

Zero-Knowledge Proofs with Low Amortized Communication from Lattice Assumptions

Ivan Damgård^{1,*} and Adriana López-Alt²

¹ Aarhus University

² New York University

Abstract. We construct zero-knowledge proofs of plaintext knowledge (PoPK) and correct multiplication (PoPC) for the Regev encryption scheme with low amortized communication complexity. Previous constructions of both PoPK and PoPC had communication cost linear in the size of the public key (roughly quadratic in the lattice dimension, ignoring logarithmic factors). Furthermore, previous constructions of PoPK suffered from one of the following weaknesses: either the message and randomness space were restricted, or there was a *super-polynomial* gap between the size of the message and randomness that an honest prover chose and the size of which an accepting verifier would be convinced. The latter weakness was also present in the existent PoPC protocols.

In contrast, $O(n)$ proofs (for lattice dimension n) in our PoPK and PoPC protocols have communication cost linear in the public key. Thus, we improve the *amortized* communication cost of each proof by a factor linear in the lattice dimension. Furthermore, we allow the message space to be \mathbb{Z}_p and the randomness distribution to be the discrete Gaussian, both of which are natural choices for the Regev encryption scheme. Finally, in our schemes there is no gap between the size of the message and randomness that an honest prover chooses and the size of which an accepting verifier is convinced.

Our constructions use the “MPC-in-the-head” technique of Ishai et al. (STOC 2007). At the heart of our constructions is a protocol for proving that a value is bounded by some publicly known bound. This uses Lagrange’s Theorem that states that any positive integer can be expressed as the sum of four squares (an idea previously used by Boudot (EUROCRYPT 2000)), as well as techniques from Cramer and Damgård (CRYPTO 2009).

1 Introduction

The problem of secure multiparty computation (MPC) [19,6,12,31] is central in the field of modern cryptography. In this problem, N parties $\mathcal{P}_1, \dots, \mathcal{P}_N$ holding

* The first author acknowledges support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which part of this work was performed; and also from the CFEM research center (supported by the Danish Strategic Research Council) within which part of this work was performed.

private inputs x_1, \dots, x_N , respectively, wish to compute a function $f(x_1, \dots, x_N)$ on their inputs without revealing any information apart from the output of the evaluation (in particular, they wish to keep their inputs secret from the other parties). Solutions to this problem abound in the literature. Many of these solutions use the circuit rerandomization technique of Beaver [3] (see e.g. [20,23,4,21,14,5,8,15], among many others). Circuit rerandomization requires players to hold (additive) secret sharings of many random triples (a, b, c) such that $c = a \cdot b$ in some finite field. Traditionally, these triples are created using zero-knowledge proofs.

Bendlin et al. [8] use zero-knowledge proofs of plaintext knowledge (PoPK) and correct multiplication (PoCM) for this purpose. To see how this is done, consider the 2-party setting as an example. To obtain an additive secret sharing of random values a, b , players \mathcal{P}_1 and \mathcal{P}_2 can each choose random values u_1, v_1 and u_2, v_2 , respectively, and define $a = u_1 + u_2$ and $b = v_1 + v_2$. Obtaining an additive secret sharing of $c = a \cdot b$ is more involved. First, notice that $c = a \cdot b = (u_1 + u_2) \cdot (v_1 + v_2) = u_1v_1 + u_1v_2 + u_2v_1 + u_2v_2$. If \mathcal{P}_1 and \mathcal{P}_2 could obtain an additive sharing of each product $u_iv_j = y_{ij} + z_{ij}$ then they could obtain a sharing for c by simply adding each of these shares: $c = (y_{11} + y_{12} + y_{21} + y_{22}) + (z_{11} + z_{12} + z_{21} + z_{22})$. Thus, the problem reduces to having \mathcal{P}_1 and \mathcal{P}_2 obtain an additive sharing of the product of their inputs m_1 and m_2 , respectively (in this case u_i and v_j).

This can be done with the following protocol. \mathcal{P}_1 encrypts his input under his public key pk and obtains a ciphertext $c_1 = \text{Enc}_{pk}(m_1; r_1)$, which he sends to \mathcal{P}_2 . Upon receiving c_1 , \mathcal{P}_2 computes a ciphertext $c_x = \text{Enc}_{pk}(x; r_x)$ of a random plaintext x and computes $c_2 = m_2 \cdot c_1 + c_x$, sends it to \mathcal{P}_1 , and outputs $-x$ as his share. If the encryption scheme has certain homomorphic properties, then $c_2 = \text{Enc}_{pk}(m_1 m_2 + x)$. \mathcal{P}_1 decrypts c_2 and outputs $m_1 m_2 + x$ as his share, thus obtaining an additive sharing of $m_1 m_2$.

However, when players are malicious, \mathcal{P}_2 needs to ensure that c_1 is a valid ciphertext and \mathcal{P}_1 needs to ensure that \mathcal{P}_2 performed the multiplication step correctly. This can be done by having \mathcal{P}_1 and \mathcal{P}_2 provide zero-knowledge proofs that they performed their respective operations correctly: \mathcal{P}_1 sends a *proof of plaintext knowledge*, proving that there exist m_1, r_1 such that $c_1 = \text{Enc}_{pk}(m_1; r_1)$, and \mathcal{P}_2 sends a *proof of correct multiplication*, proving that there exist m_2, x, r_x such that $c_2 = m_2 \cdot c_1 + \text{Enc}_{pk}(x; r_x)$.

Unfortunately, these zero-knowledge proofs can incur a large communication cost, which increases the overall communication complexity of the MPC protocol in which they are used. A key observation is that even though many triples need to be created, they can be created simultaneously. This leads to the question of whether we can lower the *amortized* communication complexity of each proof, thus lowering the *total* communication cost of all proofs. In this work, we answer this question affirmatively when the encryption scheme used is the Regev encryption scheme [29], whose security is based on the hardness of the Learning with Errors (LWE) problem.

Related Work. Bendlin et al. [8], Bendlin and Damgård [7], and Asharov et al. [2,1] give constructions of proofs of plaintext knowledge. The work of [8] shows proofs of plaintext knowledge for any “semi-homomorphic” encryption scheme, an example of which is the Regev scheme. When applied to this scheme, the communication cost of *each* proof is linear in the size of the public key (roughly quadratic in the lattice dimension, ignoring logarithmic factors). The works of [7] and [2,1] show proofs of plaintext knowledge specifically for the Regev scheme, but here again, the communication cost of each proof is linear in the size of the public key. Similarly, [8] shows proofs of correct multiplication which, when applied to the Regev encryption scheme, have communication complexity linear in the public key size per proof.

Unfortunately, the protocol of [7] only works for message space $\{0,1\}$ and randomness in $\{0,1\}^m$. Furthermore, the proofs of [8] and [2,1] suffer from the following weakness. To guarantee zero-knowledge, an honest prover must choose the message and randomness from a sufficiently small range. But in order to guarantee soundness against a cheating prover, we can only guarantee that if the verifier accepts then the message and randomness come from a much larger interval. Thus, there is a gap between the size of the witness of an honest prover and the size of which an accepting verifier will be convinced. Such a gap, which turns out to be *super-polynomial* in the security parameter, is undesirable.

Our Results and Techniques. We improve upon these results by showing proofs of plaintext knowledge and correct multiplication where the cost of $O(n)$ proofs, where n is the lattice dimension, is linear in the public key size. Thus, we improve the amortized cost of each proof by a linear factor in the lattice dimension. Furthermore, our protocol does not suffer from the weakness of [8] and [2,1]; there is no gap between the size of the witness of an honest prover and the size of which an accepting verifier is convinced. The message space in our schemes can be \mathbb{Z}_p and the probability distribution for the randomness can be the discrete Gaussian.¹

Our proof system uses the “MPC-in-the-head” technique of Ishai et al. [22], who show how to construct zero-knowledge proofs from MPC protocols. The basic idea is as follows. For an NP relation $R(x, w)$ with statement x and witness w , the prover runs an MPC protocol for the function $f_x(w) = R(x, w)$ “in his head” and commits to the view of each of the players. The verifier then outputs a subset T of the players as challenge, and the prover opens the commitments to the views of the players in T . If the views are consistent, the verifier accepts.

