

Learning with Errors in the Exponent

Özgür Dagdelen¹(✉), Sebastian Gajek^{2,3}, and Florian Göpfert⁴

¹ BridgingIT GmbH, Mannheim, Germany
oezguer.dagdelen@bridging-it.de

² NEC Research Labs, Heidelberg, Germany
sebastian.gajek@neclab.eu

³ Flensburg University of Applied Sciences, Flensburg, Germany

⁴ Technische Universität Darmstadt, Darmstadt, Germany
fgoepfert@cdc.informatik.tu-darmstadt.de

Abstract. The Snowden revelations have shown that intelligence agencies have been successful in undermining cryptography and put in question the exact security provided by the underlying intractability problem. We introduce a new class of intractability problems, called Learning with Errors in the Exponent (LWEE). We give a tight reduction from Learning with Errors (LWE) and the Representation Problem (RP) in finite groups, two seemingly unrelated problem, to LWEE. The argument holds in the classical and quantum model of computation.

Furthermore, we present the very first construction of a semantically secure public-key encryption system based on LWEE in groups of composite order. The heart of our construction is an error recovery “in the exponent” technique to handle critical propagations of noise terms.

Keywords: Lattice theory · Group theory · Public-key encryption · Intractability amplification

1 Introduction

Among the most carefully scrutinized cryptographic problems are probably the discrete logarithm in finite groups and factorization. Shor’s celebrated theorems [1, 2] curtailed for the first time the confidence of founding cryptosystems on group-theoretic assumptions. Shor showed the existence of polynomial-time solvers for integer factorization and discrete logarithm computation in the non-classical quantum computation model. Researchers have then begun to look for alternative computational problems. In this line of work Regev explored a lattice problem class known as learning with errors (LWE) [3]. Given a distribution of noisy equations $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ where e is taken from a small Gaussian error distribution, the search learning with error problem states it is hard to compute the solution \mathbf{s} whereas the decisional variant assumes it is hard to distinguish (\mathbf{a}, b) from uniformly random elements in $\mathbb{Z}_q^n \times \mathbb{Z}_q$. Several arguments flesh out LWE’s intractability [4]: First, the best known solvers run in exponential time and even quantum algorithms do not seem to help. Second,

learning with errors is a generalization of learning from parity with error, which is a well-studied problem in coding theory. Any major progress in LWE will most likely cause significant impact to known lower bounds of decoding random linear codes. Lastly and most importantly, breaking certain average-case problem instances of LWE breaks all instances of certain standard lattice problems [3, 5–7].

Taking the findings from lattices in presence of errors into account we carry on the study of noise as a *non-black box* intractability amplification technique. Specifically, we ask if noise effects the intractability of group-theoretic problems as well? If so, is non-trivial cryptography possible in such groups? The main challenge is to handle the propagation of “noise in the exponent”. Error terms require a careful treatment, because they may easily distort the cryptographic task. Apart from the theoretical interest, our work has concrete practical motivation. Recent large-scale electronic surveillance data mining programs put in question the security provided by present cryptographic mechanisms. (See also the IACR statement and mission on mass surveillance.¹) One of the problems is that many security protocols in the wild are based on a single intractability problem and we do not know the exact security. What if somebody has found a clever way to factor numbers? This already suffices to decrypt most of the TLS-protected Internet traffic and eavesdrop emails, social network activities, and voice calls.² Answering the above questions in an affirmative way advertises a novel family of computationally hard problems with strong security and robustness properties in the superposition of group and lattice theory.

1.1 Our Contribution

Blending Group and Lattice Theory. The idea of blending intractability problems is not new and is subject to several Diffie-Hellman related problems in groups of composite order which assume the hardness of the discrete log or factorization problem [8, 9]. In this work, we address the blending of group and lattice related problems, and introduce the notion of *Learning with Errors in the Exponent* (LWEE). The LWEE distribution consists of samples $(g^{\mathbf{a}}, g^{(\mathbf{a}, \mathbf{s})+e}) \in \mathbb{G}^n \times \mathbb{G}$ where \mathbf{a} is sampled uniformly from \mathbb{Z}_q^n , and $\mathbf{s} \leftarrow_R \chi_s^n$, $e \leftarrow_R \chi_e$ from some distributions χ_s, χ_e . Learning with errors in the exponent comes in two versions: The search version asks to find the secret vector \mathbf{s} while in the decisional variant one is supposed to distinguish LWEE samples from uniformly random group elements. Except for the error the assumption bears reminiscence to the representation problem RP [10]. Given a tuple of uniformly sampled elements g_1, \dots, g_ℓ, h from \mathbb{G} , the (search) representation problem (ℓ -SRP) asks to compute the “representation” $x_1, \dots, x_\ell \leftarrow \chi$ with respect to h for χ the uniform distribution such that $\prod_{i=1}^\ell g_i^{x_i} = h$. We give a tight reduction from ℓ -SRP to the search LWEE problem.

¹ <http://www.iacr.org/misc/statement-May2014.html>.

² TLS’s preferred cipher suite makes use of RSA-OAEP to transport the (master) key in the key establishment process. Once the ephemeral master key for the session is known it is possible to derive session keys and decrypt all encrypted messages.

Relations between Group and Lattice Assumptions. Looking at the decisional problem, we first define the decisional variant of the representation problem (ℓ -DRP): Given a tuple $g, g_1, \dots, g_\ell, g^{x_1}, \dots, g^{x_\ell}, h$ from \mathbb{G} , where $x_1, \dots, x_\ell \leftarrow \chi$ are sampled from some distribution χ , ℓ -DRP asks to distinguish between $\prod_{i=1}^\ell g_i^{x_i} = h$ and a randomly sampled value h in \mathbb{G} . Note, for $\ell = 1$ and uniform distribution over \mathbb{Z}_q , DRP coincides with the decisional Diffie-Hellman (DDH) problem. For $\ell > 1$, we prove in the generic group model that ℓ -DRP belongs to the class of progressively harder assumptions [11]. We then show that DRP is reducible to LWEE. This implies that if we select a group \mathbb{G} for which DDH is believed to be hard, the hardness carries over to an instantiation of LWEE in that group \mathbb{G} . It is worth mentioning that both of our reductions are tight. They hold for (potentially non-uniform) distributions χ , if the underlying RP problem is hard for representations sampled from the same distribution. Investigating the relation to lattices, we show that an algorithm solving either the search or decisional LWEE problem efficiently can be turned into a successful attacker against the search or decisional LWE problem. Our reductions are tight and hold as well for (potentially non-uniform) distribution χ if LWE is hard for secret s sampled from the same distribution.

