIV17. Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In Thuraisingham et al. [TEMX17], pages 2087–2104.
- Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996.
- Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993.
- BBB⁺17. Benedikt Bunz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. Cryptology ePrint Archive, Report 2017/1066, 2017. <https://eprint.iacr.org/2017/1066>.
- BCC⁺16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Fischlin and Coron [FC16], pages 327–357.
- BCCT12. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, January 2012.
- BCCT13. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 111–120. ACM Press, June 2013.
- BCG⁺17. Jonathan Bootle, Andrea Cerulli, Essam Ghadafi, Jens Groth, Mohammad Hajiabadi, and Sune K. Jakobsen. Linear-time zero-knowledge proofs for arithmetic circuit satisfiability. In Tsuyoshi Takagi and Thomas Peyrin,

- editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 336–365. Springer, Heidelberg, December 2017.
- BCK⁺14. Fabrice Benhamouda, Jan Camenisch, Stephan Krenn, Vadim Lyubashevsky, and Gregory Neven. Better zero-knowledge proofs for lattice encryption and their application to group signatures. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 551–572. Springer, Heidelberg, December 2014.
- BD10. Rikke Bendlin and Ivan Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 201–218. Springer, Heidelberg, February 2010.
- BDLN16. Carsten Baum, Ivan Damgård, Kasper Green Larsen, and Michael Nielsen. How to prove knowledge of small secrets. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 478–498. Springer, Heidelberg, August 2016.
- BDOP16. Carsten Baum, Ivan Damgård, Sabine Oechsner, and Chris Peikert. Efficient commitments and zero-knowledge protocols from ring-SIS with applications to lattice-based threshold cryptosystems. Cryptology ePrint Archive, Report 2016/997, 2016. <http://eprint.iacr.org/2016/997>.
- BG14. Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *CT-RSA 2014*, volume 8366 of *LNCS*, pages 28–47. Springer, Heidelberg, February 2014.
- BKLP15. Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl, editors, *ESORICS 2015, Part I*, volume 9326 of *LNCS*, pages 305–325. Springer, Heidelberg, September 2015.
- CD97. Ronald Cramer and Ivan Damgård. Linear zero-knowledge - a note on efficient zero-knowledge proofs and arguments. In *29th ACM STOC*, pages 436–445. ACM Press, May 1997.
- CDG⁺17. Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In Thuraisingham et al. [TEMX17], pages 1825–1842.
- CDK14. Ronald Cramer, Ivan Damgård, and Marcel Keller. On the amortized complexity of zero-knowledge protocols. *Journal of Cryptology*, 27(2):284–316, April 2014.
- CDXY17. Ronald Cramer, Ivan Damgård, Chaoping Xing, and Chen Yuan. Amortized complexity of zero-knowledge proofs revisited: Achieving linear soundness slack. In Coron and Nielsen [CN17], pages 479–500.
- CN17. Jean-Sébastien Coron and Jesper Buus Nielsen, editors. *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*. Springer, Heidelberg, May 2017.
- Dam10. Ivan Damgård. On Σ -protocols, 2010. <http://www.cs.au.dk/~ivan/Sigma.pdf>.
- DDL13. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 40–56. Springer, Heidelberg, August 2013.

- DL12. Ivan Damgård and Adriana López-Alt. Zero-knowledge proofs with low amortized communication from lattice assumptions. In Ivan Visconti and Roberto De Prisco, editors, *SCN 12*, volume 7485 of *LNCS*, pages 38–56. Springer, Heidelberg, September 2012.
- dPL17. Rafaël del Pino and Vadim Lyubashevsky. Amortization with fewer equations for proving knowledge of small secrets. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 365–394. Springer, Heidelberg, August 2017.
- FC16. Marc Fischlin and Jean-Sébastien Coron, editors. *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*. Springer, Heidelberg, May 2016.
- GG98. Oded Goldreich and Shafi Goldwasser. On the limits of non-approximability of lattice problems. In *30th ACM STOC*, pages 1–9. ACM Press, May 1998.
- GGI⁺15. Craig Gentry, Jens Groth, Yuval Ishai, Chris Peikert, Amit Sahai, and Adam D. Smith. Using fully homomorphic hybrid encryption to minimize non-interactive zero-knowledge proofs. *Journal of Cryptology*, 28(4):820–843, October 2015.
- GGPR13. Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
- GH98. Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Information Processing Letters*, 1998.
- GLP12. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 530–547. Springer, Heidelberg, September 2012.
- GMO16. Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. Zkboo: Faster zero-knowledge for boolean circuits. In *25th USENIX Security Symposium*, pages 1069–1083, 2016.
- GMR85. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985.
- GN08. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 31–51. Springer, Heidelberg, April 2008.
- GQ88. Louis C. Guillou and Jean-Jacques Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In C. G. Günther, editor, *EUROCRYPT’88*, volume 330 of *LNCS*, pages 123–128. Springer, Heidelberg, May 1988.
- Gro09a. Jens Groth. Linear algebra with sub-linear zero-knowledge arguments. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 192–208. Springer, Heidelberg, August 2009.
- Gro09b. Jens Groth. Linear algebra with sub-linear zero-knowledge arguments. In *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, pages 192–208, 2009.
- Gro10a. Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, December 2010.
- Gro10b. Jens Groth. A verifiable secret shuffle of homomorphic encryptions. *J. Cryptology*, 23(4):546–579, 2010.

- Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Fischlin and Coron [FC16], pages 305–326.
- GVW02. Oded Goldreich, Salil P. Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1-2):1–53, 2002.
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- IKOS07. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30. ACM Press, June 2007.
- Kil92. Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *24th ACM STOC*, pages 723–732. ACM Press, May 1992.
- KR08. Yael Tauman Kalai and Ran Raz. Interactive PCP. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 536–547. Springer, Heidelberg, July 2008.
- Lip12. Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 169–189. Springer, Heidelberg, March 2012.
- LM06. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact Knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 144–155. Springer, Heidelberg, July 2006.
- LN17. Vadim Lyubashevsky and Gregory Neven. One-shot verifiable encryption from lattices. In Coron and Nielsen [CN17], pages 293–323.
- LNSW13. San Ling, Khoa Nguyen, Damien Stehlé, and Huaxiong Wang. Improved zero-knowledge proofs of knowledge for the ISIS problem, and applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 107–124. Springer, Heidelberg, February / March 2013.
- Lyu09. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Heidelberg, December 2009.
- Lyu12. Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012.
- MR04. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th FOCS*, pages 372–381. IEEE Computer Society Press, October 2004.
- MR08. Daniele Micciancio and Oded Regev. Lattice-based cryptography. In Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen, editors, *Chapter in Post-quantum Cryptography*, pages 147–191. Springer, 2008.
- MV03. Daniele Micciancio and Salil P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 282–298. Springer, Heidelberg, August 2003.
- PHGR13. Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.

- PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Heidelberg, March 2006.
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- Sch91. Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
- Ste94. Jacques Stern. A new identification scheme based on syndrome decoding. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 13–21. Springer, Heidelberg, August 1994.
- TEMX17. Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors. *ACM CCS 17*. ACM Press, October / November 2017.

Supplementary Material

A Protocol for Linear Consistency Constraints

$$\sum_{i=1, j=1}^{m, k} \mathbf{a}_{i,j} \cdot \mathbf{w}_{u,a,i,j} + \sum_{i=1, j=1}^{m, k} \mathbf{b}_{i,j} \cdot \mathbf{w}_{u,b,i,j} + \sum_{i=1, j=1}^{m, k} \mathbf{c}_{i,j} \cdot \mathbf{w}_{u,c,i,j} = K_u \quad \text{for } u \in \{1, \dots, U\} \quad (12)$$

Without loss of generality, we pad the linear consistency relations so that U is divisible by k .

We select parameters for our protocol in section 5.

Common Reference String: Commitment key ck . The basis \mathcal{B}' for the extension field $GF(p^{2k})$, which specifies how elements should be multiplied.

Statement: Description C of an arithmetic circuit containing U linear consistency relations, specified by the vectors $\mathbf{w}_{u,a,i,j}$, $\mathbf{w}_{u,b,i,j}$ and $\mathbf{w}_{u,c,i,j}$. A basis \mathcal{B}' for the extension field $GF(p^{2k})$.

Prover's Witness: Values $\mathbf{a}_{i,j}, \mathbf{b}_{i,j}, \mathbf{c}_{i,j} \in \mathbb{Z}_p^n$, $1 \leq i \leq m$, $1 \leq j \leq k$, such that

$$\sum_{i=1, j=1}^{m, k} \mathbf{a}_{i,j} \cdot \mathbf{w}_{u,a,i,j} + \sum_{i=1, j=1}^{m, k} \mathbf{b}_{i,j} \cdot \mathbf{w}_{u,b,i,j} + \sum_{i=1, j=1}^{m, k} \mathbf{c}_{i,j} \cdot \mathbf{w}_{u,c,i,j} = K_u \quad \text{for } u \in \{1, \dots, U\}$$

Argument:

\mathcal{P} Arrange $\mathbf{a}_{i,j}, \mathbf{b}_{i,j}, \mathbf{c}_{i,j} \in \mathbb{Z}_p^n$, $1 \leq i \leq m$, $1 \leq j \leq k$, into matrices A_i, B_i and C_i in $[p]^{k \times n}$, for $1 \leq i \leq m$.