This is the same technique that was used in [7] yet we improve upon it. First, we also show how to obtain proofs of correct multiplication. But more importantly, we expand the proofs to allow the message space to be \mathbb{Z}_p (rather than bits), and allow the randomness distribution to be the discrete Gaussian (rather than bit-vectors). To achieve this, we show a protocol that allows a dealer

¹ Technically, we’ll need the Regev scheme to have perfect correctness, so the randomness distribution will be a “truncated” discrete Gaussian that is statistically close to the discrete Gaussian, where values output according to the distribution are guaranteed to be small (as opposed to small with high probability).

to prove that the secret that he secret-shared among N players is bounded by some publicly known bound B . The intuition behind this proof is as follows. Let $[s]$ denote the sharing of secret s . The dealer distributes a sharing of B , $[B]$, and the players compute sharings $[B - s]$ and $[B + s]$ by locally adding their corresponding shares. We know that $-B < s < B$ if and only if both $B - s$ and $B + s$ are positive, so the problem of proving that s is bounded by B reduces to proving that a secret s' that has been secret shared among N players is positive.

For this, we use Lagrange’s Theorem that states that any positive integer can be written as the sum of four squares (see, e.g. [16]), and moreover, that these four squares can be computed efficiently [28,24] (a similar technique was used by Boudot [9]). The dealer computes u, v, w, y such that $s' = u^2 + v^2 + w^2 + y^2$, and distributes sharings $[u], [v], [w], [y]$. The players can then locally compute shares $[u^2 + v^2 + w^2 + y^2 - s'] = [0]$, and verify that these final shares reconstruct to 0.

However, we must ensure that the values u, v, w, y are all smaller than $\sqrt{q/8}$. Otherwise we can have overflow modulo q when we square and add the four squares, which would mean that we can no longer guarantee that the sum of the four squares is positive. For this, we use techniques from Cramer and Damgård [13]. The same techniques were used in [8], yet the key difference is that we use them to bound the numbers to be squared (and thus the bound can be loose), whereas in [8] they were used to bound the secrets themselves (thus leading to the gap discussed above). The use of this technique requires our modulus q to be super-polynomial in the security parameter λ (as was also the case in [7,8,2,1]). See Section 3 for more details.

Other Applications. Recently, Brakerski et al. showed that a variant of the Regev scheme is fully homomorphic [11,10]. The zero-knowledge PoPKs shown in this work can be used to prove that a ciphertext encrypted under this Regev-based FHE scheme is well-formed.

Presentation. In Section 2, we review some background needed for our constructions. This includes the IKOS construction (Section 2.2), packed secret sharing (Section 2.3), and a protocol for verifying the consistency of secret shares (Section 2.4). In Section 3, we show a protocol that allows parties to verify that a secret that is shared among them is numerically small. In Section 4 and Section 5 we show our protocols for proofs of plaintext knowledge and proofs of correct multiplication, respectively. Due to lack of space, we defer all proofs to the full version.

2 Preliminaries

2.1 Notation

The natural security parameter in this work is λ . We let $\mathbb{Z}_q = \{-q/2, \dots, q/2\}$ and use $a \bmod q$ to denote the mapping of a into the interval $(-q/2, q/2]$. We use $[n]$ to denote the set $\{1, \dots, n\} \subset \mathbb{Z}$.

We use boldface lower-case letters to represent vectors, such as $\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{Z}_q^n$. Throughout what follows, *vectors will be assumed to be column vectors*,

unless stated otherwise. We use subscripts to denote coordinates on a vector, e.g. u_i is the i th coordinate of vector \mathbf{u} . This is to differentiate between coordinates of a vector and elements in a sequence. For the latter case, we use superscripts: $m^{(i)}$ is the i th element of sequence $m^{(1)}, \dots, m^{(k)}$. We will also sometimes use the notation $(u_i)_{i \in [n]}$ to denote the vector (u_1, \dots, u_n) . We use boldface upper-case letters to represent matrices, such as $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. For a vector $\mathbf{x} = (x_1, \dots, x_n)$ and a scalar a , we let $a\mathbf{x} = (ax_1, \dots, ax_n)$.

For a distribution χ , we denote $x \leftarrow \chi$ to be the experiment of choosing x according to χ . If S is a set, then we use $x \leftarrow S$ to denote the experiment of choosing x from the uniform distribution on S . For a randomized function f , we write $f(x ; r)$ to denote the unique output of f on input x with random coins r . Denote $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ as the group of all reals in $[0, 1)$ with addition modulo 1. For $\alpha \in \mathbb{R}^+$, Ψ_α is defined to be the distribution on \mathbb{T} of a normal variable with mean 0 and standard deviation $\alpha/\sqrt{2\pi}$, reduced modulo 1. For any probability distribution ϕ over \mathbb{T} and integer $q \in \mathbb{Z}^+$, its *discretization* $\bar{\phi}$ is the discrete distribution over \mathbb{Z}_q of the random variable $[q \cdot X_\phi] \bmod q$, where $X_\phi \leftarrow \phi$.

We use lower case π to denote MPC protocols, such as π_f , and use upper case Π to denote zero-knowledge proof protocols, such as Π_R . We use greek letters to represent shares from a secret sharing. For example, $\alpha = (\alpha^{(1)}, \alpha^{(2)}, \dots, \alpha^{(N)})$ denotes the shares $\alpha^{(i)}$ of each of the N share holders.

2.2 Overview of IKOS Construction

Let $R(x, w)$ be a NP-relation. Consider the following N -player functionality f . The public statement x is known to all players $\mathcal{P}_1, \dots, \mathcal{P}_N$. The functionality takes the entire input w from a special player \mathcal{I} called the “input client”, and outputs $R(x, w)$ to all N players. Ishai et al. [22] show how to construct a zero-knowledge proof protocol for NP-relation R from a MPC protocol π_f for the functionality f described above. We give a high-level idea of the construction. The prover runs the MPC protocol π_f “in his head” and commits to the views V_1, \dots, V_N of the N players. The verifier then chooses a subset $T \subset [N]$, and the prover opens his commitments to views $\{V_i\}_{i \in T}$. The verifier accepts iff the commitment openings are successful, the revealed views are consistent, and the output in each view is 1.

We show the formal statement of the result in Theorem 1, but first recall the security properties that the underlying MPC protocol will need to satisfy in the construction. The following definitions are taken almost verbatim from [22].

Definition 1 (Correctness). *We say that a protocol π realizes functionality f with perfect correctness if for all inputs (x, w) , the probability that the output of some player is different from the output of f is 0, where the probability is taken over the random inputs r_1, \dots, r_N .*

Definition 2 ((Statistical) t -Privacy). *Let $t \in [N]$. We say a protocol π realizes functionality f with statistical t -privacy if there exists a PPT simulator Sim such that for all inputs (x, w) and all sets of corrupted players $T \subset [N]$ with*

$|T| \leq t$, the joint view $(\text{VIEW}(P_i))_{i \in T}$ of players in T is distributed statistically close to $\text{Sim}(T, x, R_T(x, w))$.

Definition 3 (*t*-Robustness). Let $t \in [N]$. We say a protocol π realizes functionality f with perfect t -robustness if it is perfectly correct in the presence of a semi-honest adversary, and for any computationally unbounded malicious adversary corrupting \mathcal{I} and a set T of at most t players, for all inputs x , it holds that if there does not exist w such that $f(x, w) = 1$, then the probability that an uncorrupted player $P_i \notin T$ outputs 1 is 0.

Theorem 1 ([22]). Let f be the N -player functionality with input client \mathcal{I} described above. Suppose that π_f is a protocol that realizes f with perfect t -robustness (in the malicious model) and statistical t -privacy (in the semi-honest model), where $t = \Omega(\lambda)$, and $N = ct$ for some constant $c > 1$. Given π_f and an unconditionally-binding commitment scheme, it is possible to construct a computational honest-verifier zero-knowledge proof protocol $\Pi_{R, \mathcal{I}, t}$ for the NP-relation R , with negligible (in λ) soundness error.

One of the nice properties about the [22] construction is that we get *broadcast for free* because the Prover can simply send the broadcasted messages directly to the Verifier. Therefore, the communication cost of broadcasting a message is simply the size of the message. We also get *coin-flipping among the players for free* because the (honest) Verifier can simply provide the random value. Therefore, the communication cost of coin-flipping for a value is simply the size of the value. We will use these two facts in our constructions. Also, as observed by [7], if we use a commitment scheme that allows us to commit to strings with only a constant additive length increase such as those implicit in [27], then the zero-knowledge proof protocol $\Pi_{R, \mathcal{I}, t}$ (asymptotically) conserves the communication complexity of the underlying MPC protocol π_f .

Finally, using general zero-knowledge techniques, it is possible to convert the honest-verifier zero-knowledge proof protocol $\Pi_{R, \mathcal{I}, t}$ obtained from Theorem 1 into a full zero-knowledge protocol, while (asymptotically) preserving the communication complexity of the protocol. One such technique is described in [22].