A Concrete Cryptosystem. We give a first construction of a public-key encryption scheme. One may size the magnitude to which the RP and LWE intractability contribute to the security of the system. The selection of parameters (e.g., modulus, dimension) offers great flexibility to fine-tune the cryptosystem’s resilience against (quantum)-computational progress in attacking the underlying intractability problems. Concretely, one may choose the parameters to obtain short keys and ciphertexts, make the scheme post-quantum secure or immunities the scheme for the case that at some point in time either the DRP or DLWE becomes computationally tractable.

Although our construction serves the sole purpose of showcasing the feasibility of cryptosystems (in practical applications, it would be preferable to split the message information-theoretically into two shares and encrypt each share with a different encryption scheme, say El-Gamal and Regev encryption) based on “errors in the exponent”, learning with errors in the exponent is an interesting concept in its own right. We leave it open for future work to find novel applications and to study the instantiation based on the learning with errors assumptions in rings. We discuss related work in the full version [12].

1.2 Extensions and Open Problems

While learning with errors in the exponent is an interesting concept in its own right, it requires further inspection. Here we point out a few possible directions for future research:

- It would be interesting to cryptanalyze the assumption. This would help nail down concrete security parameters, in particular for the case of double-hardness where both underlying assumptions contribute to the overall security.

- We are unaware of any existential relation between the representation and learning with errors assumption neither in the classical nor quantum model of computation. In fact, any insight would require progress in solving the hidden subgroup problem (HSP) in certain finite Abelian and non-Abelian groups. Shor’s discrete-log quantum algorithm crucially relies on the HSP in Abelian groups. However, efficient quantum algorithms for the HSP in non-Abelian groups are unknown as they would give an efficient algorithm for solving the unique shortest-vector problem, being a special case of the shortest vector problem (SVP) [13].
- Clearly, building further cryptosystems based on the search or decisional variant of learning with errors in the exponent is an interesting direction.

2 Preliminaries

2.1 Notation

Random Sampling, Negligibility and Indistinguishability. If \mathcal{D} is a probability distribution, we denote by $d \leftarrow_R \mathcal{D}$ the process of sampling a value d randomly according to \mathcal{D} . If S is a set, then $s \leftarrow_R S$ means that s is sampled according to a uniform distribution over the set S . We write $[m]$ for the set $\{0, 1, \dots, m - 1\}$. The expression $\lceil x \rceil$ denotes the nearest integer to $x \in \mathbb{R}$, i.e., $\lceil x \rceil = \lfloor x + 0.5 \rfloor$.

A function $\varepsilon(\cdot)$ is called *negligible* (in the security parameter κ) if it decreases faster than any polynomial $\text{poly}(\kappa)$ for some large enough κ . An algorithm \mathcal{A} runs in probabilistic polynomial-time (PPT) if \mathcal{A} is randomized—uses internal random coins—and for any input $x \in \{0, 1\}^*$ the computation of $\mathcal{A}(x)$ terminates in at most $\text{poly}(|x|)$ steps. If the running time of an algorithm is $t' \approx t$, we mean that the distance between t' and t is negligible.

Let $X = \{X_\kappa\}_{\kappa \in \mathbb{N}}$ and $Y = \{Y_\kappa\}_{\kappa \in \mathbb{N}}$ be two distribution ensembles. We say X and Y are (t, ϵ) -computationally indistinguishable if for every PPT distinguisher \mathcal{A} with running time t , there exists a function $\epsilon(\kappa)$ such that $|\Pr[\mathcal{A}(X) = 1] - \Pr[\mathcal{A}(Y) = 1]| \leq \epsilon(\kappa)$ (and we write $X \approx_{(t, \epsilon)} Y$). If \mathcal{A} is PPT and $\epsilon(\kappa)$ is negligible, we simply say X and Y are two distribution ensembles. We say X and Y are (t, ϵ) -computationally indistinguishable (and write $X \approx_{(t, \epsilon)} Y$) if for every PPT distinguisher \mathcal{A} with running time t , there exists a function $\epsilon(\kappa)$ such that $|\Pr[\mathcal{A}(X) = 1] - \Pr[\mathcal{A}(Y) = 1]| \leq \epsilon(\kappa)$. If \mathcal{A} is PPT and $\epsilon(\kappa)$ is negligible, we simply say X and Y are (computationally) indistinguishable (and we write $X \approx Y$). We say a distribution ensemble $X = \{X_\kappa\}_{\kappa \in \mathbb{N}}$ has (high) min-entropy, if for all large enough κ , the largest probability of an element in X_κ is $2^{-\kappa}$. We say a distribution ensemble $X = \{X_\kappa\}_{\kappa \in \mathbb{N}}$ is well-spread, if for any polynomial $\text{poly}(\cdot)$ and all large enough κ , the largest probability of an element in X_κ is smaller than $\text{poly}(\kappa)$. (In other words, the max-entropy of distributions in X must vanish super-logarithmically.) Under the Gaussian distribution D_σ with parameter $\sigma > 0$, the probability of sampling an integer $x \in \mathbb{Z}$ is proportional to $\exp[-x^2/(2\sigma^2)]$.

Vectors and Matrices in the Exponent. We denote vectors by bold lower case letters and matrices by bold upper case letters. The i^{th} row of a matrix \mathbf{A} is denoted by $\mathbf{A}[i]$, the j^{th} element of a vector \mathbf{a} is denoted by a_j . To ease notation we sometimes write \mathbf{a}_i for the i^{th} row vector, and $a_{i,j}$ for the element in the i^{th} row and j^{th} column of matrix \mathbf{A} . Let \mathbb{G} be a group of order q , g a generator of \mathbb{G} , \mathbf{a} a vector in \mathbb{Z}_q^n , and \mathbf{A} a matrix in $\mathbb{Z}_q^{m \times n}$. We use the notation $g^{\mathbf{a}} \in \mathbb{G}^n$ to denote the vector $g^{\mathbf{a}} \stackrel{\text{def}}{=} (g^{a_1}, \dots, g^{a_n})$ and $g^{\mathbf{A}} \in \mathbb{G}^{m \times n}$ to denote the matrix $g^{\mathbf{A}} \stackrel{\text{def}}{=} (g^{\mathbf{a}_1}, \dots, g^{\mathbf{a}_m})^\top$.