The prover randomly selects A_0, B_0 and C_0 uniformly at random from $D_{\sigma_1}^{2k \times n}$.

The prover selects α_i, β_i and γ_i uniformly at random from $[p]^{k \times n'}$ for $1 \leq i \leq m$, and selects α_0, β_0 and γ_0 from $D_{\sigma_1}^{2k \times n'}$.

For $1 \leq i \leq m$, the prover computes

$$\begin{aligned} A_i &= \text{Com}_{ck}^* \left(\begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix}; \begin{bmatrix} \alpha_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \right) & B_i &= \text{Com}_{ck}^* \left(\begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix}; \begin{bmatrix} \beta_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \right) \\ C_i &= \text{Com}_{ck}^* \left(\begin{bmatrix} C_i \\ \mathbf{0}^{k \times n} \end{bmatrix}; \begin{bmatrix} \gamma_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \right) \end{aligned}$$

The prover also computes

$$\begin{aligned} A_0 &= \text{Com}_{ck}^* (A_0; \alpha_0) & B_0 &= \text{Com}_{ck}^* (B_0; \beta_0) \\ C_0 &= \text{Com}_{ck}^* (C_0; \gamma_0) \end{aligned}$$

The prover sends $\{A_i\}_{i=0}^m, \{B_i\}_{i=0}^m, \{C_i\}_{i=0}^m$ to the verifier.

\mathcal{V} The verifier picks $\mathbf{y} \leftarrow [p]^{2k}$, and sends \mathbf{y} to the prover.

\mathcal{P} The prover computes polynomials $a(\mathbf{X}), b(\mathbf{X}), c(\mathbf{X})$, which have matrix coefficients, in the indeterminate $\mathbf{X} \in \mathbb{Z}_q^{2k}$.

Write $w_{u,a,i}, w_{u,b,i}, w_{u,c,i} \in [p]^{2k \times n}$ as follows.

$$W_{q,a,i} = \begin{pmatrix} \mathbf{w}_{q,a,i,k} \\ \mathbf{w}_{q,a,i,k-1} \\ \vdots \\ \mathbf{w}_{q,a,i,1} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} \in [p]^{2k \times n}$$

and similarly for b and c written in place of a . Here, the final k out of $2k$ places contain zeroes.

$$A(\mathbf{X}) = A_0 + \sum_{i=1}^m M_{\mathbf{X}}^i \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \quad B(\mathbf{X}) = B_0 + \sum_{i=1}^m M_{\mathbf{X}}^i \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix}$$

$$C(\mathbf{X}) = C_0 + \sum_{i=1}^m M_{\mathbf{X}}^i \begin{bmatrix} C_i \\ \mathbf{0}^{k \times n} \end{bmatrix}$$

$$W_{a,i} = \sum_{u=1}^U (M_{\mathbf{y}}^u \bmod p) W_{u,a,i} \quad W_a(\mathbf{X}) = \sum_{i=1}^m M_{\mathbf{X}}^{m+1-i} W_{a,i}$$

$$W_{b,i} = \sum_{u=1}^U (M_{\mathbf{y}}^u \bmod p) W_{u,b,i} \quad W_b(\mathbf{X}) = \sum_{i=1}^m M_{\mathbf{X}}^{m+1-i} W_{b,i}$$

$$W_{c,i} = \sum_{u=1}^U (M_{\mathbf{y}}^u \bmod p) W_{u,c,i} \quad W_c(\mathbf{X}) = \sum_{i=1}^m M_{\mathbf{X}}^{m+1-i} W_{c,i}$$

The prover computes $a(\mathbf{X}) \odot w_a + b(\mathbf{X}) \odot w_b + c(\mathbf{X}) \odot w_c \bmod p$. The arrangements of the w constraints guarantee that when computing this expression, all of the linear constraints that we want to prove are sent to the $M_{\mathbf{X}}^{m+1}$ term, and the k th component of the field extension, resulting in the K_u terms below. Additional terms k' will be produced in other components of the $M_{\mathbf{X}}^{m+1}$ term, and other terms h_l other powers of $M_{\mathbf{X}}$. These will be committed to using commitments \mathbf{K}' and \mathbf{H}_l .

$$\begin{aligned} & a(\mathbf{X}) \odot w_a + b(\mathbf{X}) \odot w_b + c(\mathbf{X}) \odot w_c \bmod p \\ &= M_{\mathbf{X}}^{m+1} \sum_{u=1}^U M_{\mathbf{y}}^u \begin{pmatrix} 0 \\ \vdots \\ 0 \\ K_u \\ 0 \\ \vdots \\ 0 \end{pmatrix} \bmod p + M_{\mathbf{X}}^{m+1} \begin{pmatrix} k'_1 \\ \vdots \\ k'_{k-1} \\ 0 \\ k'_{k+1} \\ \vdots \\ k'_{2k} \end{pmatrix} + \sum_{l=0, l \neq m+1}^{2m} M_{\mathbf{X}}^l H_l \bmod p \end{aligned}$$

Check indices

The prover randomly selects η_l from $[p]^{2k \times n'}$ for $0 \leq l \leq 2m, l \neq m+1$.