2.3 Packed Secret Sharing

We will use the packed secret sharing technique of Franklin and Yung [17]. Similar to Shamir secret sharing over \mathbb{Z}_q [30], packed secret sharing allows a dealer to share a vector of k values $\mathbf{x} = (x_1, x_2, \dots, x_k)$ using a single random polynomial of degree at most d . To guarantee security against at most t corrupted players, we must have $d \geq t + k - 1$. The idea is to choose a random polynomial $P(\cdot)$ of degree at most d , subject to the condition $P(-j + 1) = x_j$ for $j \in [k]$. The share of player i is, as usual, the value $\alpha_i = P(i)$.

We use $[\mathbf{x}]_d$ to denote a packed secret-sharing $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_N) \in \mathbb{Z}_q^N$ for N players of the block \mathbf{x} using a polynomial of degree at most d . We call $[\mathbf{x}]_d$ a d -sharing of \mathbf{x} . We say \mathbf{x} is *correctly shared* if every honest player \mathcal{P}_i is holding a share α_i of \mathbf{x} , such that there exists a degree at most d polynomial $P(\cdot)$ with

$P(i) = \alpha_i$ for $i \in N$, and $P(-j + 1) = x_j$ for $j \in [k]$. Any (perhaps incomplete) set of shares is called d -consistent if these shares lie on a polynomial of degree at most d .

Let $\mathbf{Z} \in \mathbb{Z}_q^{m \times k}$ be a matrix of secrets. Suppose we have d -sharings of the rows of \mathbf{Z} : $[\mathbf{Z}_1]_d, \dots, [\mathbf{Z}_m]_d \in \mathbb{Z}_q^{1 \times N}$. We define $\Psi \in \mathbb{Z}_q^{m \times N}$, called a d -share matrix of \mathbf{Z} , to be a matrix

$$\Psi = \begin{bmatrix} [\mathbf{Z}_1]_d \\ \vdots \\ [\mathbf{Z}_m]_d \end{bmatrix} \in \mathbb{Z}_q^{m \times N}$$

Note that the shares held by \mathcal{P}_i are precisely the entries in the i th column vector of Ψ , denoted by $\psi^{(i)}$.

For any function $f : \mathbb{Z}_q^{m \times 1} \rightarrow \mathbb{Z}_q^{m' \times 1}$, we abuse notation and write

$$f(\Psi) = f \left(\begin{bmatrix} [\mathbf{Z}_1]_d \\ \vdots \\ [\mathbf{Z}_m]_d \end{bmatrix} \right) = \begin{bmatrix} [\mathbf{Y}_1]_{d'} \\ \vdots \\ [\mathbf{Y}_{m'}]_{d'} \end{bmatrix},$$

to signify that each player \mathcal{P}_i locally applies f to his shares of all $[\mathbf{Z}_j]_d$'s to obtain his share of each $[\mathbf{Y}_j]_{d'}$. In other words, if Ψ is the d -share matrix of \mathbf{Z} then each player locally computes $f(\psi^{(i)}) = \phi^{(i)}$, where $\Phi = [\phi^{(1)}, \dots, \phi^{(N)}] \in \mathbb{Z}_q^{m' \times N}$ is the d' -share matrix of \mathbf{Y} containing the \mathbf{Y}_j 's as rows.

It is easy to see that if $f(\mathbf{x})$ is a linear function and we define f_i to be f with its output restricted to the i th coordinate (i.e. $f(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_{m'}(\mathbf{x}))^\top$), then

$$f \left(\begin{bmatrix} [\mathbf{Z}_1]_d \\ \vdots \\ [\mathbf{Z}_m]_d \end{bmatrix} \right) = \begin{bmatrix} [f_1(\mathbf{z}^{(1)}) , \dots , f_1(\mathbf{z}^{(k)})]_d \\ \vdots \\ [f_{m'}(\mathbf{z}^{(1)}) , \dots , f_{m'}(\mathbf{z}^{(k)})]_d \end{bmatrix}$$

Note that if f is a linear function, then the sharings obtained as a result of applying f are also d -sharings. In particular, if each player \mathcal{P}_i multiplies his share vector $\psi^{(i)}$ by a matrix $\mathbf{M} \in \mathbb{Z}_q^{m' \times m}$, the player obtains a $(m' \times 1)$ -vector representing his corresponding shares of:

$$\mathbf{M}\Psi = \begin{bmatrix} [\mathbf{M}_1 \mathbf{z}^{(1)} , \dots , \mathbf{M}_1 \mathbf{z}^{(k)}]_d \\ \vdots \\ [\mathbf{M}_{m'} \mathbf{z}^{(1)} , \dots , \mathbf{M}_{m'} \mathbf{z}^{(k)}]_d \end{bmatrix} = \begin{bmatrix} \left[(\mathbf{M}_1 \mathbf{z}^{(j)})_{j \in [k]} \right]_d \\ \vdots \\ \left[(\mathbf{M}_{m'} \mathbf{z}^{(j)})_{j \in [k]} \right]_d \end{bmatrix} = \begin{bmatrix} [(\mathbf{MZ})_1]_d \\ \vdots \\ [(\mathbf{MZ})_{m'}]_d \end{bmatrix},$$

where $(\mathbf{MZ})_i$ is the i th row of the matrix \mathbf{MZ} .

Parameters. We discuss requirements on the parameters of the scheme. We let $N = c_1 t$ for $c_1 > 2$, satisfying the requirements of the IKOS construction. In order to guarantee privacy of the secret shares, we must have $d \geq t + k - 1$. We

will sometimes use $(d/2)$ -shares, so we assume $d/2 \geq t + k - 1$. Furthermore, we must have enough honest players so that their shares alone can determine a polynomial of degree d (in case corrupt players do not send their shares for reconstruction). We therefore need $N - t \geq d \geq d/2 \geq t + k - 1$. For our choice of N this yields $k \leq (c_1 - 2)t + 1$. Thus, we assume $k = \Theta(t)$. Also, in order to have enough evaluation points, we must have $q > k + N$. Henceforth, we will use this choice of parameters.

2.4 Verifying Consistency of Shares

We now describe a protocol that can be used by N parties to check that their shares are d -consistent. Security is guaranteed if at most $t < N/2$ parties are corrupted. Players check $N - 2t$ sets of shares at a time. More formally, let $\mathbf{Z} \in \mathbb{Z}_q^{(N-2t) \times k}$ be a matrix of secrets, and suppose d -shares $[\mathbf{Z}_1]_d, \dots, [\mathbf{Z}_{N-2t}]_d$ of the rows of \mathbf{Z} are distributed among the N players. The players want to verify that each sharing is d -consistent without revealing their individual shares. Beerliová-Trubíniová and Hirt [4] describe a protocol in which the N parties can perform this check when they hold N sharings (as opposed to $N - 2t$, as described here) and sharing $[\mathbf{Z}_i]_d$ was created by player \mathcal{P}_i . Bendlin and Damgård [7] extend this protocol to the case when all the shares were prepared by a (possibly corrupt) input client \mathcal{I} . We describe the protocol of [7] in Figure 1. In the protocol, all players receive as common input a *hyper-invertible* matrix $\mathbf{M} \in \mathbb{Z}_q^{N \times (N-t)}$ for $q > 2N$. Informally, a hyper-invertible matrix is a matrix such that every square submatrix of \mathbf{M} is invertible. Beerliová-Trubíniová and Hirt [4] show how such matrices can be constructed.

Lemma 1. *The protocol π_{CHECK} described in Figure 1 allows N players, at most t of which are corrupted, to verify with zero error probability that $(N - 2t)$ pack-sharings, each of $k = \Theta(t)$ secrets in \mathbb{Z}_q , are d -consistent (for $d \geq t + k - 1$). It is t -private in the presence of a semi-honest adversary, t -robust in the presence of a malicious adversary, and has communication complexity $N(N + t) \log q$.*

2.5 Regev Encryption Scheme

Before presenting the Regev encryption scheme [29], we first introduce the hardness assumption on which its security is based. For positive integers $n = n(\lambda)$ and $q = q(\lambda) \geq 2$, a vector $\mathbf{s} \in \mathbb{Z}_q^n$, and a probability distribution χ on \mathbb{Z}_q , let $A_{\mathbf{s}, \chi}$ be the distribution obtained by choosing $\mathbf{a} \leftarrow \mathbb{Z}_q^n$ and $x \leftarrow \chi$, and outputting $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + x) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.