Computations in the Exponent. Given $g^{\mathbf{a}}$ and \mathbf{b} , the inner product of vectors \mathbf{a} and \mathbf{b} in the exponent, denoted by $g^{(\mathbf{a}, \mathbf{b})}$, is

$$\prod_{i=1}^n (g^{a_i})^{b_i} = \prod_{i=1}^n g^{a_i \cdot b_i} = g^{\sum_{i=1}^n a_i \cdot b_i} = g^{(\mathbf{a}, \mathbf{b})} .$$

Likewise, a matrix-vector product in the exponent, given a vector \mathbf{v} and $g^{\mathbf{A}}$ for a matrix $\mathbf{A} = (\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_n)$ can be performed by $\prod_{i=1}^n (g^{\mathbf{a}_i})^{v_i} = \prod_{i=1}^n g^{\mathbf{a}_i \cdot v_i} = g^{\sum_{i=1}^n \mathbf{a}_i \cdot v_i} = g^{\mathbf{A}\mathbf{v}}$. Adding (and subtracting) in the exponent is computed via element-wise multiplication (and division) of the group elements $g^{\mathbf{a}} \cdot g^{\mathbf{b}} = g^{\mathbf{a}+\mathbf{b}}$.

Quadratic Residuosity. The Legendre symbol verifies whether an integer $a \in \mathbb{Z}_p$ is a quadratic residue modulo a prime p , i.e., $x^2 \equiv a \pmod p$ for some x . If $\mathbb{L}(a, p) := a^{(p-1)/2} = 1$, this is the case; otherwise $\mathbb{L}(a, p) = -1$. More generally, for $n \geq 2$, we define $\mathbb{L}(a, p)_n := a^{(p-1)/\gcd(n, p-1)}$. If the modulus N is of the form $N = p_1 \cdots p_k$ where the p_i are odd primes, one uses its generalization, namely the Jacobi symbol, which is defined as $\mathbb{J}(a, N) = \prod_{i=1}^k \mathbb{L}(a, p_i)$. Note that $\mathbb{J}(a, N) = 1$ does not imply that a is a quadratic residue modulo N . However, if $\mathbb{J}(a, N) = -1$, a is certainly not. The set of quadratic residues modulo N is denoted by $\mathbb{QR}_N := \{a^2 : a \in \mathbb{Z}_N^*\}$. By \mathbb{J}_N we denote the subgroup of all elements from \mathbb{Z}_N^* with Jacobi symbol 1, i.e., $\mathbb{J}_N = \{a \in \mathbb{Z}_N^* : \mathbb{J}(a, N) = 1\}$. Note that \mathbb{QR}_N is a subgroup of \mathbb{J}_N . It is widely believed that one cannot efficiently decide whether an element $a \in \mathbb{J}_N$ is a quadratic residue modulo N if the prime factors of N are unknown (For more details, full version).

2.2 Standard Group-Theoretic Problems

We will make use of the rank hiding assumption introduced by Naor and Segev [14] (and later extended by Agrawal et al. [15]).³ It was proven to be equivalent to the $\text{DDH}_{\mathbb{G}, \chi}$ assumption for groups of prime order and uniform χ [14].

Definition 1 (Rank Hiding). *Let \mathbb{G} be a group of order q with generator g , and $i, j, n, m \in \mathbb{N}$ satisfying $i \neq j$ and $i, j \geq 1$. The Rank Hiding problem $(\text{RH}_{\mathbb{G}, i, j, m, n})$ is (t, ϵ) -hard if*

$$\{(\mathbb{G}, q, g, g^{\mathbf{M}}) : \mathbf{M} \leftarrow_R \text{Rk}_i(\mathbb{Z}_q^{m \times n})\} \approx_{(t, \epsilon)} \{(\mathbb{G}, q, g, g^{\mathbf{M}}) : \mathbf{M} \leftarrow_R \text{Rk}_j(\mathbb{Z}_q^{m \times n})\}$$

³ The assumption was first introduced by Boneh et al. [16] under the Matrix DDH assumption.

where $\text{Rk}_k(\mathbb{Z}_q^{m \times n})$ returns an $m \times n$ matrix uniformly random from $\mathbb{Z}_q^{n \times m}$ with rank $k \leq \min(n, m)$.

2.3 Representation Problem

The representation problem in a group \mathbb{G} assumes that given l random group elements $g_1, \dots, g_l \in \mathbb{G}$ and $h \in \mathbb{G}$ it is hard to find a representation $\mathbf{x} \in \mathbb{Z}_q^\ell$ such that $h = \prod_{i=1}^\ell g_i^{x_i}$ holds. Brands shows an electronic cash system based on the problem. Recently, the assumption was extensively applied to show leakage resiliency [15, 17, 18].

We now state a more general version of the search representation problem where vector $\mathbf{x} \leftarrow_R \chi^\ell$ is sampled from a distribution χ with (at least) min-entropy and where an adversary is given $m \geq 1$ samples instead of a single one.

Definition 2 (Search Representation Problem). *Let χ be a distribution over \mathbb{Z}_q , and ℓ, m be integers. Sample $\mathbf{M} \leftarrow_R \mathbb{Z}_q^{m \times \ell}$ and $\mathbf{x} \leftarrow_R \chi^\ell$. The Search Representation Problem ($\text{SRP}_{\mathbb{G}, \chi, \ell, m}$) is (t, ϵ) -hard if any algorithm \mathcal{A} , running in time t , upon input $(g, g^{\mathbf{M}}, g^{\mathbf{x}}, g^{\mathbf{M}\mathbf{x}})$, outputs $\mathbf{x}' \in \mathbb{Z}_q^\ell$ such that $g^{\mathbf{M}\mathbf{x}'} = g^{\mathbf{M}\mathbf{x}}$ with probability at most ϵ . If χ is the uniform distribution, we sometimes skip χ in the index and say that $\text{SRP}_{\mathbb{G}, \ell, m}$ is (t, ϵ) -hard.*

Brands proves the equivalence of the representation problem and the discrete logarithm problem for uniform χ and $m = 1$. It is easy to verify that the reduction holds for every distribution for which the discrete logarithm problem holds.

