Threshold Decryption and Zero-Knowledge Proofs for Lattice-Based Cryptosystems

Rikke Bendlin and Ivan Damgård

Department of Computer Science, Aarhus University
{rikkeb, ivan}@cs.au.dk

Abstract. We present a variant of Regev's cryptosystem first presented in [Reg05], but with a new choice of parameters. By a recent classical reduction by Peikert we prove the scheme semantically secure based on the worst-case lattice problem GapSVP. From this we construct a threshold cryptosystem which has a very efficient and non-interactive decryption protocol. We prove the threshold cryptosystem secure against passive adversaries corrupting all but one of the players, and againts active adversaries corrupting less than one third of the players. We also describe how one can build a distributed key generation protocol. In the final part of the paper, we show how one can, in zero-knowledge - prove knowledge of the plaintext contained in a given ciphertext from Regev's original cryptosystem or our variant. The proof is of size only a constant times the size of a ciphertext.

1 Introduction

Cryptography based on lattice problems is one of the most important examples of techniques holding promise for public-key cryptography that is secure even under quantum attacks and are also interesting in that they can be based on worst-case complexity assumptions. Recently, these techniques have become much more efficient after it has been realized that one can base the actual cryptosystem on the learning with error problem (LWE), and then argue that the (variant of the) LWE problem used is as hard as some lattice related problem, typically computing the shortest vector in a lattice. In the LWE problem, the adversary must compute a secret vector s with entries in some field or ring, given only the inner product of s with some public vectors where, however, some noise has been added to the products. As mentioned, basing a cryptosystem on LWE can lead to quite efficent cryptosystems, see, e.g., [Reg05],[PVW08],[MR08],[Pei09].

As lattice-based cryptography moves closer to practice, it becomes an important research question to investigate whether these cryptosystems can provide the same "extra" functionality we have come to expect from well-known public-key cryptosystems based on factoring or discrete logarithms. For instance, can we have threshold versions of these systems? In other words, we want to share the private key among a set of servers and efficently decrypt a ciphertext while revealing nothing but the plaintext to the adversary. And furthermore, can one prove, in zero-knowledge and efficiently, knowledge of the plaintext contained in a given ciphertext?

In this paper we construct such a threshold cryptosystem, based on a variant of Regev's system [Reg05]. We show our scheme semantically secure based on a worst-case lattice problem using a recent reduction of Peikert[PVW08]. To the best of our knowledge, it is the first lattice-based threshold cryptosystem. We need to use a larger modulus than Regev, thus making ciphertexts larger, on the other hand we get a very efficient and non-interactive decryption protocol: each player needs only to do local computation and announce a single element from the underlying ring. The basic version of the protocol is secure against a passive adversary corrupting all but one of the players. For a small number of players, we show an equally efficent version secure against a malicious adversary corrupting less than a third of the players. We also describe a distributed protocol for generating keys.

Various improvements of Regev's original cryptosystem have been made since its first appearence, e.g. in [PVW08] and [MR08]. Our threshold cryptosystem can be generalized in the same way, but we stick to Regev's original approach here for simplicity.

In the final part of the paper, we present a zero-knowledge protocol for proving knowledge of the plaintext contained in a given ciphertext, for Regev's original cryptosystem as well as our variant. The proof is much more efficient than what generic methods would give us: it has size only a constant times the size of the ciphertext, and the computation required is comparable to what is required to encrypt and decrypt. The protocol is based on the construction from [IKOS07] of zero-knowledge from multiparty computation protocols. Whereas this paradigm has perhaps been perceived primarily as a theoretical tool, we show here that it can also be highly relevant in practice.

2 Preliminaries

When writing $x \in_R S$ we mean that x is chosen uniformly at random from the set S. Equivalently $x \in_\chi S$ means choosing x from the set S according to the distribution χ.

Given a probability distribution χ on $\mathbb{Z}_q$, let n be some integer and $\mathbf{s} \in \mathbb{Z}_q^n$. We define $A_{\mathbf{s},\chi}$ as the distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ obtained by choosing $\mathbf{a} \in_R \mathbb{Z}_q^n$, $e \in_\chi \mathbb{Z}_q$ and outputting $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$. We define the decisional Learning With Errors (LWE) problem as being able to distinguish between a sample from $A_{s,\chi}$ and the uniform distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ with non-negligible probability. We define the search LWE problem as given a sample from $A_{s,\chi}$ finding s with non-negligible probability.

By $\overline{\Psi}_\alpha$ we denote a discrete Gaussian distribution on $\mathbb{Z}_q$ with mean 0 and standard deviation $\frac{q\alpha}{\sqrt{2\pi}}$. Likewise Ψ_α is a continuous Gaussian distribution on $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ with mean 0 and standard deviation $\frac{\alpha}{\sqrt{2\pi}}$. By χ^{*k} we denote the distribution given by summing k independent samples from χ. Note in particular that when $\chi = \overline{\Psi}_\alpha$ we have that $\chi^{*k} = \overline{\Psi}_{\sqrt{k}\cdot\alpha}$. This follows immediately from $\overline{\Psi}_\alpha$ behaving as an ordinary normal distribution.

3 Cryptosystem

We first present the underlying cryptosystem which was proposed first in [Reg05], but with a new choice of parameters better suited for the distributed decryption protocol given later.

3.1 Description

Let n be the security parameter of the cryptosystem. Then the main parameter is an integer q which is chosen as $q = 2^{O(n)}$. More specifically q will not be a prime but a B-smooth number where B is of polynomial size. That is $q = \prod p_i$ is a product of prime numbers $p_1, \ldots, p_k$, where $p_i < B$ and also $p_i > u$, the number of players in the distributed decryption protocol. The latter requirement on the primes is necessary in order to do secret sharing over the the ring $\mathbb{Z}_q$, more on this later. We also need an integer m which will be chosen to be $O(n^3)$. Finally, we need a distribution χ on $\mathbb{Z}_q$ which will be taken to be the discrete Gaussian distribution $\overline{\Psi}_\alpha$, where $\alpha = q^\beta$ for some $\beta < 1/4$.