Learning with Errors (LWE $_{n,q,\chi}$ and dLWE $_{n,q,\chi}$). The Learning with Errors problem LWE $_{n,q,\chi}$ is defined as follows. Given $m = \text{poly}(n)$ samples chosen according to $A_{\mathbf{s}, \chi}$ for *uniformly chosen* $\mathbf{s} \in \mathbb{Z}_q^n$, output \mathbf{s} with noticeable probability. The Decisional Learning with Errors problem dLWE $_{n,q,\chi}$ is to distinguish (with non-negligible advantage) $m = \text{poly}(n)$ samples chosen according to $A_{\mathbf{s}, \chi}$ for *uniformly chosen* $\mathbf{s} \in \mathbb{Z}_q^n$, from m samples chosen uniformly at random from

Protocol π_{Check} between parties $(\mathcal{P}_1, \dots, \mathcal{P}_N)$ to verify d -consistency of shares.

Common input: hyper-invertible matrix $\mathbf{M} \in \mathbb{Z}_q^{N \times (N-t)}$

Input to \mathcal{P}_i : corresponding shares of $[\mathbf{Z}_1]_d, \dots, [\mathbf{Z}_{(N-2t)}]_d$.

1. Input client \mathcal{I} chooses and d -shares random vectors in $\mathbb{Z}_q^{1 \times k}$. Let $[\mathbf{Z}_{N-2t+1}]_d, \dots, [\mathbf{Z}_{N-t}]_d$ be the resulting shares. Augment matrix \mathbf{Z} with rows $\mathbf{Z}_{N-2t+1}, \dots, \mathbf{Z}_{N-t}$ to obtain matrix $\mathbf{Z}' \in \mathbb{Z}_q^{(N-t) \times k}$. Let $\Psi \in \mathbb{Z}_q^{(N-t) \times N}$ be the d -share matrix of \mathbf{Z}' .
2. Players locally compute:

$$\Phi = \mathbf{M}\Psi = \begin{bmatrix} [(\mathbf{M}\mathbf{Z}')_1]_d \\ \vdots \\ [(\mathbf{M}\mathbf{Z}')_N]_d \end{bmatrix} \in \mathbb{Z}_q^{N \times N}$$

3. The players reconstruct the resulting shares, each towards a different player: player \mathcal{P}_i receives Φ_i . Each player verifies that the shares he receives are d -consistent and broadcasts “ABORT” if he finds a fault, and otherwise broadcasts “OK”.
4. If all players broadcast “OK” then the players conclude that the initial shares were d -consistent.

Fig. 1. Protocol π_{CHECK} to verify consistency of shares

$\mathbb{Z}_q^n \times \mathbb{Z}_q$. In other words, if $d\text{LWE}_{n,q,\chi}$ is hard then $A_{s,\chi}$ is pseudorandom. We will use $\chi = \tilde{\Psi}_\alpha$ and in this case, we write $\text{LWE}_{n,q,\alpha}$ to mean $\text{LWE}_{n,q,\tilde{\Psi}_\alpha}$.

Discrete Gaussian Distribution. We present an elementary fact that shows that the discrete Gaussian distribution with standard deviation r outputs an element x with $\|x\| \leq r\sqrt{n}$ with high probability.

Lemma 2 (see [25], Theorem 4.4). *Let $n \in \mathbb{N}$. For any real number $r > \omega(\sqrt{\log n})$, we have $\Pr_{\mathbf{x} \leftarrow D_{\mathbb{Z}^n, r}}[\|\mathbf{x}\| > r\sqrt{n}] \leq 2^{-n+1}$.*

Using Lemma 2 together with the fact that for all $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_\infty \geq \|\mathbf{x}\|/\sqrt{n}$ we arrive at the following bound.

Lemma 3. *Let $n \in \mathbb{N}$. For any real number $r > \omega(\sqrt{\log n})$, we have $\Pr_{\mathbf{x} \leftarrow D_{\mathbb{Z}^n, r}}[\|\mathbf{x}\|_\infty > r] \leq 2^{-n+1}$.*

This allows us to define a *truncated* Gaussian distribution that always outputs (with probability 1) elements with ℓ_∞ norm less than r . Simply define the truncated Gaussian $\overline{D}_{\mathbb{Z}^n, r}$ over \mathbb{Z}^n with standard deviation r to sample a vector according to the discrete Gaussian $D_{\mathbb{Z}^n, r}$ and repeat the sampling if the vector has ℓ_∞ norm greater than r . We will use the truncated discrete Gaussian in our schemes to ensure that samples are bounded by r in each coordinate (and can thus ensure perfect correctness), but state security in terms of the discrete

Gaussian. Since the distributions are statistically close, all results stated using the discrete Gaussian also hold when using the truncated distribution.

We present a generalized version of the Regev encryption scheme [29] (with the modifications of [18]), using the truncated discrete Gaussian (as above). The scheme is parametrized by integers $n = n(\lambda), m = m(\lambda) > n, q = q(\lambda), r = r(\lambda)$, and $p = p(\lambda) < q$. The message space is $\mathcal{M} = \mathbb{Z}_p$, the ciphertext space is $\mathcal{C} = (\mathbb{Z}_q^n, \mathbb{Z}_q)$. All operations are performed over \mathbb{Z}_q .

- **KeyGen**(1^n): Output $sk = \mathbf{s}, pk = (\mathbf{A}, \mathbf{b})$, where $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{x} \leftarrow \chi^m$, $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{x} \in \mathbb{Z}_q^m$.
- **Enc** $_{pk}(m)$: Output (\mathbf{u}, c) , where $\mathbf{r} \leftarrow \overline{D}_{\mathbb{Z}_q^{m,r}}$, $\mathbf{u} = \mathbf{A}\mathbf{r} \in \mathbb{Z}_q^{n \times 1}$, $c = \mathbf{b}^\top \mathbf{r} + m \cdot \lfloor q/p \rfloor \in \mathbb{Z}_q$.
- **Dec** $_{sk}(\mathbf{u}, c)$: Output $m = \lfloor (c - \mathbf{s}^\top \mathbf{u}) \cdot p/q \rfloor$.

Theorem 2 ([29,18]). *Let $q \geq 5prm, \alpha \leq 1/(p \cdot r \sqrt{m} \cdot \omega(\sqrt{\log \lambda}))$, $\chi = \overline{\Psi}_\alpha, m \geq 2(n+1) \log q + \omega(\log \lambda)$. With this choice of parameters, the Regev encryption scheme is correct and IND-CPA-secure, assuming $\text{LWE}_{n,q,\chi}$ is hard.*

Parameters and Worst-case Guarantees. Our construction requires the modulus q to be super-polynomial in the security parameter λ . More specifically, we require $\sqrt{q/8} > 2^{\omega(\log \lambda)} \cdot m \cdot \max(p/2, r)$. We can use any choice of parameters that satisfies this constraint and keeps the cryptosystem secure.

One option is to let the dimension of the lattice be our security parameter, ie. $n = \lambda$ and set our modulus q to be exponential in the lattice dimension n . Peikert [26] showed that for such a large q , $\text{LWE}_{n,q,\alpha}$ is as hard as $\text{GapSVP}_{\tilde{O}(n/\alpha)}$ if q is a product of primes, each of polynomial size. The works of [7,8] use this choice of parameters.

Another possible choice is to let $n = \lambda^{1/\epsilon}$ for some $\epsilon \in (0, 1)$ (e.g. $n = \lambda^2$), $p, r, m = \text{poly}(\lambda)$ and let $q = 2^{n^\epsilon}$ be subexponential in the lattice dimension n . In this case, we can rely on Regev's quantum reduction [29] to $\text{GapSVP}_{\tilde{O}(n/\alpha)}$ or Peikert's classical reduction [26] to $\text{GapSVP}_{\zeta, \gamma}$ where $\gamma(n) \geq n/(\alpha \sqrt{\log n})$, $\zeta(n) \geq \gamma(n)$ and $q \geq \zeta \cdot \omega(\sqrt{\log n/n})$. The work of [2,1] uses this choice of parameters.

3 Verifying that Secrets are Numerically Small

At the heart of our constructions of proofs of plaintext knowledge and correct multiplication, we will use a protocol that allows a dealer (in our case the input client \mathcal{I}) to prove to the players that the secret that he secret-shared among them is bounded by some publicly known bound B . Formally, let $\mathbf{R} \in \mathbb{Z}_q^{m \times k}$ be a matrix of secrets. And suppose that a dealer has distributed d -sharings of the rows of \mathbf{R} : $[\mathbf{R}_1]_d, \dots, [\mathbf{R}_m]_d$ between N players. We show a protocol π_{VERSM} that allows the dealer to prove to each player \mathcal{P}_i , without revealing \mathbf{R} , that all secrets in \mathbf{R} are smaller than $B \ll q/2$.