To establish relations to the learning with errors in the exponent problem (cf. Sect. 3.2), we need a decisional variant of the representation problem. To our surprise, the decisional version has not been defined before, although the assumption is a natural generalization of the decisional Diffie-Hellman problem to ℓ -tuples (similar in spirit as the ℓ -linear problem in \mathbb{G} [11]). Given ℓ random group elements $g_1, \dots, g_\ell \in \mathbb{G}$ together with $h \in \mathbb{G}$ and $g^{x_1}, \dots, g^{x_\ell} \in \mathbb{G}$ where $x_1, \dots, x_\ell \leftarrow_R \mathbb{Z}_q^*$, it is hard to decide if $h = \prod_{i=1}^\ell g_i^{x_i}$ or h is a random group element in \mathbb{G} . Our definition below generalizes this problem to the case, where $m \geq 1$ samples are given to an adversary and x_1, \dots, x_ℓ are sampled from any min-entropy distribution χ .

Definition 3 (Decisional Representation Problem). *Let χ be a distribution over \mathbb{Z}_q^* , and ℓ, m be integers. Sample $\mathbf{M} \leftarrow_R \mathbb{Z}_q^{m \times \ell}$, $\mathbf{h} \leftarrow_R \mathbb{Z}_q^m$, and $\mathbf{x} \leftarrow_R \chi^\ell$. The Decisional Representation ($\text{DRP}_{\mathbb{G}, \chi, \ell, m}$) problem is (t, ϵ) -hard if*

$$(g, g^{\mathbf{M}}, g^{\mathbf{x}}, g^{\mathbf{M}\mathbf{x}}) \approx_{(t, \epsilon)} (g, g^{\mathbf{M}}, g^{\mathbf{x}}, g^{\mathbf{h}}).$$

If χ is the uniform distribution over \mathbb{Z}_q^ , we say $\text{DRP}_{\mathbb{G}, \ell, m}$ is (t, ϵ) -hard.*

Remark 1. $\text{DRP}_{\mathbb{G}, \chi, \ell, m}$ can be stated in the framework of the Matrix-DDH assumption recently introduced by Escala et al. [19] and thus we put another class of hardness problems to the arsenal of their expressive framework.

We now give evidence that the family of $\text{DRP}_{\mathbb{G},\chi,\ell,m}$ problems is a class of progressively harder problems (with increasing ℓ). Proofs of following propositions can be found in the full version.

Proposition 1. *If $\text{DRP}_{\mathbb{G},\chi,\ell,m}$ is (t, ϵ) -hard, then for any $\ell, m \geq 1$ with $t' \approx t$ and distribution χ with min-entropy $\text{DRP}_{\mathbb{G},\chi,\ell+1,m}$ is (t', ϵ) -hard.*

Proposition 2. *In the generic group model $\text{DRP}_{\mathbb{G},\chi,\ell+1,m}$ is hard for distribution χ with minimal entropy, even in presence of a $\text{DRP}_{\mathbb{G},\chi,\ell,m}$ -oracle.*

Remark 2. $\text{DRP}_{\mathbb{G},\chi,1,1}$ -problem with χ being the uniform distribution over \mathbb{Z}_q coincides with the decisional Diffie-Hellman (DDH) problem. Hence, we obtain the corollary that for uniform distributions χ , the decisional Diffie-Hellman problem implies the representation problem $\text{DRP}_{\mathbb{G},\chi,\ell,1}$ for $\ell \geq 1$. In fact, Proposition 1 suggests a stronger argument. Assuming the decisional Diffie-Hellman problem holds for well-spread and min-entropy distributions χ , then the $\text{DRP}_{\mathbb{G},\chi,\ell,1}$ holds for χ and $\ell \geq 1$.

While Propositions 1 and 2 show that the DRP problem progressively increases with ℓ , the following proposition states that the problem remains hard with increasing number of samples m . More precisely, we show that $\text{DRP}_{\mathbb{G},\chi,\ell,m+1}$ is hard as long as $\text{DRP}_{\mathbb{G},\chi,\ell,m}$ and the Rank Hiding problem $\text{RH}_{\mathbb{G},m,m+1,m+1,2\ell+1}$ (cf. Definition 1) is hard. The proof is given in the full version.

Proposition 3. *If $\text{RH}_{\mathbb{G},m,m+1,m+1,2\ell+1}$ is (t, ϵ) -hard and $\text{DRP}_{\mathbb{G},\chi,\ell,m}$ is (t', ϵ') -hard in a cyclic group \mathbb{G} of order q , then for any distribution χ_e and any $m > 0$ with $t' \approx t$ and $\epsilon'' \leq (1 - \epsilon)^{-1}\epsilon'$ $\text{DRP}_{\mathbb{G},\chi,\ell,m+1}$ is (t, ϵ'') -hard.*

2.4 Learning with Errors

The learning with errors assumption comes as a search and decision lattice problem. Given a system of m linear equations with random coefficients $\mathbf{a}_i \in \mathbb{Z}_q^n$ in the n indeterminates \mathbf{s} sampled from some distribution χ_s and biased with some error e_i from the error distribution χ_e , it is hard to compute vector \mathbf{s} or distinguish the solution $b_i = \sum_i^n \mathbf{a}_i \mathbf{s} + e_i$ from a uniform element in \mathbb{Z}_q .

Definition 4 (Learning with Errors). *Let n, m, q be integers and χ_e, χ_s be distributions over \mathbb{Z} . For $\mathbf{s} \leftarrow_R \chi_s$, define the LWE distribution $L_{n,q,\chi_e}^{\text{LWE}}$ to be the distribution over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ obtained such that one first draws $\mathbf{a} \leftarrow_R \mathbb{Z}_q^n$ uniformly, $e \leftarrow_R \chi_e$ and returns $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ with $b = \langle \mathbf{a}, \mathbf{s} \rangle + e$. Let (\mathbf{a}_i, b_i) be samples from $L_{n,q,\chi_e}^{\text{LWE}}$ and $c_i \leftarrow_R \mathbb{Z}_q$ for $0 \leq i < m = \text{poly}(\kappa)$.*

- The Search Learning With Errors ($\text{SLWE}_{n,m,q,\chi_e}(\chi_s)$) problem is (t, ϵ) -hard if any algorithm \mathcal{A} , running in time t , upon input $(\mathbf{a}_i, b_i)_{i \in [m]}$, outputs \mathbf{s} with probability at most ϵ .
- The Decisional Learning with Errors ($\text{DLWE}_{n,m,q,\chi_e}(\chi_s)$) problem is (t, ϵ) -hard if

$$(\mathbf{a}_i, b_i)_{i \in [m]} \approx_{(t, \epsilon)} (\mathbf{a}_i, c_i)_{i \in [m]}$$

for a random secret $\mathbf{s} \leftarrow_R \chi_s$.