The prover computes $\mathbf{H}_l = \text{Com}_{ck}^*(H_l; \eta_l)$ for $0 \leq l \leq 2m, l \neq m+1$.

The prover selects $\kappa' \in [p]^{2k \times n'}$ by choosing rows uniformly at random from $[p]^n$, but setting the k th row to zero. Write

$$K' = \begin{bmatrix} k'_1 \\ \vdots \\ k'_{k-2} \\ 0 \\ k'_k \\ \vdots \\ k'_{2k} \end{bmatrix}$$

The prover computes $\mathbf{K}' = \text{Com}_{ck}^*(K'; \kappa')$. By definition, the k th commitment making up \mathbf{K}' is a trivial one.

The prover sends $\mathbf{H}_0, \dots, \mathbf{H}_m, \mathbf{H}_{m+2}, \dots, \mathbf{H}_{2m}$ and \mathbf{K}' to the verifier.

\mathcal{V} The verifier picks $\mathbf{x} \leftarrow [p]^{2k}$, and sends \mathbf{x} to the prover.

\mathcal{P} The prover computes the following values modulo p .

$$\begin{aligned} A &= A_0 + \sum_{i=1}^m M_{\mathbf{x}}^i \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \pmod p & \alpha &= \alpha_0 + \sum_{i=1}^m M_{\mathbf{x}}^i \begin{bmatrix} \alpha_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \pmod p \\ B &= B_0 + \sum_{i=1}^m M_{\mathbf{x}}^i \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \pmod p & \beta &= \beta_0 + \sum_{i=1}^m M_{\mathbf{x}}^i \begin{bmatrix} \beta_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \pmod p \\ C &= C_0 + \sum_{i=1}^m M_{\mathbf{x}}^i \begin{bmatrix} C_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \pmod p & \gamma &= \gamma_0 + \sum_{i=1}^m M_{\mathbf{x}}^i \begin{bmatrix} \gamma_i \\ \mathbf{0}^{k \times n'} \end{bmatrix} \pmod p \\ W_a &= \sum_{i=1}^m M_{\mathbf{x}}^{m+1-i} W_{a,i} \pmod p & W_b &= \sum_{i=1}^m M_{\mathbf{x}}^{m+1-i} W_{b,i} \pmod p \\ W_c &= \sum_{i=1}^m M_{\mathbf{x}}^{m+1-i} W_{c,i} \pmod p \end{aligned}$$

Note that $A \equiv A(\mathbf{x}) \pmod p$ and similarly for B and C . Set

$$D = (A \odot W_a + B \odot W_b + C \odot W_c \pmod p)$$

$$- \sum_{u=1}^U M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^u \begin{pmatrix} 0 \\ \vdots \\ 0 \\ K_u \\ 0 \\ \vdots \\ 0 \end{pmatrix} \pmod p - (M_{\mathbf{x}}^{m+1} \pmod p) K' - \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \pmod p) H_l$$

The prover randomly selects $\delta \leftarrow D_{\sigma_2}^{2k \times n'}$ and computes $\mathbf{D} = \text{Com}_{ck}^*(D; \delta)$.
The prover randomly selects E from $p \cdot D_{\sigma_3}^{2k \times n}$. Note that E is 0 modulo p .
The prover randomly selects $\epsilon \leftarrow D_{\sigma_4}^{2k \times n}$ and computes $\mathbf{E} = \text{Com}_{ck}^*(E; \epsilon)$.
The prover computes the following value modulo p .

$$\rho = (M_{\mathbf{x}}^{m+1} \bmod p)\kappa + \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \bmod p)\eta_l + \delta$$

The prover sends \mathbf{D} and \mathbf{E} to the verifier.

\mathcal{V} The verifier picks $\mathbf{z} \leftarrow [p]^{2k}$, and sends \mathbf{z} to the prover.

\mathcal{P} The prover computes $\bar{D} = (M_{\mathbf{z}} \bmod p)D + E$ and $\bar{\delta} = (M_{\mathbf{z}} \bmod p)\delta + \epsilon$.