We first have the dealer compute and distribute a sharing $[\mathbf{b}]_d$ of $\mathbf{b} = (B, \dots, B) \in \mathbb{Z}_q^k$. Players can then compute

$$\begin{bmatrix} [\mathbf{b}]_d \\ \vdots \\ [\mathbf{b}]_d \end{bmatrix} - \begin{bmatrix} [\mathbf{R}_1]_d \\ \vdots \\ [\mathbf{R}_m]_d \end{bmatrix} = \begin{bmatrix} [\mathbf{b} - \mathbf{R}_1]_d \\ \vdots \\ [\mathbf{b} - \mathbf{R}_m]_d \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} [\mathbf{b}]_d \\ \vdots \\ [\mathbf{b}]_d \end{bmatrix} + \begin{bmatrix} [\mathbf{R}_1]_d \\ \vdots \\ [\mathbf{R}_m]_d \end{bmatrix} = \begin{bmatrix} [\mathbf{b} + \mathbf{R}_1]_d \\ \vdots \\ [\mathbf{b} + \mathbf{R}_m]_d \end{bmatrix}$$

Proving that each secret is bounded by B (and thus lies between $-B$ and B) reduces to proving that all the secrets that are pack-shared by each $[\mathbf{b} - \mathbf{R}_i]_d$ and $[\mathbf{b} + \mathbf{R}_i]_d$ for $i \in [m]$, are positive. We thus show a subroutine, described in Figure 3 that allows a dealer to prove that secrets that are pack-shared among players are positive. To do this, we follow an idea of Boudot [9] and use Lagrange’s Four-Square Theorem, which states that *every positive number can be written as the sum of four squares* (see e.g. [16]). Moreover, these four squares can be efficiently computed [28,24]. Suppose the dealer has pack-shared a secret vector $\mathbf{z} \in \mathbb{Z}_q^{1 \times k}$. For each coordinate z_j for $j \in [k]$, the dealer finds the four numbers u_j, v_j, w_j, y_j such that $z_j = u_j^2 + v_j^2 + w_j^2 + y_j^2$. We let $\tilde{\mathbf{u}}, \tilde{\mathbf{v}}, \tilde{\mathbf{w}}, \tilde{\mathbf{y}}$ be the vectors with u_j, v_j, w_j, y_j as the j th coordinate, respectively. The dealer $(d/2)$ -shares each of these vectors $[\tilde{\mathbf{u}}]_{d/2}, [\tilde{\mathbf{v}}]_{d/2}, [\tilde{\mathbf{w}}]_{d/2}, [\tilde{\mathbf{y}}]_{d/2}$. Similarly, we let $\mathbf{u}, \mathbf{v}, \mathbf{w}, \mathbf{y}$ be the vectors with $u_j^2, v_j^2, w_j^2, y_j^2$ as the j th coordinate, respectively. Players can locally compute sharings $[\mathbf{u}]_d, [\mathbf{v}]_d, [\mathbf{w}]_d, [\mathbf{y}]_d$ by squaring their corresponding shares of $[\tilde{\mathbf{u}}]_{d/2}, [\tilde{\mathbf{v}}]_{d/2}, [\tilde{\mathbf{w}}]_{d/2}, [\tilde{\mathbf{y}}]_{d/2}$. Each player then computes,

$$[\mathbf{z}]_d - [\mathbf{u}]_d - [\mathbf{v}]_d - [\mathbf{w}]_d - [\mathbf{y}]_d = [\mathbf{z} - \mathbf{u} - \mathbf{v} - \mathbf{w} - \mathbf{y}]_d = [\mathbf{0}]_d$$

and together they check that the result is indeed a pack-sharing of the vector $\mathbf{0} \in \mathbb{Z}^k$.

However, suppose that a cheating dealer chooses $|u_j| > \sqrt{q/2}$. Then $|u_j^2| > q$ and we have wrap-around modulo q , which means that the cheating dealer could convince the players that a secret z_j is positive, without this being true. To ensure this does not happen, we have the dealer prove that each of u_j, v_j, w_j, y_j is bounded by some bound B' , which although larger than B , is certainly much smaller than $\sqrt{q/2}$ (in fact, we will need $B' < \sqrt{q/8}$ so that we don’t have overflow when adding the four squares).

Our protocol for verifying that numbers are bounded by B' uses techniques from Cramer and Damgård [13]. Players check τ shares at a time, where τ should be thought of as the “local security parameter” for the protocol π_{VERBND} . The players compute a linear combination of their shares (with some noise added) and reconstruct the result, such that if the secrets resulting from this reconstruction are “not too big” then the original secrets (i.e. the entries in \mathbf{R}) are also small. To ensure that the reconstructed result does not reveal \mathbf{R} , we let the added noise be in an interval that is a factor of 2^τ larger than the entries in \mathbf{R} . To guarantee that π_{VERBND} has statistical (in λ) t -privacy, we set $\tau = \omega(\log \lambda)$. The final bound that we are able to prove is $B' = 2^{2\tau+1}mB$. We will thus need to ensure that $\sqrt{q/8} > 2^{2\tau+1}mB$.

We give full descriptions of the protocol π_{VERSM} in Figure 2, of the subroutine to verify that secrets are positive in Figure 3, and the subroutine to verify that numbers are bounded by B' in Figure 4.

Protocol π_{VERSM} between parties $(\mathcal{P}_1, \dots, \mathcal{P}_N)$ and input client \mathcal{I} .Common input: bound B Input to \mathcal{I} : $\mathbf{R} \in \mathbb{Z}_q^{m \times k}$.Input to \mathcal{P}_i : Corresponding shares of $[\mathbf{R}_1]_d, \dots, [\mathbf{R}_m]_d$.

1. \mathcal{I} prepares a d -sharing of $\mathbf{b} = (B, \dots, B) \in \mathbb{Z}_q^k$: $[\mathbf{b}]_d$. \mathcal{I} gives each player its corresponding shares.
2. Players run the subroutine π_{VERPOS} (see Figure 3) with

$$\begin{bmatrix} [\mathbf{b}]_d \\ \vdots \\ [\mathbf{b}]_d \end{bmatrix} - \begin{bmatrix} [\mathbf{R}_1]_d \\ \vdots \\ [\mathbf{R}_m]_d \end{bmatrix} = \begin{bmatrix} [\mathbf{b} - \mathbf{R}_1]_d \\ \vdots \\ [\mathbf{b} - \mathbf{R}_m]_d \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} [\mathbf{b}]_d \\ \vdots \\ [\mathbf{b}]_d \end{bmatrix} + \begin{bmatrix} [\mathbf{R}_1]_d \\ \vdots \\ [\mathbf{R}_m]_d \end{bmatrix} = \begin{bmatrix} [\mathbf{b} + \mathbf{R}_1]_d \\ \vdots \\ [\mathbf{b} + \mathbf{R}_m]_d \end{bmatrix}$$

Fig. 2. Protocol π_{VERSM} to verify that secrets are numerically small**Subroutine π_{VERPOS} between parties $(\mathcal{P}_1, \dots, \mathcal{P}_N)$ and input client \mathcal{I} , to verify that secrets are positive.**Input to \mathcal{I} : $\mathbf{Z} \in \mathbb{Z}_q^{m \times k}$.Input to \mathcal{P}_i : Corresponding shares of $[\mathbf{Z}_1]_d, \dots, [\mathbf{Z}_m]_d$.

1. For each entry $z_i^{(j)}$ of \mathbf{Z} (for $i \in [m], j \in [k]$), the dealer finds the four numbers $u_{ij}, v_{ij}, w_{ij}, y_{ij}$ such that $z_i^{(j)} = u_{ij}^2 + v_{ij}^2 + w_{ij}^2 + y_{ij}^2$. Define $\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{W}}, \tilde{\mathbf{Y}}$ to be the matrices with $u_{ij}, v_{ij}, w_{ij}, y_{ij}$ as the (i, j) th entry, respectively. Similarly, define $\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Y}$ to be the matrices with $u_{ij}^2, v_{ij}^2, w_{ij}^2, y_{ij}^2$ as the (i, j) th entry, respectively.
2. \mathcal{I} computes and distributes $(d/2)$ -sharings of the rows of $\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{W}}, \tilde{\mathbf{Y}}$: $[\tilde{\mathbf{U}}_i]_{d/2}, [\tilde{\mathbf{V}}_i]_{d/2}, [\tilde{\mathbf{W}}_i]_{d/2}, [\tilde{\mathbf{Y}}_i]_{d/2}$, for $i \in [m]$.
3. Players run protocol π_{CHECK} from Section 2.4 with the shares $[\tilde{\mathbf{U}}_i]_{d/2}, [\tilde{\mathbf{V}}_i]_{d/2}, [\tilde{\mathbf{W}}_i]_{d/2}, [\tilde{\mathbf{Y}}_i]_{d/2}$, for $i \in [m]$ (a total of $4m/(N-t)$ times) to verify that these shares are $d/2$ -consistent.
4. \mathcal{I} and the players run the subroutine π_{VERBND} (see Figure 4) with the shares $[\tilde{\mathbf{U}}_i]_{d/2}, [\tilde{\mathbf{V}}_i]_{d/2}, [\tilde{\mathbf{W}}_i]_{d/2}, [\tilde{\mathbf{Y}}_i]_{d/2}$, for $i \in [m]$ (a total of $4m/\tau$ times), to verify that each of the $u_{ij}, v_{ij}, w_{ij}, y_{ij}$ is bounded by $B' < \sqrt{q/8}$.
5. For each row $i \in [m]$, players locally compute d -sharings $[\mathbf{U}_i]_d, [\mathbf{V}_i]_d, [\mathbf{W}_i]_d, [\mathbf{Y}_i]_d$ by squaring their corresponding shares of $[\tilde{\mathbf{U}}_i]_{d/2}, [\tilde{\mathbf{V}}_i]_{d/2}, [\tilde{\mathbf{W}}_i]_{d/2}, [\tilde{\mathbf{Y}}_i]_{d/2}$.
6. For each row $i \in [m]$, players locally compute