If χ_s is the uniform distribution over \mathbb{Z}_q , we simply write $\text{LWE}_{n,m,q,\chi_e}$.

A typical distribution for the error is a discrete Gaussian distribution with an appropriate standard deviation. There are several proposals for the distribution of the secret. While the uniform distribution is the most standard one, it is shown that setting $\chi_s = \chi_e$, known as the “normal form”, retains the hardness of LWE [20, 21]. We also note that the learning with errors problem where the error is scaled by a constant α relatively prime to q is as hard as the original definition [22]. The “scaled” LWE distribution then returns (\mathbf{a}, b) with $\mathbf{a} \leftarrow_R \mathbb{Z}_q^n$ and $b = \langle \mathbf{a}, \mathbf{s} \rangle + \alpha e$.

3 Learning with Errors in the Exponent

3.1 Definition

For self-containment, the assumption is stated both as a search and decision problem over a group \mathbb{G} of order q , and exponents sampled from distributions χ_e, χ_s over \mathbb{Z} . We demonstrate the versatility and general utility of the decisional version in Sect. 4.

Definition 5 (Learning with Errors in the Exponent). *Let \mathbb{G} be a group of order q where g is a generator of \mathbb{G} . Let n, m, q be integers and χ_e, χ_s be distributions over \mathbb{Z} . For any fixed vector $\mathbf{s} \in \mathbb{Z}_q^n$, define the LWEE distribution $L_{\mathbb{G}, n, q, \chi_e}^{\text{LWEE}}$ to be the distribution over $\mathbb{G}^n \times \mathbb{G}$ obtained such that one first draws vector $\mathbf{a} \leftarrow_R \mathbb{Z}_q^n$ uniformly, $e \leftarrow_R \chi_e$ and returns $(g^{\mathbf{a}}, g^b) \in \mathbb{G}^n \times \mathbb{G}$ with $b = \langle \mathbf{a}, \mathbf{s} \rangle + e$. Let $(g^{\mathbf{a}^i}, g^{b_i})$ be samples from $L_{\mathbb{G}, n, q, \chi_e}^{\text{LWEE}}$ and c_i be uniformly sampled from \mathbb{Z}_q^* for $0 \leq i < m = \text{poly}(\kappa)$.*

- *The Search Learning With Errors in the Exponent (SLWEE $_{\mathbb{G}, n, m, q, \chi_e}(\chi_s)$) problem is (t, ϵ) -hard if any algorithm \mathcal{A} , running in time t , upon input $(g^{\mathbf{a}^i}, g^{b_i})_{i \in [m]}$, outputs \mathbf{s} with probability at most ϵ .*
- *The Decision Learning With Errors in the Exponent (DLWEE $_{\mathbb{G}, n, m, q, \chi_e}(\chi_s)$) problem is (t, ϵ) -hard if $(g^{\mathbf{a}^i}, g^{b_i})_{i \in [m]} \approx_{(t, \epsilon)} (g^{\mathbf{a}^i}, g^{c_i})_{i \in [m]}$ for a random secret $\mathbf{s} \leftarrow_R \chi_s^n$. If χ_s is the uniform distribution over \mathbb{Z}_q , we write DLWEE $_{\mathbb{G}, n, m, q, \chi_e}$.*

We let $\text{Adv}_{\mathbb{G}, n, m, q, \chi_e, \chi_s}^{\text{DLWEE/SLWEE}}(t)$ denote a bound on the value ϵ for which the decisional/search LWEE problem is (t, ϵ) -hard.

One may interpret learning with errors in the exponent in two ways. One way is to implant an error term from a distribution χ_e into the Diffie-Hellman exponent. Another way to look at LWEE is as compressing an LWE instance within some group \mathbb{G} of order q .

3.2 Relations to Group and Lattice Problems

We connect the representation and learning with errors problem to learning with errors in the exponent. The essence is that there exist tight reductions from the search (resp. decision) learning with errors in the exponent problem to

either the search (resp. decision) representation problem and the search (resp. decision) learning with errors problem. This has several interesting property preserving implications. As a corollary we infer that for appropriate parameter choices LWEE preserves the *hardness* and *robustness* properties of the representation and/or learning with errors problem. Essentially then LWEE boils down to the security of either of the two underlying problems. This way, the cryptosystem can be instantiated to leverage leakage resistance and post-quantum hardness thanks to LWEE [3, 23]. On the flip side, the cryptosystem may offer short instance sizes through the underlying RP problem (when instantiated on elliptic curves). Of particular interest for many emerging applications is the partnering of the two hardness assumptions. One may choose parameters such that both RP and LWEE hold. We call the case *double-hard*, which appeals to provide in some sense hedged security.

Following four propositions summarize our main results. Proofs appear in the full version.