The prover runs $\text{Rej}(\bar{D}/p, D/p, \sigma_3, e)$.

The prover runs $\text{Rej}(\bar{\delta}, \delta, \sigma_2, e)$. Check SD

The prover runs $\text{Rej}((A||\alpha||B||\beta||C||\gamma), (A||\alpha||B||\beta||C||\gamma) - (A_0||\alpha_0||B_0||\beta_0||C_0||\gamma_0), \sigma_1, e)$.

The prover sends $A, \alpha, B, \beta, C, \gamma, \rho, \bar{D}$ and $\bar{\delta}$ to the verifier.

\mathcal{V} The prover and the verifier engage in a proof-of-knowledge including every commitment sent from the prover to the verifier.

The verifier accepts if

$$\begin{aligned} \mathbf{A} &= \sum_{i=0}^m (M_{\mathbf{x}}^i \bmod p)\mathbf{A}_i & \mathbf{B} &= \sum_{i=0}^m (M_{\mathbf{x}}^i \bmod p)\mathbf{B}_i \\ \mathbf{C} &= \sum_{i=0}^m (M_{\mathbf{x}}^i \bmod p)\mathbf{C}_i \end{aligned}$$

$$\begin{aligned} & \text{Com}_{ck}^*(A \odot W_a + B \odot W_b + C \odot W_c \bmod p; \rho) \\ &= \sum_{u=1}^U M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^u \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \text{Com}_{ck}(K_u; \mathbf{0}) \\ 0 \\ \vdots \\ 0 \end{pmatrix} \bmod p + (M_{\mathbf{x}}^{m+1} \bmod p)\mathbf{K}' \\ & \quad + \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \bmod p)\mathbf{H}_l + \mathbf{D} \end{aligned}$$

$$\begin{aligned} \bar{D} &= (M_{\mathbf{z}} \bmod p)\mathbf{D} + \mathbf{E} & \bar{D} &= 0 \bmod p \\ \|(A||\alpha||B||\beta||C||\gamma)\|_2 &\leq 2\sqrt{6kn}\sigma_1 & \|\rho\|_2 &\leq 2\sqrt{k}\sigma_2 \\ \|\bar{D}\|_2 &\leq 2\sqrt{k}\sigma_3 p & \|\bar{\delta}\|_2 &\leq 2\sqrt{k}\sigma_4 \end{aligned}$$

and the proof-of-knowledge is accepting.

Sizes of Standard Deviations.

$$\begin{aligned}\sigma_1 &= 48\sqrt{knkmp^2}, & \sigma_2 &= 24\sqrt{2knk}p^2(2m+1), \\ \sigma_3 &= 24\sqrt{2kn}(1+k+2mkp), & \sigma_4 &= 24\sqrt{2kn}(2m+1)kp^2\end{aligned}$$

B Proof of Theorem 3

Before proving the security of Theorem 3, we state a variation of the Schwarz-Zippel Lemma from [?]. This will be used when analysing the knowledge soundness of both the Hadamard product proof and the double-shift proof described in the next section.

Lemma 6. *Let \mathbb{F} be a field. Let P be a function of the following form, where p_{0,i_0} are constant values, and $p_{1,i_1}(Z_0), \dots, p_{u,i_u}(Z_0, \dots, Z_{u-1})$ are arbitrary functions and not necessarily polynomials.*

$$\begin{aligned}P(Z_0, \dots, Z_u) &= \sum_{i_0=-d_0}^{d_0} p_{0,i_0} Z_0^{i_0} + \sum_{i_1=-d_1, i_1 \neq 0}^{d_1} p_{1,i_1}(Z_0) Z_1^{i_1} \\ &+ \dots + \sum_{i_u=-d_u, i_u \neq 0}^{d_u} p_{u,i_u}(Z_0, \dots, Z_{u-1}) Z_u^{i_u}\end{aligned}$$

Let S be a finite subset of \mathbb{F} . Let z_0, \dots, z_u be selected at random independently and uniformly from S . Let \mathbf{E} be the event that at least one value among p_{0,i_0} or $p_{s,i_s}(z_0, \dots, z_{s-1})$ is not zero.