$$[\mathbf{Z}_i]_d - [\mathbf{U}_i]_d - [\mathbf{V}_i]_d - [\mathbf{W}_i]_d - [\mathbf{Y}_i]_d = [\mathbf{Z}_i - \mathbf{U}_i - \mathbf{V}_i - \mathbf{W}_i - \mathbf{Y}_i]_d$$

and check that the result is a pack-sharing of the vector $\mathbf{0} \in \mathbb{Z}^{1 \times k}$.**Fig. 3.** Subroutine π_{VERPOS} to verify that secrets are positive

Subroutine π_{VERBND} between parties $(\mathcal{P}_1, \dots, \mathcal{P}_N)$ and input client \mathcal{I} , to verify that numbers are bounded by $B' = 2^{2\tau+1}mB$.

Common input: bound B

Input to \mathcal{I} : $\mathbf{Z}' \in \mathbb{Z}_q^{\tau \times k}$.

Input to \mathcal{P}_i : Corresponding shares of $[\mathbf{Z}'_1]_d, \dots, [\mathbf{Z}'_\tau]_d$ (that are known to be d -consistent).

1. \mathcal{I} chooses $\mathbf{X} \leftarrow [-2^\tau mB, 2^\tau mB]^{(2\tau-1) \times k}$, and prepares d -sharings of the rows of \mathbf{X} : $[\mathbf{X}_1]_d, \dots, [\mathbf{X}_{2\tau-1}]_d$. \mathcal{I} gives each player its corresponding shares.
2. Players $\mathcal{P}_1, \dots, \mathcal{P}_N$ coin-flip for a random vector $\mathbf{e} \in \{0, 1\}^{\tau \times 1}$.
3. Define matrix $\mathbf{M}_\mathbf{e}$ to be the $(2\tau - 1) \times \tau$ matrix with its (i, j) -th entry defined by $m_{\mathbf{e},i}^{(j)} = e_{i-j+1}$ for $1 \leq i - j + 1 \leq \lambda$. Each player locally computes

$$\begin{bmatrix} [(\mathbf{M}_\mathbf{e}\mathbf{Z}')_1]_{d'} \\ \vdots \\ [(\mathbf{M}_\mathbf{e}\mathbf{Z}')_{2\tau-1}]_{d'} \end{bmatrix} + \begin{bmatrix} [\mathbf{X}_1]_{d'} \\ \vdots \\ [\mathbf{X}_{2\tau-1}]_{d'} \end{bmatrix} = \begin{bmatrix} [(\mathbf{M}_\mathbf{e}\mathbf{Z}' + \mathbf{X})_1]_{d'} \\ \vdots \\ [(\mathbf{M}_\mathbf{e}\mathbf{Z}' + \mathbf{X})_{2\tau-1}]_{d'} \end{bmatrix}$$

4. Players reconstruct $\mathbf{M}_\mathbf{e}\mathbf{Z}' + \mathbf{X}$ row by row and check that all its entries are bounded by $2^{2\tau+1}mB$.

Fig. 4. Subroutine π_{VERBND} to verify that numbers are bounded by $B' = 2^{2\tau+1}mB < \sqrt{q}/8$

We set $N = \Theta(t)$ as is required for the IKOS construction and for privacy (see Section 2.3), and analyze the communication complexity of the π_{VERSM} protocol. Each share has size at most $\log q$. Each execution of π_{VERBND} has communication cost $O(\tau N \log q)$: sharing the \mathbf{X}_i 's has communication cost $(2\tau - 1)N \log q$, the coin-flipping of \mathbf{e} has communication cost τ since we'll use this MPC protocol inside the IKOS construction, and reconstructing $\mathbf{M}_\mathbf{e}\mathbf{Z}' + \mathbf{X}$ has communication cost $(2\tau - 1)N \log q$. The subroutine π_{VERPOS} (Figure 3) has communication complexity $O(mN \log q)$: sharing of the rows of $\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Y}$ has cost $4mN \log q$, the total cost of running π_{CHECK} is $(N(N + t) \log q) \cdot 4m/(N - 2t) = O(mN \log q)$, the total cost of running π_{VERBND} is $O(\tau N \log q) \cdot 4m/\tau = O(mN \log q)$, and the final reconstruction has cost $mN \log q$. Finally, the communication complexity of protocol π_{VERSM} is $O(mN \log q)$: sharing \mathbf{b} has communication cost $N \log q$, and we run the subroutine π_{VERPOS} twice.

Lemma 4. *Let n, m, r, q, N, t, k be as in Theorem 2 and Section 2.3, and let B be some publicly-known bound. If $\tau = \omega(\log \lambda)$ and $\sqrt{q}/8 > 2^{2\tau+1}mB$ then the protocol π_{VERSM} described in Figure 2 allows N players to verify, with negligible error probability in λ , that all entries in a secret matrix $\mathbf{R} \in \mathbb{Z}_q^{m \times k}$ are bounded by B . It has statistical t -privacy in the presence of a semi-honest adversary, perfect t -robustness in the presence of a malicious adversary, and communication complexity $O(mN \log q)$.*

4 Proofs of Plaintext Knowledge

We wish to show a zero-knowledge proof protocol that allows a prover to prove that he knows the plaintexts of k different ciphertexts, each encrypted under the same public key. We show how to do this for the Regev encryption scheme described in Section 2.5. More formally, we show a zero-knowledge proof protocol for the following relation:

$$\begin{aligned}
 R_{\text{PoPK}} = \{ (x, w) \mid & x = ((\mathbf{A}, \mathbf{b}), (\mathbf{u}^{(1)}, c^{(1)}), \dots, (\mathbf{u}^{(k)}, c^{(k)})), \\
 & w = ((m^{(1)}, \mathbf{r}^{(1)}), \dots, (m^{(k)}, \mathbf{r}^{(k)})) \text{ s.t.} \\
 & \forall j \in [k] : (\mathbf{u}^{(j)}, c^{(j)}) = \text{Enc}_{(\mathbf{A}, \mathbf{b})}(m^{(j)}; \mathbf{r}^{(j)}) \\
 & \text{and } |m^{(j)}| \leq p/2, \|\mathbf{r}^{(j)}\|_\infty < r \}
 \end{aligned}$$

We create protocol Π_{PoPK} for relation R_{PoPK} using the ‘‘MPC-in-the-head’’ technique of [22] described in Section 2.2. We let f_{PoPK} be the N -party functionality that takes the entire input w from \mathcal{I} and outputs $R_{\text{PoPK}}(x, w)$ to all N players. In Figure 5, we show our construction of a t -robust and t -private N -party protocol, π_{PoPK} , realizing functionality f_{PoPK} . The idea is to have \mathcal{I} pack secret-share the messages, as well as pack secret-share each coordinate of the randomness vectors. The players then locally run the encryption algorithm on their shares, reconstruct the resulting shares, and check that the reconstructed secrets are indeed the claimed ciphertexts. The input client \mathcal{I} also needs to prove that the messages and randomness come from the correct spaces. For example, he would need to show that the magnitude of each message is less than $p/2$ (since the message space is \mathbb{Z}_p), and that each coordinate of each randomness vector is at most r (since we are using the truncated Gaussian distribution described in Section 2.5). For this, we will use the protocol π_{VERSM} described in Section 3.