Proposition 4. *If $\text{SRP}_{\mathbb{G}, \chi_s, \ell, m}$ is (t, ϵ) -hard in a cyclic group \mathbb{G} of order q , then for any distribution χ_e and any number of samples $m > 0$ $\text{SLWEE}_{\mathbb{G}, \ell, m, q, \chi_e}(\chi_s)$ is (t', ϵ) -hard with $t' \approx t$.*

Proposition 5. *If $\text{SLWE}_{n, m, q, \chi_e}(\chi_s)$ is (t, ϵ) -hard, then for any cyclic group \mathbb{G} of order q with known (or efficiently computable) generator $\text{SLWEE}_{\mathbb{G}, n, m, q, \chi_e}(\chi_s)$ is (t', ϵ) -hard with $t' \approx t$.*

Proposition 6. *If $\text{DRP}_{\mathbb{G}, \chi_s, \ell, m}$ is (t, ϵ) -hard in a cyclic group \mathbb{G} of order q , then for any distribution χ_e and any number of samples $m > 0$ $\text{DLWEE}_{\mathbb{G}, \ell, m, \chi_e}(\chi_s)$ is (t', ϵ) -hard with $t' \approx t$.*

Proposition 7. *If $\text{DLWE}_{n, m, q, \chi_e}(\chi_s)$ is (t, ϵ) -hard, then for any cyclic group \mathbb{G} of order q with known (or efficiently computable) generator $\text{DLWEE}_{\mathbb{G}, n, m, \chi_e}(\chi_s)$ is (t', ϵ) -hard with $t' \approx t$.*

3.3 On the Generic Hardness of LWEE

With Propositions 4–7 in our toolbox we conjecture LWEE to be harder than either of the underlying RP or LWE problems. The argument is heuristic and based on what is known about the hardness of each intractability problem (see full version for more details).

Fix parameters such that RP and LWE problem instances give κ bits security. The only obvious known approach today to solve the LWEE instance is to first compute the discrete logarithm of samples (g^{a_i}, g^{b_i}) and then solve the LWE problem for samples (\mathbf{a}_i, b_i) . Note that an adversary must solve $n^2 + n$ many discrete logarithms because the secret vector \mathbf{s} is information-theoretically hidden, if less than n samples of LWE are known. Solving $N := n^2 + n$ discrete logarithms in generic groups of order q takes time $\sqrt{2Nq}$ while computing a single discrete logarithm takes time $\sqrt{\pi q/2}$ [24, 25].⁴ In fact, this bound is proven to be optimal

⁴ Solving N -many discrete logarithms is easier than applying N times a DL solver for a single instance.

in the generic group model [26]. Note, parameters for LWEE are chosen such that computing a single discrete logarithm takes time 2^κ . Hence, in order to solve the LWEE instance for $N = \mathcal{O}(\kappa^2)$, one requires time $\frac{2}{\sqrt{\pi}}\sqrt{N} \cdot 2^\kappa + 2^\kappa > 2^{\kappa+2\log(\kappa)}$. This shows that generically the concrete instance of LWEE is logarithmically harder in the security parameter κ .

4 Public-Key Encryption from LWEE

4.1 The High-Level Idea

The idea behind our scheme is reminiscent of Regev’s public-key encryption scheme. In a nutshell, the public key is an LWEE instance $(g^{\mathbf{A}}, g^{\mathbf{A}\mathbf{s}+\mathbf{x}}) \in \mathbb{G}^{n \times n} \times \mathbb{G}^n$. Similarly to [27, 28] and as opposed to Regev [3], for efficiency reason we avoid the use of the leftover hash lemma –instead we impose one further LWEE instance– and make use of a square matrix A . Ciphertexts consist of two LWEE instances $C = (\mathbf{c}_0, c_1)$ where $\mathbf{c}_0 = g^{\mathbf{A}\mathbf{r}+\mathbf{e}_0}$ encapsulates a random key $\mathbf{r} \in \mathbb{Z}_q^n$ and $c_1 = g^{(\mathbf{b}, \mathbf{r})+\mathbf{e}_1} \cdot g^{\alpha\mu}$ encrypts the message μ (we discuss the exact value of α below). The tricky part is the decryption algorithm. All known LWE-based encryption schemes require some technique to deal with the noise terms. Otherwise, decryption is prone to err. Regev’s technique ensures small error terms. One simply rounds $c_1 - \mathbf{c}_0\mathbf{s}$ to some reference value c_b indicating the encryption of bit b . While rounding splendidly works on integers, the technique fails in our setting.

Our approach explores a considerably different path. Instead of rounding, we synthesize the pesky error terms. To this end, we adapt the trapdoor technique of Joye and Libert [29] and recover partial bits of the discrete logarithm (by making use of the Pohlig-Hellman algorithm [30]). The main idea is to tweak the modulus in a smart way. Given composite modulus $N = pq$ with p', q' , such that $p = 2^k p' + 1$ and $q = 2^k q' + 1$ are prime, there exists an efficient algorithm for recovering the k least significant bits of the discrete logarithm. We choose the parameters so that the sum of all error terms in the exponent is (with high probability) at most $2^{k-\ell}$. This leads to a “gap” between error bits and those bits covered by the discrete log instance. We plant the message in this gap by shifting it to the $2^{k-\ell}$ ’s bit, where ℓ is the size of the message we want to decrypt. Hence, we choose $\alpha = 2^{k-\ell}$ in our construction to shift the message bits accordingly. We leave it as an interesting open problem to instantiate the scheme in prime order groups.

4.2 Our Construction

The scheme is parameterized by positive integers $n, k, \ell < k$ and Gaussian parameters σ_s, σ_e .

KeyGen: Sample prime numbers p' and q' , such that $p = 2^k p' + 1$ and $q = 2^k q' + 1$ are prime. Set $N = pq$ and $M = 2^k p' q'$. Sample $\mathbf{s} \leftarrow_R \mathcal{D}_{\sigma_s}^n$, $\mathbf{A} \leftarrow_R \mathbb{Z}_M^{n \times n}$ and $\mathbf{x} \leftarrow_R \mathcal{D}_{\sigma_e}^n$ and compute $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{x}$. Sample $g \in \mathbb{J}_N \setminus \mathbb{QR}_N$ of order M . The public key consists of $\text{pk} = (g, g^{\mathbf{A}}, g^{\mathbf{b}}, N)$, and the secret key of $\text{sk} = (p, \mathbf{s})$.

Algorithm 1.