The cryptosystem is now defined as follows:

- **Secret key:** Choose $\mathbf{s} \in_R \mathbb{Z}_q^n$. The secret key is then $\mathbf{s}$.
- **Public key:** Choose m vectors $\mathbf{a}_1, \ldots, \mathbf{a}_m \in_R \mathbb{Z}_q^n$, m elements $e_1, \ldots, e_m \in_\chi \mathbb{Z}_q$. The public key is then given by $(\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s}\rangle + e_i)_{i=1}^m$.
- **Encryption:** Choose a random set S among all the subsets of $[m]$. Given a bit γ, the encryption of γ is given by $(\sum_{i\in S} \mathbf{a}_i, \gamma \cdot \lfloor \frac{q}{2} \rfloor + \sum_{i\in S} b_i)$.
- **Decryption:** Given a ciphertext $(\mathbf{a}, b)$, calculate $b - \langle \mathbf{a}, \mathbf{s}\rangle$ and determine whether it is closer to 0, the encrypted bit being 0, or closer to $\frac{q}{2}$, the encrypted bit being 1.

3.2 Correctness

The correctness of the decryption protocol is given in the following theorem.

Theorem 1 (Correctness). *If for any $k \in \{0, 1, \ldots, m\}$ it holds that*

$$\Pr_{e\sim\chi^{*k}}(|e| \geq \sqrt[3]{q}) \leq 2^{-O(n)}$$

then the decryption protocol will give correct output except with negligible probability.

A similar theorem is proved in [Reg05] for Regev's original choice of parameters. The intuition is clear, if the noise added is not too big, we will be able to decrypt to the right bit. The correctness with the new parameters follows from the following claim.

Claim (Correctness). For the choice of parameters made, for any $k \in \{0, 1, \ldots, m\}$ and $e \sim \chi^{*k}$ it holds that

$$\Pr_{e \sim \chi^{*k}}(|e| \geq \sqrt[3]{q}) \leq 2^{-O(n)}$$

Proof. We will prove this using the Chebyshev inequality, but first we will reduce the problem from $\overline{\Psi}_\alpha$ to Ψ_α. Given $e \sim \overline{\Psi}_\alpha^{*k}$ we have that $e = \sum_{i=1}^{k} \lfloor qx_i \rceil \pmod q$, where $x_i \sim \Psi_\alpha$. The value of e is at most $k < m < \sqrt[3]{q}/2$ away from $\sum_{i=1}^{k} qx_i \pmod q$, so it is sufficient to prove that $|\sum_{i=1}^{k} qx_i \pmod q| < \sqrt[3]{q}/2$ unless with negligible probability. Since $\sum_{i=1}^{k} qx_i \pmod q$ comes from a distribution with standard deviation approximately $\sqrt{k} \cdot q^\beta$ and mean 0 we get the following result from Chebyshev's inequality, with $m = n^3$ and $t = \frac{\sqrt[3]{q}}{2\sqrt{m}\sqrt[4]{q}} \geq \frac{\sqrt[3]{q}}{\sqrt[20]{q}\sqrt[4]{q}} = \sqrt[30]{q}$.

$$\Pr(|e| \geq \sqrt[3]{q}/2) \leq \Pr(|e| \geq t \cdot \sqrt{k}\sqrt[4]{q}) \leq \frac{1}{t^2}$$

We see that $1/t^2 \leq 1/\sqrt[15]{q}$ is in fact negligibly small.

Note that the inequalities used above are not very tight, especially the Chebyshev inequality. Therefore in practice one would expect to be able to choose much better parameters, for instance a bigger standard deviation on the distribution used. This would in turn give us security reductions to the hardness of somewhat bigger lattice problem instances. Furthermore the claim is actually stronger than what is needed for the original decryption protocol to be correct, but we will need this stronger result in the proofs of the distributed decryption protocols described below.

3.3 Security

The security of the cryptosystem is given by the following theorem.

Theorem 2 (Security). *The cryptosystem is semantically secure under the assumption that* GapSVP *is hard in the worst case.*

Below we will sketch the ideas of the proof. It boils down to showing how the proofs given in [Reg05] can be adjusted to the new choice of parameters.

Proof. The proof of security given in [Reg05] is based on the property that distinguishing between encryptions of 0 and 1 is at least as hard as distinguishing public keys from randomly chosen elements in $\mathbb{Z}_q^n \times \mathbb{Z}_q$. The latter problem being the decision LWE problem. The proof of the reduction does not depend heavily on the values of the parameters, and is therefore still valid with the new choice of parameters.

The decision LWE is then further reduced to search LWE. This reduction in [Reg05] heavily relies on the fact that q is chosen to be polynomial in that it does exhaustive search over all elements in Z_q. But in fact the same idea can be used when q is exponential in size, but B-smooth with B polynomial. The idea being to do the reduction modulo each of the primes p_i in q, and recombine the solutions to a full solution modulo q using the Chinese Remainder Theorem.

The last step is to reduce search LWE to standard lattice problems. Since q is chosen to be exponentially large we can use the reduction to GapSVP made in [PVW08].

This is another advantage of choosing an exponentially large q: With the original choice of a polynomial q the reductions to lattice problems are either a quantum reduction as in [Reg05] or a reduction to a special variant of GapSVP, the hardness of which is not completely understood.