Then

$$\Pr[\{P(z_0, \dots, z_u) = 0\} \wedge \mathbf{E}] \leq \frac{\sum_{t=0}^u (2d_t + 1)}{|S|}$$

Proof (of Theorem 3). To prove completeness, first note that by the homomorphic property of the commitment scheme, careful inspection shows that an honest prover's messages will satisfy the verification equations. The protocol can only fail if the prover aborts, or if the norms of $A, \alpha, B, \beta, \rho, \bar{D}$ or $\bar{\delta}$ do not satisfy the bounds checked by the verifier. By Lemma 1 and the choice of σ_1 and σ_2 , we know that when the prover does not abort, the elements of A, α, B, β are distributed as D_{σ_1} , and the elements of $\bar{D}, \bar{\delta}$ are distributed as D_{σ_2} . Hence, by Lemma 3, with high probability, all of the norms satisfy the bounds checked by the verifier.

For special honest verifier zero-knowledge we start by picking random challenges \mathbf{x}, \mathbf{y} and \mathbf{z} from $[p]^{2k}$. Now, select the elements of A, α, B, β from D_{σ_1} , the elements of ρ from D_{σ_3} , the elements of \bar{D} from $p \cdot D_{\sigma_4}$, and the elements of $\bar{\delta}$ from D_{σ_5} . By Lemma 1, these values can only be distinguished from those in a real proof with very small probability, because the rejection sampling used

in a real proof ensures that these values have distributions close to the discrete gaussians given above. The simulator can create random commitments $\{\mathbf{A}_i\}_{i=1}^m, \{\mathbf{B}_i\}_{i=1}^m, \{\mathbf{C}_i\}_{i=1}^m, \mathbf{H}_0, \dots, \mathbf{H}_{2m}$. These are distributed just as in a real proof. Now, observe that the simulator can compute all terms in the first verification equation except for \mathbf{A}_0 , which is determined completely by this equation. Therefore, the simulator uses the first verification equation to compute \mathbf{A}_0 . Similarly, the simulator uses the second verification equation to compute \mathbf{B}_{m+1} , then the third verification equation to compute \mathbf{D} , and last, the fourth verification equation to compute \mathbf{E} . All terms in the simulated transcript are distributed as in a real proof, conditional on the verification equations being satisfied. Finally, we invoke the simulator for the proof-of-knowledge protocol. The obtained transcript is indistinguishable from an honestly generated one.

We prove that the arithmetic circuit protocol has knowledge soundness. First, note that by the soundness of the proof-of-knowledge protocol given in Figure 1, there exist openings $A_i, \alpha_i, B_i, \beta_i, C_i, \gamma_i, H_l, \eta_l, D, \delta, E, \epsilon$ to all of the commitments made by the prover during the protocol. By Theorem 1, and the parameters set in Figure 2, all of the extracted openings can be bounded by P' multiplied by the soundness slack from the protocol in Figure 1. Therefore, we can be sure that none of the quantities derived from the commitment openings below fall outside the bounds where the commitment scheme is binding.

By the bounds on D and δ , and since $\mathbf{z} \in [p]^{2k}$, the final verification equation and the binding property of the commitment scheme, we deduce that $\bar{D} = 0 \pmod p$ and $\bar{D} = (M_{\mathbf{z}} \pmod p)D + E$, which implies that $M_{\mathbf{z}}D + E = 0 \pmod p$. This relation modulo p can also be viewed as a relation over $GF(p^{2k})$. Suppose that D represents a non-zero element of $GF(p^{2k})$. Since $\mathbf{z} \pmod p$ is distributed as a uniformly random element of $GF(p)^{2k}$, then the probability that $M_{\mathbf{z}}D + E$ is zero over $GF(p^{2k})$ is $1/p^{2k}$. Therefore, we conclude that $D = 0 \pmod p$. Now consider the first and second verification equations. Again, by the binding property of the commitment scheme, and the size bounds on the openings, we deduce that

$$\begin{aligned} A &= A_0 + \sum_{i=1}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \\ B &= B_{m+1} + \sum_{i=1}^m (M_{\mathbf{x}}^{m+1-i} \pmod p) \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \end{aligned}$$

Similarly, from the third verification equation, we deduce that

$$\begin{bmatrix} A \\ \mathbf{0}^{k \times n} \end{bmatrix} \circ \begin{bmatrix} B \\ \mathbf{0}^{k \times n} \end{bmatrix} \pmod p = \sum_{i=1}^m (M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} C_i \\ C'_i \end{bmatrix} + \sum_{l=0, l \neq m+1}^{2m} (M_{\mathbf{x}}^l \pmod p) H_l + D$$

and hence, since $D = 0 \pmod p$, we see that

$$\begin{aligned} & \left(A_0 + \sum_{i=1}^m M_{\mathbf{x}}^i M_{\mathbf{y}}^i \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \right) \circ \left(B_{m+1} + \sum_{i=1}^m M_{\mathbf{x}}^{m+1-i} \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \right) \\ &= M_{\mathbf{x}}^{m+1} \sum_{i=1}^m M_{\mathbf{y}}^i \begin{bmatrix} C_i \\ C'_i \end{bmatrix} + \sum_{l=0, l \neq m+1}^{2m} M_{\mathbf{x}}^l H_l \pmod p \end{aligned}$$