Input: Generator g of a group with order $p - 1 = 2^k p'$, p and k
Output: k least significant bits of $\log_g(h)$

```

begin
   $a = 0, B = 1;$ 
  for  $i \in \{1, \dots, k\}$  do
     $z \leftarrow \mathbb{L}(h, p)_{2^i} \bmod p;$ 
     $t \leftarrow \mathbb{L}(g, p)_{2^i}^a \bmod p;$ 
    if  $z \neq t$  then
       $a \leftarrow a + B;$ 
    end
     $B \leftarrow 2B;$ 
  end
  return  $a$ 
end
    
```

Encrypt(pk, μ): To encrypt ℓ bits $\mu \in \{0, 1, \dots, 2^\ell - 1\}$ given public key pk choose $\mathbf{r} \leftarrow_R \mathcal{D}_{\sigma_s}^n$, $\mathbf{e}_0 \leftarrow_R \mathcal{D}_{\sigma_e}^n$ and $e_1 \leftarrow_R \mathcal{D}_{\sigma_e}$. Use $g^{\mathbf{A}}$, \mathbf{r} and \mathbf{e}_0 to compute $g^{\mathbf{A}\mathbf{r} + \mathbf{e}_0}$, and $g^{\mathbf{b}}$, \mathbf{r} and e_1 to compute $g^{\langle \mathbf{b}, \mathbf{r} \rangle + e_1}$. The ciphertext is \mathbf{c}_0, c_1 with

$$\mathbf{c}_0 = g^{\mathbf{A}\mathbf{r} + \mathbf{e}_0}, \quad c_1 = g^{\langle \mathbf{b}, \mathbf{r} \rangle + e_1} \cdot g^{2^{k-\ell} \mu}.$$

Decrypt(sk, (\mathbf{c}_0, c_1)): To decrypt the ciphertext (\mathbf{c}_0, c_1) given secret key $\text{sk} = (p, \mathbf{s})$, first compute $g^{\langle \mathbf{s}, \mathbf{A}\mathbf{r} + \mathbf{e}_0 \rangle}$ and then $h = c_1 / g^{\langle \mathbf{s}, \mathbf{A}\mathbf{r} + \mathbf{e}_0 \rangle}$. Run Algorithm 1 to synthesize $v = \log_g(h) \bmod 2^k$ and return $\lfloor \frac{v}{2^{k-\ell-1}} \rfloor$.

4.3 Correctness

To show correctness of our construction we build upon two facts. First, Algorithm 1 synthesizes the k least significant bits of a discrete logarithm. The algorithm's correctness for a modulus being a multiple of 2^k is proven in [29, Section 3.2]. Second, noise in the exponent does not overlap with the message. To this end, we bound the size of the noise with following lemma.

Lemma 1 (adapted from [28, Lemma 3.1]). *Let c, T be positive integers such that*

$$\sigma_s \cdot \sigma_e \leq \frac{\pi}{c} \frac{T}{\sqrt{n \ln(2/\delta)}} \quad \text{and} \quad \left(c \cdot \exp\left(\frac{1-c^2}{2}\right) \right)^{2n} \leq 2^{-40}.$$

For $\mathbf{x}, \mathbf{s} \leftarrow_R \mathcal{D}_{\sigma_s}^n$, $\mathbf{r}, \mathbf{r}_0 \leftarrow_R \mathcal{D}_{\sigma_e}^n$, $e_1 \leftarrow_R \mathcal{D}_{\sigma_e}$, we have $|\langle \mathbf{x}, \mathbf{r} \rangle - \langle \mathbf{s}, \mathbf{e}_0 \rangle + e_1| < T$ with probability at least $1 - \delta - 2^{-40}$.

We are now ready to prove the following theorem.

Theorem 1. *Let c, T be as in Lemma 1. Then, the decryption is correct with probability at least $1 - \delta - 2^{-40}$.*

4.4 Ciphertext Indistinguishability

Theorem 2. *Let $\mathbb{G} = \langle g \rangle$ be the cyclic group of composite order generated by g . If the decisional LWEE problem $\text{DLWEE}_{\mathbb{G},n,n+1,q,D_{\sigma_e}}(D_{\sigma_s})$ is (t, ϵ) -hard, then the above cryptosystem is $(t, 2\epsilon)$ -indistinguishable against chosen plaintext attacks.*

Proof. In a high level, our proof works as follows. Instead of showing IND-CPA security via a direct argument we show that the distribution $(\mathbf{pk}, \mathbf{c}_0, c_1)$ is indistinguishable from the uniform distribution over $(\mathbb{G}^{n \times n} \times \mathbb{G}^{2n+1})$. That is, a ciphertext (\mathbf{c}_0, c_1) under public key \mathbf{pk} appears completely random to an adversary. This holds, in particular, in the IND-CPA experiment when the adversary chooses the underlying plaintext. We prove the theorem via a series of hybrid arguments, Hybrid_0 to Hybrid_2 , where in each consecutive argument we make some slight changes with the provision that the adversary notices the changes with negligible probability only. In the following, we use the abbreviations $\mathbf{u} = \mathbf{A}\mathbf{r} + \mathbf{e}_0$ and $v = \langle \mathbf{b}, \mathbf{r} \rangle + e_1 + 2^{k-\ell}\mu$.

Hybrid₀: In this hybrid we consider the original distribution of the tuple

$$(\mathbf{pk}, (\mathbf{c}_0, \mathbf{c}_1)) = (g^{\mathbf{A}}, g^{\mathbf{b}}, g^{\mathbf{u}}, g^v).$$

Hybrid₁: In this hybrid we modify the distribution and claim

$$(g^{\mathbf{A}}, g^{\mathbf{b}}, g^{\mathbf{u}}, g^v) \approx_c (g^{\mathbf{A}'}, g^{\mathbf{b}'}, g^{\mathbf{A}'\mathbf{r}+\mathbf{e}_0}, g^{\langle \mathbf{b}', \mathbf{r} \rangle + e_1} \cdot g^{2^{k-\ell}\mu})$$

for a uniformly sampled elements $g^{\mathbf{A}'}, g^{\mathbf{b}'} \in \mathbb{G}^{n \times n} \times \mathbb{G}^n$. We argue that any successful algorithm distinguishing between Hybrid_0 and Hybrid_1 can be easily turned into a successful distinguisher \mathcal{B} in the $\text{DLWEE}_{\mathbb{G},n,n,q,D_{\sigma_e}}(D_{\sigma_s})$ problem. The DLWEE-adversary \mathcal{B} is given as challenge the tuple $(g^{\mathbf{A}'}, g^{\mathbf{b}'})$ and is asked to decide whether there exist vectors $\mathbf{s} \leftarrow_R D_{\sigma_s}$, $\mathbf{x} \leftarrow_R D_{\sigma_e}^n$ such that $g^{\mathbf{b}'} = g^{\mathbf{A}'\mathbf{T}\mathbf{s}+\mathbf{x}}$ or $g^{\mathbf{b}'}$ was sampled uniformly from \mathbb{G}^n .