4 Distributed Decryption (Passive Adversaries)

In this section we present a distributed decryption protocol for the above cryptosystem involving u players which is secure against a static, passive adversary corrupting up to $t = u - 1$ players. That is, we assume the adversary is able to see all messages and internal data of a corrupted player, but the player still follows the protocol. The adversary must choose which players to corrupt at the start of the protocol.

We assume that communication is synchronous and that the client has access to a broadcast channel to all players. Private channels between players are not necessary since there is no interaction between players in the protocol. We assume the adversary sees all communication between the client and the players.

We use Shamir secret sharing over $\mathbb{Z}_q$ as described in [Sha79] to make secret sharings of various values in the protocol. Normally Shamir secret sharing is done over a field, but since q is not a prime $\mathbb{Z}_q$ is only a ring. This turns out not to be a problem with the choice made of the prime factors in q. The only thing that is needed is that one can do Lagrange interpolation over the points $1, \ldots, u$ which in turn boils down to being able to invert elements in this range. We chose $q = \prod p_i$, where $p_i > u$, so obviously invertion of the points needed is possible.

We furthermore make use of the concept of pseudorandom secret sharing (PRSS) described in [CDI05]. PRSS will enable the players to non-interactively share a common random value from some interval. The idea is as follows. For each subset A of size t of the players we associate a key $K_A \in_R \mathbb{Z}_q$. Such a key is given to player i exactly if $i \notin A$. Assume we are given a pseudorandom function ϕ that given a key and a ciphertext as input, will output values in the interval $[-\sqrt{q}, \sqrt{q}]$. A player can now compute $\phi_{K_A}(c)$ for all K_A he has been given, and afterwards take an appropriate linear combination of the results. This will result in all players having a Shamir share of the common random value $x = \sum_A \phi_{K_A}(c)$. Since $|A| = t$ there are $\binom{u}{t}$ possibilities for A, so x will be in the interval $\left[-\binom{u}{t}\sqrt{q}, \binom{u}{t}\sqrt{q}\right]$. We note that $\binom{u}{t} = u$ for our choice of t (but we will consider other choices later).

The protocol and proofs will be given in the setting of the Universal Composability (UC) framework proposed by Canetti. For details of this see [Can01].

4.1 Key Generation and Distribution

We assume for now that generation and distribution of keys and key-shares to players are handled by some trusted party. This is described by the functionality F_{KeyGen}.

Functionality F_{KeyGen}

1. When receiving "start" from all honest players, choose the secret key $\mathbf{s}$ and construct the public key $(\mathbf{a}_i, b_i)_{i=1}^m$ as described in section 3. Furthermore for each subset A of size t of the players, choose key $K_A \in_R \mathbb{Z}_q$.
2. For each entry j in the secret key make a share $s_{i,j}$ for each player i. We write $[\mathbf{s}]$ as short for the set of shares in $\mathbf{s}$. To each player i privately send to him his shares from $[\mathbf{s}]$ and all keys K_A where $i \notin A$.
3. Finally send the public key to all players and the adversary.

4.2 Decryption Protocol

We now describe the decryption protocol. To make things more easily describable we introduce a client, who is the party receiving the ciphertext in the first place, and who wants to decrypt with help from the players.

Protocol *Decrypt*

1. Each player sends "start" to F_{KeyGen} and stores the public key, the share of the secret key and the keys K_A received.
2. When receiving a ciphertext $c = (\mathbf{a}, b)$, the client broadcasts c to all players.
3. The players compute $[e'] = [b - \langle \mathbf{a}, \mathbf{s} \rangle] = [e + \lfloor \frac{q}{2} \rfloor \cdot \gamma]$. Since $(\mathbf{a}, b)$ is public this is a linear operation on $\mathbf{s}$ and only requires the players to locally do the same linear operation on their shares. Then $\phi_{K_A}(c)$ is computed for all the keys K_A the player received and the player takes an appropriate linear combination of the result to obtain a sharing $[x] = [\sum_A \phi_{K_A}(c)]$. Finally the players compute $[x + e']$, and send all these shares to the client.
4. Having received all the shares of $[x + e']$ the client reconstructs $x + e'$, checks whether it is closer to 0 or to $q/2$, and outputs 0 or 1 accordingly.

4.3 Security

To prove security we wish to be able to implement the following functionality.

Functionality $F_{KeyGen-and-Decrypt}$

1. Upon receiving "start" from all honest players, choose the secret key and construct the public key to be used. Send the public key to all players, the client and the adversary.
2. Hereafter on receiving "decrypt $(\mathbf{a}, b)$" from the client, send "decrypt $(\mathbf{a}, b)$" to all players and the adversary.
3. In the next round, send "result γ" to the client and the adversary, where γ is the bit corresponding to the given ciphertext.

The security is now given by the following theorem.

Theorem 3 (Security). *When given access to the functionality F_{KeyGen} and assuming that ϕ is a pseudo-random function, the protocol Decrypt securely implements $F_{KeyGen-and-Decrypt}$. The adversary is assumed to be passive and static, corrupting up to $t = u - 1$ of the players.*

Proof. We abbreviate $F_{KeyGen-and-Decrypt}$ by F_{KG-D} in the following. To prove security we must construct a simulator to work on top of the ideal functionality F_{KG-D}, such that an adversary playing with either the simulator and ideal functionality or the real world decryption protocol cannot tell in which case he is. We denote by *Adv* the adversary communicating with the real decryption protocol and must show that we can simulate everything *Adv* sees. The simulation proceeds as follows.

1. Let B denote the set of players corrupted by *Adv*. When receiving "start" to F_{KeyGen} send "start" to F_{KG-D}. Upon receiving the public key, compute a sharing of $\mathbf{0}$, the zero-vector in $\mathbb{Z}_q^n$, to simulate sharing the secret key. Also choose the necessary keys K_A. Then send to the adversary the public key, the shares of the secret key corresponding to B, and the keys K_A that should be send to players in B.