This can be viewed as a polynomial relation over $GF(p^{2k})$. Since \mathbf{x}, \mathbf{y} are sampled uniformly at random from $[p]^{2k}$, they are uniform in $GF(p^{2k})$, and therefore, by Lemma 6 over $GF(p^{2k})$, whether the prover aborts or not, there is probability $2m/p^{2k}$ that this relation is satisfied for random \mathbf{x} and \mathbf{y} if there is not also an identity of polynomials. Therefore, this is an upper bound on the probability on the probability that the prover does not abort, and the polynomial relations are also satisfied. Comparing coefficients of $M_{\mathbf{x}}^{m+1}$, we see that $\begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \circ \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} = \begin{bmatrix} C_i \\ C'_i \end{bmatrix}$ for $1 \leq i \leq m$. Note that we know that the matrices on the left hand side contain zeroes in the bottom half by construction; \mathbf{A}_i and \mathbf{B}_i consist of trivial commitments in the final k components.

By our choice of basis for $GF(p^{2k})$, this implies that $A_i \circ B_i = C_i \pmod p$. Hence, the prover knows secret values satisfying $N = nmk$ multiplicative relations. \square

C Proof of Theorem 4

Proof. To prove completeness, first note that by the homomorphic property of the commitment scheme, careful inspection shows that an honest prover's messages will satisfy the verification equations. The protocol can only fail if the prover aborts, or if the norms of $A, \alpha, B, \beta, C, \gamma, \rho, \bar{D}$ or $\bar{\delta}$ do not satisfy the bounds checked by the verifier. By Lemma 1 and the choice of σ_3 and σ_4 , we know that when the prover does not abort, the elements of $A, \alpha, B, \beta, C, \gamma$ are distributed as D_{σ_3} , and the elements of $\bar{D}, \bar{\delta}$ are distributed as D_{σ_4} . Hence, by Lemma 3, with high probability, all of the norms satisfy the bounds checked by the verifier. For special honest verifier zero-knowledge we start by picking random challenges \mathbf{x}, \mathbf{y} and \mathbf{z} from $[p]^{2k}$. Now, select the elements of $A, \alpha, B, \beta, C, \gamma$ from D_{σ_1} , the elements of ρ from D_{σ_2} , the elements of \bar{D} from $p \cdot D_{\sigma_3}$, and the elements of $\bar{\delta}$ from D_{σ_4} . By Lemma 1, these values can only be distinguished from those in a real proof with very small probability, because the rejection sampling used in a real proof ensures that these values have distributions close to the discrete Gaussians given above. The simulator can create random commitments $\mathbf{A}_1, \dots, \mathbf{A}_m, \mathbf{B}_1, \dots, \mathbf{B}_m, \mathbf{C}_1, \mathbf{C}_m, \mathbf{H}_0, \dots, \mathbf{H}_{2m}, \mathbf{K}'$. These are distributed just as in a real proof. Now, observe that the simulator can compute all terms in the first verification equation except for \mathbf{A}_0 , which is determined completely by this equation. Therefore, the simulator uses the first verification equation to compute \mathbf{A}_0 . Similarly, the simulator uses the second verification equation

to compute \mathbf{B}_0 , then the third verification equation to compute \mathbf{C}_0 . Next, the simulator uses the fourth verification equation to compute \mathbf{D} , and last, the fifth verification equation to compute \mathbf{E} . All terms in the simulated transcript are distributed as in a real proof, conditional on the verification equations being satisfied. Finally, we invoke the simulator for the proof-of-knowledge protocol. The obtained transcript is indistinguishable from an honestly generated one. We prove that the linear constraint argument has knowledge soundness. First, note that by the soundness of the proof-of-knowledge protocol given in Figure 1, there exist openings $A_i, \alpha_i, B_i, \beta_i, C_i, \gamma_i, H_l, \eta_l, D, \delta, E, \epsilon$ to all of the commitments made by the prover during the protocol. By Theorem 1, and the parameters set in Figure 2, all of the extracted openings can be bounded by P' multiplied by the soundness slack from the protocol in Figure 1. Therefore, we can be sure that none of the quantities derived from the commitment openings below fall outside the bounds where the commitment scheme is binding. By the bounds on D and δ , and since $\mathbf{z} \in [p]^{2k}$, the final verification equation and the binding property of the commitment scheme, we deduce that $\bar{D} = 0 \pmod p$ and $\bar{D} = (M_{\mathbf{z}} \pmod p)D + E$, which implies that $M_{\mathbf{z}}D + E = 0 \pmod p$. This relation modulo p can also be viewed as a relation over $GF(p^{2k})$. Suppose that D represents a non-zero element of $GF(p^{2k})$. Since $\mathbf{z} \pmod p$ is distributed as a uniformly random element of $GF(p)^{2k}$, then the probability that $M_{\mathbf{z}}D + E$ is zero over $GF(p^{2k})$ is $1/p^{2k}$. Therefore, we conclude that $D = 0 \pmod p$. Now consider the first, second and third verification equations. Again, by the binding property of the commitment scheme, and the size bounds on the openings, we deduce that