We set $t = \Theta(k)$ and $N = \Theta(t)$ as is required for the IKOS construction and for privacy (see Section 2.3), and analyze the communication complexity of our protocol π_{PoPK} (see Figure 5). Since each share has size $\log q$, step 1 has communication cost $(m + 1)N \log q = O(mk \log q)$. We run π_{CHECK} $m + 1/(N - 2t) = O(m/k)$ times, so step 2 has communication cost $N(N + t) \log q(m/k) = O(mk \log q)$. The reconstruction in step 3 has cost $2nN \log q$ and running protocol π_{VERSM} has cost $2mN \log q$ so the total cost of step 3 and of π_{PoPK} is $O(mk \log q)$.

Our techniques are similar to those of Bendlin and Damgård [7]. However, our protocol π_{VERSM} for proving that a secret is small (see Section 3) allows us to prove soundness for message space \mathbb{Z}_p and randomness sampled from the discrete Gaussian, whereas the construction of [7] only worked for bit messages and bit-vector randomness. Finally, our use of packed secret sharing allows us to achieve a better amortized communication complexity. The protocol of [7] has complexity $O(nm \log q)$ per proof, whereas we achieve an amortized complexity of $O(m \log q)$ per proof.

Lemma 5. *Let n, m, r, p, q, N, t, k be as in Lemma 4 with $B = \max(p/2, r)$. The protocol π_{PoPK} described in Figure 5 realizes f_{PoPK} with statistical t -privacy in the presence of a semi-honest adversary and perfect t -robustness in the presence of a malicious adversary, and has communication complexity $O(mk \log q)$.*

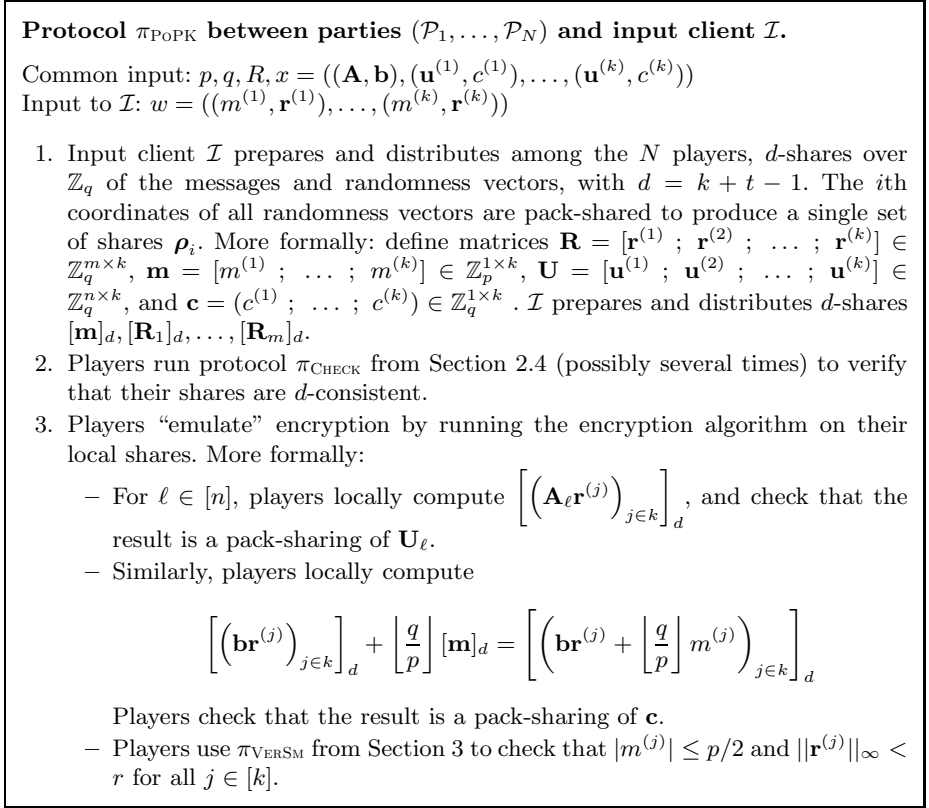


Fig. 5. MPC protocol π_{PoPK} that realizes f_{PoPK}

Putting together Lemma 5 with Theorem 1 yields the following theorem.

Theorem 3. *Let n, m, r, p, q be as in Lemma 4 with $B = \max(p/2, r)$. Given an unconditionally-binding commitment scheme, it is possible to construct a computational zero-knowledge proof protocol Π_{PoPK} for relation R_{PoPK} with negligible (in λ) soundness error and amortized communication complexity $O(m \log q)$ per proof.*

5 Proofs of Correct Multiplication

In this section we show proofs for correct multiplication for the Regev encryption scheme. In our protocol, the prover performs k proofs at a time, all under the same public key. More formally, we give a zero-knowledge proof protocol for the following relation:

Protocol π_{POCM} between parties $(\mathcal{P}_1, \dots, \mathcal{P}_N)$ and input client \mathcal{I} .

Common input: $p, q, R, x = ((\mathbf{A}, \mathbf{b}), (\mathbf{u}^{(1)}, c^{(1)}, \mathbf{v}^{(1)}, e^{(1)}), \dots, (\mathbf{u}^{(k)}, c^{(k)}, \mathbf{v}^{(k)}, e^{(k)}))$
 Input to \mathcal{I} : $w = ((m^{(1)}, \mathbf{r}^{(1)}, x^{(1)}), \dots, (m^{(k)}, \mathbf{r}^{(k)}, x^{(k)}))$

1. Input client \mathcal{I} prepares and distributes among the N players, d -shares over \mathbb{Z}_q of the messages and randomness vectors, with $d = k + t - 1$. The i th coordinates of all randomness vectors are packed shared to produce a single set of shares.. More formally: define matrices $\mathbf{R} = [\mathbf{r}^{(1)} ; \mathbf{r}^{(2)} ; \dots ; \mathbf{r}^{(k)}] \in \mathbb{Z}_q^{m \times k}$, $\mathbf{m} = [m^{(1)} ; \dots ; m^{(k)}] \in \mathbb{Z}_p^{1 \times k}$, $\mathbf{x} = [x^{(1)} ; \dots ; x^{(k)}] \in \mathbb{Z}_q^{1 \times k}$, $\mathbf{U} = [\mathbf{u}^{(1)} ; \mathbf{u}^{(2)} ; \dots ; \mathbf{u}^{(k)}] \in \mathbb{Z}_q^{n \times k}$, $\mathbf{c} = (c^{(1)}, \dots, c^{(k)}) \in \mathbb{Z}_q^{1 \times k}$, $\mathbf{V} = [\mathbf{v}^{(1)} ; \mathbf{v}^{(2)} ; \dots ; \mathbf{v}^{(k)}] \in \mathbb{Z}_q^{n \times k}$, and $\mathbf{e} = (e^{(1)} ; \dots ; e^{(k)}) \in \mathbb{Z}_q^{1 \times k}$. \mathcal{I} prepares and distributes $(d/2)$ -share $[\mathbf{m}]_d$ and d -shares $[\mathbf{x}]_d, [\mathbf{R}_1]_d, \dots, [\mathbf{R}_m]_d$. \mathcal{I} also prepares and *broadcasts* $(d/2)$ -shares $[\mathbf{c}]_{d/2}, [\mathbf{U}_1]_{d/2}, \dots, [\mathbf{U}_n]_{d/2}$.
2. Players run protocol π_{CHECK} from Section 2.4 (possibly several times) to verify that shares $[\mathbf{x}]_d, [\mathbf{R}_1]_d, \dots, [\mathbf{R}_m]_d$ are d -consistent, and share $[\mathbf{m}]_d$ is $(d/2)$ -consistent. They also check locally that $\mathbf{c}, \mathbf{U}_1, \dots, \mathbf{U}_m$ are correctly shared.
3. Players “emulate” correct computation of each $(\mathbf{v}^{(i)}, c^{(i)})$. More formally:
 - For $\ell \in [n]$, players locally compute $\left[\left(\mathbf{A}_\ell \mathbf{r}^{(j)} \right)_{j \in [k]} \right]_d$. They also locally compute $\left[\left(\mathbf{b} \mathbf{r}^{(j)} + \left\lfloor \frac{q}{p} \right\rfloor x^{(j)} \right)_{j \in [k]} \right]_d$.
 - For $\ell \in [n]$, players locally compute

$$[\mathbf{m}]_{d/2} [\mathbf{U}_\ell]_{d/2} + \left[\left(\mathbf{A}_\ell \mathbf{r}^{(j)} \right)_{j \in [k]} \right]_d = \left[\left(u_\ell^{(j)} m^{(j)} + \mathbf{A}_\ell \mathbf{r}^{(j)} \right)_{j \in [k]} \right]_d$$

Players check that the result is a pack-sharing of \mathbf{V}_ℓ .