2. When receiving "decrypt $(\mathbf{a}, b)$" from F_{KG-D}, the ciphertext is sent to Adv for each player in B. When "result γ" is received in the next round, we have to simulate the shares of $x + e'$ that honest players would send. To play the role of x, we form a value y as the sum of those $\phi_{K_A}(c)$ where the adversary knows K_A, and one uniformly random value from $[-\sqrt{q}, \sqrt{q}]$ for each K_A that adversary does not know. The idea is to let $y + \lfloor \frac{q}{2} \rfloor \cdot \gamma$ play the role of the value $x + e + \lfloor \frac{q}{2} \rfloor \cdot \gamma$ that would be revealed in the real protocol. Note that from the shares and keys given to the adversary, we can compute the shares corrupted players would send to the client. Using Lagrange interpolation, we can compute a polynomial f of degree at most t that is consistent with these shares and has $f(0) = y + \lfloor \frac{q}{2} \rfloor \cdot \gamma$. We use this polynomial to compute shares for the honest players and give these to the adversary.

The final thing is to prove that no environment is able to distinguish between the real decryption protocol and the simulation presented above. This basically comes down to proving that the decryption protocol is able to recover the bit encrypted and that the distributions of the shares sent to the adversary in both cases are computationally indistinguishable.

The shares of the secret key in step 1 are distributed the same in both cases beacuse of the security of the underlying secret sharing scheme used. The keys K_A are also obviously distributed identically in the two cases.

Next, note that in both simulation and real protocol, the shares revealed in the decryption step follow deterministically from the information sent in step 1 and the values $y + \lfloor \frac{q}{2} \rfloor \cdot \gamma$, $x + e + \lfloor \frac{q}{2} \rfloor \cdot \gamma$ used in simulation, respectively real protocol. It is therefore enough to show that these values are computationally indistinguishable in the view of the adversary. For this, note that in the real protocol the adversary is not given all keys K_A, and so, by pseudorandomness of ϕ and construction of y, $y + e + \lfloor \frac{q}{2} \rfloor \cdot \gamma$ is computationally indistinguishable from the $x + e + \lfloor \frac{q}{2} \rfloor \cdot \gamma$ in the view of the adversary. Second since y is a sum including at least one value that is uniform in an interval of size $2\sqrt{q}$, which is exponentially larger than the interval $[-\sqrt[3]{q}, \sqrt[3]{q}]$ in which e is distributed, we find that $y + \lfloor \frac{q}{2} \rfloor \cdot \gamma$ is statistically indistinguishable from $y + e + \lfloor \frac{q}{2} \rfloor \cdot \gamma$.

Finally in both the simulated and the real run the client will output the correctly decrypted value. This is obvious in the simulated case and in the real world it follows from Lemma 1 below.

Lemma 1 (Correctness). *Let $\binom{u}{t} < \frac{1}{4}\sqrt{q} - 1$. Assume that for any $k \in \{0, 1, ...m\}$, χ^{*k} satisfies that*

$$\Pr_{e \sim \chi^{*k}} [|e| \geq \lfloor \sqrt[3]{q} \rfloor] \leq 2^{-O(n)}.$$

Then the error probability when decrypting is negligible.

Proof. Given an encryption of 0 the result which is reconstructed is given by $b - \langle \mathbf{a}, \mathbf{s} \rangle = e + x = \sum_{i \in S} e_i + x$. The distribution of e is exactly given by $\chi^{*|S|}$, therefore according to our assumption $|e| < \lfloor \sqrt[3]{q} \rfloor$ with probability at least $1 - 2^{-O(n)}$. Since $\binom{u}{t} < \frac{1}{4}\sqrt{q} - 1$ according to our assumption, we have that $|x| < \frac{q}{4} - \sqrt[3]{q}$. Combined we get that $|e + x| < \frac{q}{4}$ with probability at least $1 - 2^{-O(n)}$. In this case the result is closer to 0 than $\frac{q}{2}$ and the decryption is correct. A similar proof can be done for an ecryption of 1.

The assumptions in the lemma are fulfilled according to the claim in section 3. We note that the correctness puts an upper bound on the possible number of players, which is also to be expected, since there is a limit to how much random noise can be added before an encryption of 0 turns into an encryption of 1. Note though that when $t = u - 1$, as is the case in the passive case, we have $\binom{u}{t} = u$. So here the number of players is bounded by approximately $\sqrt{q}$ which is still quite a big number.

5 Distributed Decryption for Stronger Adversaries

The protocol for doing distributed decryption against a passive adversary corrupting less than $t = u - 1$ players, can easily be turned into a protocol secure against a stronger adversary. First, if the adversary is semi-honest, i.e. corrupted players follow the protocol but may stop at any point, exactly the same protocol will be secure, if $t < u/2$. The proof is the same, one just notes that at least $t + 1$ players will always complete the protocol.

If the adversary is active, again almost the same protocol and proof applies, if we assume $t < u/3$. The only significant difference to the protocol is that the client must use standard methods for error correction to reconstruct $x + e'$ at the end of the decryption since some players may lie about their shares. This is possible exactly when $t < u/3$.

It should be noted that both variants of the protocol are only feasible to execute for a small number of players, since the number of keys K_A we must give to each player increases exponentially with u whenever t is a constant fraction of u. However, in most realistic applications of threshold cryptography, one indeed expects the number of players to be small.

6 Distributed Key Generation

In this section we will describe how to do key generation and distribution against an active adversary, although some details have been omitted. In some of the parts involving interaction between the players, we will have to assume private communication channels between players.