$$\begin{aligned} A &= A_0 + \sum_{i=1}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} A_i \\ \mathbf{0}^{k \times n} \end{bmatrix} & B &= B_0 + \sum_{i=1}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} B_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \\ C &= C_0 + \sum_{i=1}^m (M_{\mathbf{x}}^i M_{\mathbf{y}}^i \pmod p) \begin{bmatrix} C_i \\ \mathbf{0}^{k \times n} \end{bmatrix} \end{aligned}$$

Similarly, from the third verification equation, and the fact that $D = 0 \pmod p$, we deduce that

$$\begin{aligned} A \odot W_a + B \odot W_b + C \odot W_c &= \sum_{u=1}^U M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^u \begin{pmatrix} 0 \\ \vdots \\ 0 \\ K_u \\ 0 \\ \vdots \\ 0 \end{pmatrix} + M_{\mathbf{x}}^{m+1} \begin{pmatrix} k'_1 \\ \vdots \\ k'_{k-1} \\ 0 \\ k'_{k+1} \\ \vdots \\ k'_{2k} \end{pmatrix} \\ &+ \sum_{l=0, l \neq m+1}^{2m} M_{\mathbf{x}}^l H_l \pmod p \end{aligned}$$

This can be viewed as a polynomial relation over $GF(p^{2k})$. Since \mathbf{x}, \mathbf{y} are sampled uniformly at random from $[p]^{2k}$, they are uniform in $GF(p^{2k})$, and therefore, by

Lemma 6, whether the prover aborts or not, there is probability $2m/p^{2k}$ that this relation is satisfied for random \mathbf{x} and \mathbf{y} if there is not also an identity of polynomials. Hence, if the prover does not abort, this is an upper bound on the probability that they satisfy these relations. Comparing coefficients of $M_{\mathbf{x}}$ implies that

$$\begin{pmatrix} \mathbf{a}_{i,1} \\ \mathbf{a}_{i,2} \\ \vdots \\ \mathbf{a}_{i,k} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} \odot \begin{pmatrix} \mathbf{w}_{a,i,k} \\ \mathbf{w}_{a,i,2} \\ \vdots \\ \mathbf{w}_{a,i,1} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} + \begin{pmatrix} \mathbf{b}_{i,1} \\ \mathbf{b}_{i,2} \\ \vdots \\ \mathbf{b}_{i,k} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} \odot \begin{pmatrix} \mathbf{w}_{b,i,k} \\ \mathbf{w}_{b,i,2} \\ \vdots \\ \mathbf{w}_{b,i,1} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} + \begin{pmatrix} \mathbf{c}_{i,1} \\ \mathbf{c}_{i,2} \\ \vdots \\ \mathbf{c}_{i,k} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} \odot \begin{pmatrix} \mathbf{w}_{c,i,k} \\ \mathbf{w}_{c,i,2} \\ \vdots \\ \mathbf{w}_{c,i,1} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix}$$

can be written as

$$\sum_{u=1}^U M_{\mathbf{x}}^{m+1} M_{\mathbf{y}}^u \begin{pmatrix} 0 \\ \vdots \\ 0 \\ K_u \\ 0 \\ \vdots \\ 0 \end{pmatrix} + M_{\mathbf{x}}^{m+1} \begin{pmatrix} k'_1 \\ \vdots \\ k'_{k-1} \\ 0 \\ k'_{k+1} \\ \vdots \\ k'_{2k} \end{pmatrix} + \sum_{l=0, l \neq m+1}^{2m} M_{\mathbf{x}}^l H_l \pmod p$$

Note that we know that the A , B and C vectors contain k zeroes, and that there is a zero in the vector of k'_j by construction rather than by any security property of the protocol; the prover only makes commitments to the components which are non-zero. By our choice of basis for $GF(p^{2k})$, comparing coefficients of $M_{\mathbf{y}}$ and looking at the k th component shows that each linear consistency equation must be satisfied. Hence, the prover knows secret values satisfying U linear consistency relations. \square