- Players locally compute

$$\begin{aligned} & [\mathbf{m}]_{d/2} [\mathbf{c}]_{d/2} + \left[\left(\mathbf{b} \mathbf{r}^{(j)} + \left\lfloor \frac{q}{p} \right\rfloor x^{(j)} \right)_{j \in [k]} \right]_d \\ &= \left[\left(c^{(j)} m^{(j)} + \mathbf{b} \mathbf{r}^{(j)} + \left\lfloor \frac{q}{p} \right\rfloor x^{(j)} \right)_{j \in [k]} \right]_d \end{aligned}$$

Players check that the result is a pack-sharing of \mathbf{e} .

- Players use π_{VERSM} from Section 3 to check that $|m^{(j)}| \leq p/2$, $|x^{(j)}| \leq p/2$ and $\|\mathbf{r}^{(j)}\|_\infty < r$ for all $j \in [k]$.

Fig. 6. MPC protocol π_{POCM} that realizes f_{POCM}

$$\begin{aligned} R_{\text{POCM}} = \{ (x, w) \mid & x = ((\mathbf{A}, \mathbf{b}), (\mathbf{u}^{(1)}, c^{(1)}, \mathbf{v}^{(1)}, e^{(1)}), \dots, (\mathbf{u}^{(k)}, c^{(k)}, \mathbf{v}^{(k)}, e^{(k)})), \\ & w = ((m^{(1)}, \mathbf{r}^{(1)}, x^{(1)}), \dots, (m^{(k)}, \mathbf{r}^{(k)}, x^{(k)})) \text{ s.t.} \\ & \forall j \in [k]: (\mathbf{v}^{(j)}, e^{(j)}) = m^{(j)} (\mathbf{u}^{(j)}, c^{(j)}) + \text{Enc}_{(\mathbf{A}, \mathbf{b})}(x^{(j)}; \mathbf{r}^{(j)}) \\ & \text{and } |m^{(j)}| \leq p/2, |x^{(j)}| \leq p/2, \|\mathbf{r}^{(j)}\|_\infty < r \} \end{aligned}$$

As in Section 4, we create protocol Π_{PoCM} for relation R_{PoCM} using the “MPC-in-the-head” technique of [22], described in Section 2.2. We let f_{PoCM} be the N -party functionality that takes the entire input w from \mathcal{I} and outputs $R_{\text{PoCM}}(x, w)$ to all N players. In Figure 6, we show our construction of a t -robust and t -private N -party protocol, π_{PoCM} , realizing functionality f_{PoCM} . Again, the idea is to have \mathcal{I} pack secret-share the messages, as well as pack secret-share each coordinate of the randomness vectors. The players then locally emulate the encryption of the random message and perform the multiplication, then reconstruct the resulting shares, and check that the reconstructed secrets are indeed the claimed ciphertexts. As before, the input client \mathcal{I} also needs to prove that the messages and randomness come from the correct spaces. We again use the protocol π_{VERSM} described in Section 3 for this purpose.

We set $t = \theta(k)$ and $N = \theta(t)$ as is required for the IKOS construction and for privacy (see Section 2.3), and analyze the communication complexity of π_{PoCM} described in Figure 6. Since each share has size $\log q$, step 1 has communication cost $2(m+1)N \log q = O(mk \log q)$. We run π_{CHECK} $m + 1/(N - 2t) = O(m/k)$ times, so step 2 has communication cost $N(N + t) \log q(m/k) = O(mk \log q)$. The reconstruction in step 3 has cost $2nN \log q$ and running protocol π_{VERSM} has cost $2mN \log q$ so the total cost of step 3 and of π_{PoPK} is $O(mk \log q)$.

Lemma 6. *Let n, m, r, p, q, N, t, k be as in Lemma 4 with $B = \max(p/2, r)$. The protocol π_{PoCM} described in Figure 6 realizes f_{PoCM} with statistical t -privacy in the presence of a semi-honest adversary and perfect t -robustness in the presence of a malicious adversary, and has communication complexity $O(mk \log q)$.*

Putting Lemma 6 together with Theorem 1 yields the following theorem.

Theorem 4. *Let n, m, r, p, q be as in Lemma 4 with $B = \max(p/2, r)$. Given an unconditionally-binding commitment scheme, it is possible to construct a computational zero-knowledge proof protocol Π_{PoCM} for relation R_{PoCM} with negligible (in λ) soundness error and amortized communication complexity $O(m \log q)$ per proof.*

References

1. Asharov, G., Jain, A., López-Alt, A., Tromer, E., Vaikuntanathan, V., Wichs, D.: Multiparty Computation with Low Communication, Computation and Interaction via Threshold FHE. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 483–501. Springer, Heidelberg (2012)
2. Asharov, G., Jain, A., Wichs, D.: Multiparty computation with low communication, computation and interaction via threshold fhe. Cryptology ePrint Archive: Report 2011/613 (2011)
3. Beaver, D.: Efficient Multiparty Protocols Using Circuit Randomization. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 420–432. Springer, Heidelberg (1992)
4. Beerliová-Trubíniová, Z., Hirt, M.: Perfectly-Secure MPC with Linear Communication Complexity. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 213–230. Springer, Heidelberg (2008)

5. Beerliová-Trubíniová, Z., Hirt, M., Nielsen, J.B.: On the theoretical gap between synchronous and asynchronous mpc protocols. In: Richa, A.W., Guerraoui, R. (eds.) PODC, pp. 211–218. ACM (2010)
6. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: STOC, pp. 1–10 (1988)
7. Bendlin, R., Damgård, I.: Threshold Decryption and Zero-Knowledge Proofs for Lattice-Based Cryptosystems. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 201–218. Springer, Heidelberg (2010)
8. Bendlin, R., Damgård, I., Orlandi, C., Zakarias, S.: Semi-homomorphic Encryption and Multiparty Computation. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 169–188. Springer, Heidelberg (2011)
9. Boudot, F.: Efficient Proofs that a Committed Number Lies in an Interval. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 431–444. Springer, Heidelberg (2000)
10. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: Goldwasser, S. (ed.) ITCS, pp. 309–325. ACM (2012)
11. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) lwe. In: Ostrovsky, R. (ed.) FOCS, pp. 97–106. IEEE (2011)
12. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols (extended abstract). In: STOC, pp. 11–19 (1988)
13. Cramer, R., Damgård, I.: On the Amortized Complexity of Zero-Knowledge Protocols. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 177–191. Springer, Heidelberg (2009)
14. Damgård, I., Orlandi, C.: Multiparty Computation for Dishonest Majority: From Passive to Active Security at Low Cost. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 558–576. Springer, Heidelberg (2010)
15. Damgård, I., Pastro, V., Smart, N.P., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. IACR Cryptology ePrint Archive, 2011:535 (2011)
16. Fine, B., Rosenberger, G.: Number Theory: An Introduction via the Distribution of Primes. Birkhäuser (2006)
17. Franklin, M.K., Yung, M.: Communication complexity of secure computation (extended abstract). In: STOC, pp. 699–710. ACM (1992)
18. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Dwork, C. (ed.) STOC, pp. 197–206. ACM (2008)
19. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or a completeness theorem for protocols with honest majority. In: STOC, pp. 218–229 (1987)
20. Hirt, M., Maurer, U.M.: Robustness for Free in Unconditional Multi-party Computation. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 101–118. Springer, Heidelberg (2001)
21. Hirt, M., Nielsen, J.B., Przydatek, B.: Asynchronous Multi-Party Computation with Quadratic Communication. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 473–485. Springer, Heidelberg (2008)
22. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: Johnson, D.S., Feige, U. (eds.) STOC, pp. 21–30. ACM (2007)

23. Katz, J., Koo, C.-Y.: Round-Efficient Secure Computation in Point-to-Point Networks. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 311–328. Springer, Heidelberg (2007)
24. Lipmaa, H.: On Diophantine Complexity and Statistical Zero-Knowledge Arguments. In: Laih, C.-S. (ed.) ASIACRYPT 2003. LNCS, vol. 2894, pp. 398–415. Springer, Heidelberg (2003)
25. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.* 37(1), 267–302 (2007)
26. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Mitzenmacher, M. (ed.) STOC, pp. 333–342. ACM (2009)
27. Peikert, C., Vaikuntanathan, V., Waters, B.: A Framework for Efficient and Composable Oblivious Transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008)
28. Rabin, M.O., Shallit, J.O.: Randomized algorithms in number theory. *Communications on Pure and Applied Mathematics* 39(S1), S239–S259 (1986)
29. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) STOC, pp. 84–93. ACM (2005)
30. Shamir, A.: How to share a secret. *Commun. ACM* 22(11), 612–613 (1979)
31. Yao, A.C.-C.: Protocols for secure computations (extended abstract). In: FOCS, pp. 160–164 (1982)