The main idea to generate the secret key is simply to let players choose random vectors, share them among each other and letting the final secret key be the sum of the individually chosen vectors. To generate the keys K_A for doing PRSS we simply let one player not in A choose K_A and distribute it to all players also not in A. For generating the public key $(\mathbf{a}_i, b_i)_{i=1}^m$ the tricky part is that shares of non-uniformly distributed values are to be generated and distributed. For this we use the technique of non-interactive verifiable secret sharing (NIVSS) described in [CDI05] which builds on top of PRSS described earlier. The a_i's in the public key will be generated the same way the secret key is generated.

In the protocol for doing NIVSS, once again a common input is given and keys K_A similar to PRSS. Furthermore a dealer holds all keys and the value v to be shared. Each player computes a preliminary random share using PRSS, and the dealer can then compute the secret s determined by the preliminary shares. Also knowing the value v the dealer then broadcasts $v - s$, and each player can now compute their share of v. To use this to securely generate the noise variables e_i in the public key, the NIVSS protocol will be performed u times such that each player gets to act as dealer. This way each player j gets to choose a noise contribution $e_{i,j}$ to the final noise e_i which is then simply $\sum_j e_{i,j}$. To furthermore guard against active adversaries the pseudorandom function used will be outputting values from a much smaller interval than $\mathbb{Z}_q$[1], effectively restricting the size of the noise contribution $e_{i,j}$ an adversary j might choose to add to the final noise e_i.

[1] This smaller interval could for instance be $[-\sqrt[3]{q}, \sqrt[3]{q}]$ as was used in the distributed decryption protocol. This change will mean little to an honest player.

6.1 Key Generation Protocol

We now describe the key generation protocol.

Protocol *KeyGeneration*

1. To generate and distribute the secret key each player j chooses a vector $\mathbf{s}_j \in_R \mathbb{Z}_q^n$ and sends shares to all players. Each player then calculates his share of the final secret key $\mathbf{s} = \sum_j \mathbf{s}_j$.
2. To generate and distribute the keys K_A for PRSS, for each subset A some player not in A will choose the key K_A. Using techniques for verifiable secret sharing he shares K_A among the players also not in A. Finally all players not in A send their shares to all players not in A, such that all the players that was supposed to get K_A can reconstruct.
3. To generate the public key each player j first chooses m vectors $\mathbf{a}_{1,j}, \ldots, \mathbf{a}_{m,j} \in_R \mathbb{Z}_q^n$ and share these among all players. Each player then calculates his share of the final $\mathbf{a}_i = \sum_j \mathbf{a}_{i,j}$, and the players open and reconstruct the $\mathbf{a}_i$'s.
4. Each player j then chooses noise contributions $e_{i,j}$ according to the distribution $\overline{\Psi}_\alpha$ and by using the NIVSS protocol described above and in [CDI05] shares these values among the players. The final noise variables e_i is now given by $\sum_j e_{i,j}$, and each player can locally compute their share of e_i.
5. Finally each player can compute their shares of the b_i's in the public key $(\mathbf{a}_i, b_i)_{i=1}^m$ locally, send them to all other players, and the players open to reconstruct the public key. The public key is sent to the adversary.

6.2 Security

We now give a discussion of the security of the key generation protocol, a formal proof is left for the full version of the paper. In the case of an active adversary the assumption is that less than $t = u/3$ of the players are corrupted.

It is clear from the security of the secret sharing scheme that the generation of the secret key will be correct. Furthermore it will not leak any information about the keys.

Concerning the keys K_A used for PRSS we first note that if we are in the case of a passive adversary the above protocol is overly complicated. In this case the selected player could simply send the key K_A to all other players who was supposed to receive K_A. The problem with this approach in case of an active adversary is that the adversary may not send the same K_A to all players, leading to the players not in A not agreeing on the key. Therefore we use verifiable secret sharing to make sure that this is not possible for the adversary. Furthermore we note that if the adversary gets to see the key K_A, which he does exactly if he is not in A, then the distribution of K_A does not matter. Finally the adversary gets no information about the keys K_A that he was not supposed to get, since he does not participate in the generation of these.

For the public key $(\mathbf{a}_i, b_i)_{i=1}^m$ the security and correctness issues are concerning the generation of the $\mathbf{a}_i$'s and the b_i's. The $\mathbf{a}_i$'s will be generated correctly, and since all players have already committed to their contributions before the vectors are revealed, a corrupt player have no way of choosing his contribution in a more clever way than just following the protocol. For the b_i's, as already mentioned in the protocol, it is straight forward for each player to calculate $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i$, when the e_i's have been generated and distributed. The security comes from the security of the underlying secret sharing scheme.

We look at the generation of e_i. The e_i's generated by the above protocol is actually not distributed according to exactly the right distribution, but instead according to $\overline{\Psi}_{\sqrt{u}\cdot\alpha}$. This though, does not affect the system much. If one wants e_i distributed exactly according to $\overline{\Psi}_\alpha$ a smaller error distribution depending on the number of players can be used. Another factor contributes to not having exactly the correct distribution, namely the restriction on the output interval of the pseudorandom function used for NIVSS. Again this does not matter much to an honest player, since the probability of sampling outside of the interval is very small.

We first note that the generation of the public key does not leak any information about the secret key or the noise added to the b_i's. This is true in both the passive and active case and is due to the security of the underlying secret sharing scheme.

We should also argue that the keys generated can be used securely for encryption. In the case of a passive adversary this requirement is fulfilled, since the generated keys are distributed (almost) according to the original cryptosystem which is assumed to be secure. For an active adversary the best he can do is to get rid of as much noise as possible, that is the noise he already knows, but then he ends up with keys distributed according to the original cryptosystem which again is assumed to be secure. Finally the restriction on the interval from which the noise contributions can be taken ensures that decryption is still possible, even in the presence of an active adversary who might not choose his noise contribution from the distribution intended.

7 Zero-Knowledge Proof of Plaintext Knowledge

In this section, we consider Regev's original cryptosystem, where the random choices and plaintext are binary and q is a prime. All arithmetic in this section is modulo q. In the appendix we describe a slightly more complicated scheme that works for our variant

We define the relation R_{Regev} as the set of pairs $\{x, w\}$ such that $x = ((\mathbf{a}_i, b_i)_{i=1..m}, (\mathbf{a}, b))$, and $w = (r_1, ..., r_m, \gamma)$ such that $(\mathbf{a}, b) = (\sum_{i=1}^m r_i\mathbf{a}_i, \gamma \cdot \lfloor \frac{q}{2} \rfloor + \sum_{i=1}^m r_i b_i)$. The language L_{Regev} will be the set of x for which there exist w with $(x, w) \in R_{Regev}$. Our goal is to build a zero-knowledge interactive proof for L_{Regev} which is also a proof of knowledge for R_{Regev}. In other words, the prover demonstrates that the ciphertext is well-formed and that he knows the plaintext and random coins that were used to form it.

We will use the technique from [IKOS07] where it was shown how to construct zero-knowledge proofs from multiparty computation protocols. We briefly sketch the idea: Assume we have a multiparty computation protocol π for input client I, players $P_1, ..., P_u$ and output client O, where I gets the prover's secret witness as input, shares it among the players, who then carry out a secure computation that verifies whether the witness is valid with respect to the public common input. The players send their results to O who outputs 1 or 0 accordingly. The protocol must be secure against a malicious adversary corrupting the clients and/or up to t of the other players. The prover now emulates π "in his head" and commits to the views of all players. Here, a view consists of the inputs and random coins of the player, and all received messages. The verifier selects a random subset of players among those that π can tolerate as corrupted sets[2]. The prover must open the corresponding commitments and the verifier checks that these views are consistent with each other and with the protocol and accepts or rejects accordingly.

The intuition is that the protocol is zero-knowlegde since π is secure even if the set chosen by the verifier is corrupted, and hence no information on the secret witness is released. The protocol

[2] The protocol must be secure against a corrupt I, but the verifier is of course allowed to "open" I.

is sound since if the witness is invalid, the prover must introduce some inconsistency to make it seem that π accepts the witness.

Indeed, it is shown in [IKOS07] that if π implements the function that checks the witness with perfect sercurity and if both u and t are $\theta(n)$, then the resulting two-party protocol has soundness error $2^{-\Omega(n)}$. It is honest verifier zero-knowledge, and can be made zero-knowledge in general, e.g., by generating the verifier's choice of subset to corrupt via a suitable coinflip protocol.

We make a couple of observations that are helpful in constructing a protocol π for our purposes: first, while broadcast is usually considered an expensive resource, it is virtually for free in this setting - any information π would broadcast can just be sent to the verifier immediately, as he would see it anyway no matter what subset is chosen. This was already noted in [IKOS07]. Second, π does not have to guarantee termination, in the following sense: suppose all players broadcast some message in some round of π, and then all honest players decide (using the same procedure) whether to abort or continue. Suppose further that if all players have behaved honestly so far, we will never abort, and that further π has perfect correctness and privacy conditioned on the event that we do not abort. In this case, we can simply ask the verifier to reject if the prover sends a set of broadcast messages that would cause an abort. This will not hurt the honest prover, but will force a cheating prover to (claim that) he lets the virtual players behave such that π terminates.

In view of the above, all we have to do is to build an efficient protocol π that checks $r_1, ..., r_m, \gamma$ against $(\mathbf{a}_i, b_i)_{i=1..m}$ and $(\mathbf{a}, b)$. In order to do this, we need to borrow two tools from the design of efficient multiparty protocols, namely Packed Secret-Sharing[FY92] and Hyper-Invertible Matrices[BTH08], which we describe below.

7.1 Packed Secret Sharing

Packed Secret-Sharing is a generalization of standard Shamir sharing where secret values are assigned to more than one interpolation point. In other words, the secret to share is in fact a vector $(x_1, ..., x_\ell) \in \mathbb{Z}_q$. To do the sharing, we construct a random polynomial f of degree at most d, such that $f(0) = x_1, f(-1) = x_2, ..., f(-\ell+1) = x_\ell$. The shares are, as usual, $f(1), ..., f(u)$. To make this possible, and to guarantee privacy against t corrupted players, d must be at least $t + \ell - 1$. In our case, we will choose $\ell = n + 1$, and t to be $\theta(n)$. Furthermore, we will need that there are sufficiently many honest players such their shares alone can determine a polynomial of degree $2d$, i.e., $u - t \geq 2(t + n + 1)$. This shows that we can indeed choose u to be $\theta(n)$, as promised above.

Note that to ensure we have enough distinct evaluation points, we need that if q is a prime, it must be larger than $\ell + u = n + 1 + u$ which is $\theta(n)$ or, in our construction of q for the threshold scheme, the smallest prime factor must be larger than $\ell + u$. This is already satisfied by the schemes as they stand.

We will write $[\mathbf{z}]_d$ for a set of shares determining a packed sharing of the block $\mathbf{z}$ using a polynomial of degree d.

Note that if players locally add respectively multiply their shares of blocks $\mathbf{z}, \mathbf{z'}$, this results in shares in the coordinate-wise sum respectively product, i.e., we have $[\mathbf{z}]_d + [\mathbf{z'}]_d = [\mathbf{z}+\mathbf{z'}]_d$, and $[\mathbf{z}]_d * [\mathbf{z'}]_d = [\mathbf{z} * \mathbf{z'}]_{2d}$, where $*$ denotes the coordinate-wise product.

7.2 Hyper-Invertible Matrices

A hyper-invertible matrix M (with entries in $\mathbb{Z}_q$) has the property that any square submatrix of M is invertible. Such matrices can be constructed from Van der Monde matrices and were used in [BTH08] to check consistency of secrets sharings with zero error probability. We briefly explain how this works:

Suppose M is a matrix with u rows and $u-t$ columns. Suppose the players hold $u-t$ sets of shares $[\mathbf{z}_1], ..., [\mathbf{z}_{u-t}]$, and we want to check that each set of shares is consistent with a polynomial of degree at most e. The players can locally compute u new sets of shares,

$$[M(\mathbf{z}_1, ..., \mathbf{z}_{2v})_1], ..., [M(\mathbf{z}_1, ..., \mathbf{z}_u)_u] := [\mathbf{y}_1], ..., [\mathbf{y}_u],$$

simply by multiplying M on the vector of $u-t$ shares that they hold (thinking of the shares as a column vector). Assume now that for $i = 1..u$, each player sends his share in $[\mathbf{y}_i]$ to P_i. This allows P_i to check that the shares he receives are e-consistent, i.e., on a polynomial of degree at most e. P_i can now broadcast whether his check was OK or not.

We can now see that if all players are happy, this means in particular that all honest players were happy, and that they therefore agree with all honest players on the set of $u-t$ e-consistent shares that they checked, i.e., $\{[\mathbf{y}_j]\}_{j\in H}$, where H is the set of honest players, are all e-consistent. Let M_H be the matrix we get from M by only taking the rows corresponding to players in H. This matrix is invertible by assumption on M, so we can obtain $[\mathbf{z}_1], ..., [\mathbf{z}_{u-t}]$ as a linear function defined by M_H^{-1} of the the shares in $\{[\mathbf{y}_j]\}_{j\in H}$, and hence the $[\mathbf{z}_i]$'s are all e-consistent as well.

Furthermore, if it is important that the shared information is kept secret, one can arrange the input shares such that only $[\mathbf{z}_1], ..., [\mathbf{z}_{u-2t}]$ contains information we want to protect, while $[\mathbf{z}_{u-2t+1}], ..., [\mathbf{z}_{u-t}]$ are chosen randomly using polynomials of degree at most e. These t random sets of shares will randomize the t sets of shares seen by corrupt players, again by hyper-invertibility of M. This also means that we do not need, for instance, $[\mathbf{z}_1]$ to be a *random* sharing of $\mathbf{z}_1$ to be able to hide it.

Note also that this method can be used to also check if $\mathbf{z}_1, ..., \mathbf{z}_{u-t}$ all satisfy some fixed condition, as long as what the condition asks is that each $\mathbf{z}_i$ satisfies some linear equation. For instance, we might want to check that $\mathbf{z}_i = (0, ..., 0)$ for all i. This is done by having players verify that all $\mathbf{y}_i$ satisfy the same condition.

Regarding the complexity, it is easy to see that a set of shares of total size T bits can be verified while keeping the shared information perfectly private by sending $O(T)$ bits and creating random shares of size $O(T)$ bits.

7.3 The Multiparty Protocol

Recall that the secret witness to be checked consists of binary values $r_1, ..., r_m, \gamma$ where $(\mathbf{a}, b) = (\sum_{i=1}^m r_i\mathbf{a}_i, \gamma \cdot \lfloor \frac{q}{2} \rfloor + \sum_{i=1}^m r_i b_i)$, and where the public information is public key $(\mathbf{a}_i, b_i)_{i=1..m}$ and ciphertext $(\mathbf{a}, b)$. For any $z \in \mathbb{Z}_q$, we set $\overline{z} = (z, z, ..., z)$, a vector of length $n+1$. The protocol works as follows:

Protocol VerifyCiphertext

1. The input client I send shares $[\overline{r_i}]_d$, $i = 1..m$ and $[\overline{\gamma}]_d$ to the players. In addition, it also send random shares as required for the verifications below using the hyper-invertible matrix M.
2. Verify that $[\overline{r_1}]_d, ..., [\overline{r_m}]_d, [\overline{\gamma}]_d$ are d-consistent and that in each block shared, all $n+1$ entries are equal. If any player broadcasts "not OK", the protocol aborts.
3. Compute, using local multiplications, $[\overline{r_i(1-r_i)}]_{2d}$ for $i = 1..m$ and $[\overline{\gamma(1-\gamma)}]_{2d}$.

4. Form sharings of the public vectors: $[(\mathbf{a}_i, b_i)]_d$, $i = 1..m$, $[(0,...0,\lfloor\frac{q}{2}\rfloor)]_d$, and $[(\mathbf{a}, b)]_d$ (using some default choice of polynomial of degree at most d). We then emulate the encryption on the shared values: compute, using local computation,

$$[\sum_{i=1}^{m} \overline{r_i} * (\mathbf{a}_i, b_i)]_{2d} + [(0,...,0,\lfloor\frac{q}{2}\rfloor) * \overline{\gamma}]_{2d} = [(\sum_{i=1}^{m} r_i\mathbf{a}_i, \sum_{i=1}^{m} r_i b_i + \gamma\lfloor\frac{q}{2}\rfloor)]_{2d}$$

From this, we locally subtract shares of the ciphertext $[(\mathbf{a}, b)]_d$, so we get

$$[(\sum_{i=1}^{m} r_i\mathbf{a}_i - \mathbf{a}, \sum_{i=1}^{m} r_i b_i + \gamma\lfloor\frac{q}{2}\rfloor) - b]_{2d} := [(\mathbf{z}, v)]_{2d}$$

5. Verify that $[\overline{r_1(1-r_1)}]_{2d}, ..., [\overline{r_m(1-r_m)}]_{2d}$, $[\overline{\gamma(1-\gamma)}]_{2d}$ and $[(\mathbf{z}, v)]_{2d}$ are indeed $2d$-consistent sharings of all-zero blocks. If any player broadcasts "not OK", the protocol aborts. This ensures that the r_i's and γ are binary, and that encryption results in the claimed ciphertext.

Since the verifications of shares works with zero error probability, it is clear that if the protocol terminates successfully, we are guaranteed that the shared values determine the correct ciphertext. No information on the secret is released, since the only communication is what is required for the verification of sharings, and we already argued above that these release no information on the shared values that we verify.

Regarding complexity, it is clear from inspection of the protocol that it is completely determined by the total size T of the sharings $[\overline{r_i}]_d$, $i = 1..m$ and $[\overline{\gamma}]_d$, in particular, the total size of communication is $O(T)$. We have that T is $O(mu \log q)$ which is $O(mn \log q)$. Note that the size of the ciphertext is also $O(mn \log q)$.

It is described in [IKOS07] how to transform this protocol to a zero-knowledge proof using an unconditionally binding commitment scheme. If this scheme allows us to commit to strings with an additive length increase that is independent of the string length, we can preserve the efficiency of the multiparty protocol. An unconditionally hiding commitment scheme is also needed, for the verifier to commit to his challenge. This gives us:

Theorem 4. *Given an unconditionally binding and an unconditionally hiding commitment scheme with constant additive overhead, using protocol VerifyCiphertext in the construction from [IKOS07] produces a two-party zero-knowledge proof for L_{Regev}. The protocol has communication complexity $O(mn \log q)$ bits and error probability $2^{-\Omega(n)}$.*

We can base the commitment schemes needed on lattice problems, thus using assumptions we would need anyway. An efficient unconditionally binding scheme follows from the cryptosystem in[PVW08], while an unconditionally hiding scheme can be based on any collision intractable hash function [DPP98], which can in turn be based on lattice assumptions.

In [IKOS07], it was not shown that their construction is a proof of knowledge for R_{Regev}. However, for the honest verifier zero-knowledge version of the protocol, one can do a rewinding argument to show that it is indeed a proof of knowledge with negligible knowledge error. If we go to the version that is zero-knowledge in general, things are different, since the construction from [IKOS07] has the verifier commit to his challenges, which means rewinding the prover is not possible unless the extractor can equivocate these commitments.

However, in the common reference string model, we can easily make the protocol be a proof of knowledge for R_{Regev}, by having a public key for a commitment scheme placed in the reference string, e.g., a public key for the cryptosystem from [PVW08], and the prover uses these for committing to the views. If the extractor knows the corresponding secret key, it can extract all committed views without rewinding and easily compute the secret.

References

[BTH08] Zuzana Beerliová-Trubíniová and Martin Hirt. Perfectly-secure mpc with linear communication complexity. In *TCC*, pages 213–230, 2008.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145, 2001.

[CDI05] Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In *TCC*, pages 342–362, 2005.

[DPP98] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. Statistical secrecy and multibit commitments. *IEEE Transactions on Information Theory*, 44(3):1143–1151, 1998.

[FY92] Matthew K. Franklin and Moti Yung. Communication complexity of secure computation (extended abstract). In *STOC*, pages 699–710, 1992.

[IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *STOC*, pages 21–30, 2007.

[MR08] Daniele Micciancio and Oded Regev. Lattice-based cryptography. In D. J. Bernstein and J. Buchmann, editors, *Post-quantum Cryprography*. Springer, 2008.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *STOC*, pages 333–342, 2009.

[PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.

[Sha79] Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.

A Zero-Knowledge Proof when q is not prime

The only part of the multiparty protocol underlying our zero-knowledge proof that does not work when q is not a prime is the step where it is verified that the r_i are binary, essentially by verifying that $r_i(1-r_i) \bmod q = 0$. Of course, this check is not good if q is not prime. We sketch a procedure that can be used instead, but has only statistical security:

The input client I supplies $[\overline{r_i}]_d$ and it is checked as in the original protocol that a block has been shared where all entries are equal. Note that if the sharing was correctly formed, it would be the case that $\mathbf{r}' = 2\overline{r_i} - (1, ..., 1)$ would be $(1, ..., 1)$ or $(-1, ..., -1)$. I also supplies a sharing $[\mathbf{z}]_d = [(z_1, ..., z_{n+1}]_d$ such that all z_i are randomly chosen to be 1 or -1. Finally, a public random challenge is generated: $\mathbf{v} = (v_1, ..., v_{n+1})$, where each v_i is 0 or 1. (when transforming this to a 2-party protocol, we let the verifier generate the challenge). We compute (locally)

$$[\mathbf{r}' * \mathbf{z} * \mathbf{v} + \mathbf{z} * (\overline{1} - \mathbf{v})]_{3d}.$$

And finally we add a random degree $3d$-sharing of the all-zero block and open the result. The opened block must contain only 1's and -1's. Put another way, the opening shows us, in each coordinate position, an entry from $\mathbf{z}$ or from $\mathbf{r}' * \mathbf{z}$ and they must all be ± 1.

For privacy, the intuition is that by random choice of $\mathbf{z}$, $\mathbf{r}' * \mathbf{z}$ has no information on $\mathbf{r}'$ and neither does $\mathbf{z}$ so seeing, for each index i, the i'th entry of $\mathbf{r}' * \mathbf{z}$ or $\mathbf{z}$ reveals nothing on $\mathbf{r}'$.

For correctness, if there is just a single position in which both $\mathbf{z}$ and $\mathbf{r}' * \mathbf{z}$ are ± 1, $\mathbf{r}'$ will be ± 1 in that position too, and this implies that the original r_i was 0 or 1. On the other hand, if no such position exists, the honest players will accept $[\overline{r_i}]$ with probability only 2^{-n-1}, by the assumed randomness of $\mathbf{v}$.