Let $\Pr[\text{Hybrid}_i(t)]$ denote the probability of any algorithm with runtime t to win the IND-CPA experiment in hybrid i . Then, we have

$$\Pr[\text{Hybrid}_0(t)] \leq \Pr[\text{Hybrid}_1(t)] + \text{Adv}_{\mathbb{G},n,n,q,D_{\sigma_e},D_{\sigma_s}}^{\text{DLWEE}}(t).$$

Hybrid₂: In this hybrid we modify the distribution and claim

$$(g^{\mathbf{A}'}, g^{\mathbf{b}'}, g^{\mathbf{A}'\mathbf{r}+\mathbf{e}_0}, g^{\langle \mathbf{b}', \mathbf{r} \rangle + e_1} \cdot g^{2^{k-1}\mu}) \approx_c (g^{\mathbf{A}''}, g^{\mathbf{b}''}, g^{\mathbf{u}'}, g^{v'} \cdot g^{2^{k-1}\mu})$$

for a uniformly sampled elements $g^{\mathbf{A}''}, g^{\mathbf{b}''}, g^{\mathbf{u}'}, g^{v'} \cdot g^\mu \in \mathbb{G}^{(n+1) \times n} \times \mathbb{G}^{n+1}$. We argue that any successful algorithm distinguishing between Hybrid_1 and Hybrid_2 can be easily turned into a successful distinguisher \mathcal{B} against the $\text{DLWEE}_{\mathbb{G},n,n+1,q,D_{\sigma_e}}(D_{\sigma_s})$ problem. Note that $g^{\mathbf{b}'}, g^{\langle \mathbf{b}', \mathbf{r} \rangle + e_1}$ is an additional sample from the LWEE distribution from which $g^{\mathbf{A}'}, g^{\mathbf{A}'\mathbf{r}+\mathbf{e}_0}$ is sampled.

We have

$$\Pr[\text{Hybrid}_1(t)] \leq \Pr[\text{Hybrid}_2(t)] + \text{Adv}_{\mathbb{G},n,n+1,q,D_{\sigma_e},D_{\sigma_s}}^{\text{DLWEE}}(t).$$

Note that now all exponents are uniformly distributed, and, in particular, independent of μ and thus, independent of b in the IND-CPA game. Hence, any algorithm has in Hybrid_2 exactly a success probability of $1/2$.

This completes the proof of semantic security.

4.5 Candidate Instantiations of Our Encryption Scheme

We give three possible instantiations to derive a system with short key sizes, post-quantum security or double hardness. Throughout this section we instantiate our scheme such that the encryption scheme from Sect. 4.2 encrypts only a single bit. Nonetheless, parameters can easily be upscaled to many bits.

Table 1. Key sizes in kilobytes (kB) for our encryption scheme basing security on DRP or LWE, respectively.

Sizes/Security	DRP-based instantiation			LWE-based instantiation		
	80-bit	128-bit	256-bit	80-bit	128-bit	256-bit
Public-key size	0.565 kB	1.500 kB	7.500 kB	235 kB	417 kB	1233 kB
Secret-key size	0.212 kB	0.563 kB	2.813 kB	0.976 kB	1.302 kB	2.237 kB
Ciphertext size	0.283 kB	0.750 kB	3.750 kB	0.980 kB	1.306 kB	2.241 kB

The Classical Way. Here, we instantiate our encryption scheme such that the underlying DRP is intractable, and neglecting the hardness of the underlying LWE. In the full version, we recall some groups where we believe DRP is hard to solve. Our encryption scheme works in the group $\mathbb{J}_N := \{x \in \mathbb{Z}_N : \mathbb{J}(x, N) = 1\}$ for $N = pq$ with p, q being k -safe primes. In fact, we can even take safe primes p, q (i.e., $k = 1$) since we do not need any noise in the exponent if we neglect the underlying LWE hardness. Thus, we embed the message to the least significant bit in the exponent. For this reason, we can sample $g \leftarrow_R \mathbb{J}_N / \mathbb{Q}\mathbb{R}_N$ where $\langle g \rangle$ has order $2p'q'$. Since the LWE instance within LWEE is not an issue here we select $n = m = 1$, $\sigma_s = \infty$ and $\sigma_e = 0$.

We obtain 80-bit security for the underlying DRP problem if we choose safe primes p and q such that $\log p = \log q = 565$ (see full version for more details). Table 1 lists possible key sizes for our encryption scheme. Recall that the public key consists of $\text{pk} = (g, g^{\mathbf{A}}, g^{\mathbf{b}}, k, N)$ (i.e., 4 group elements if we fix $k = 1$) and the secret key of $\text{sk} = (p, \mathbf{s})$.

The Post-Quantum Way. Here we give example instantiations of our encryption scheme when it is based on a presumably quantum-resistant LWEE assumption. That is, we select parameters such that the underlying LWE assumption is intractable without relying on the hardness of DRP. For this, we modify the scheme slightly by choosing fixed values for p' and q' instead of sampling. A good choice is $k = 15$, since it allows to choose $p' = 2$ and $q' = 5$, which are very small prime numbers such that $2^k p' + 1$ and $2^k q' + 1$ are prime. For the LWE modulus, this leads to $M = 2^k p' q' = 327680$. Like Lindner and Peikert [28], we choose the Gaussian parameter such that the probability of decoding errors is bounded by 1%. We choose furthermore the same parameter for error and secret distribution (i.e. $\sigma_s = \sigma_e = \sigma$), since a standard argument reduces LWE with arbitrary secret

to LWE with secret chosen according to the error distribution. For this choice of k , p' and q' , we obtain 80-bit security by choosing $n = 240$ and $\sigma = 33.98$. Table 1 lists the key sizes when our encryption scheme is instantiated such that its security is based on LWE only (see full version for more information about the concrete hardness of LWE).

The Hardest Way (Double-Hardness). The most secure instantiation of our encryption is such that even if one of the problems DRP or LWE is efficiently solvable at some point, our encryption scheme remains semantically secure. Selecting parameters for double hardness, however, is non-trivial.

To select appropriate parameters for the case of double hardness, we apply the following approach: For a given security level (say $\kappa = 80$), we select N such that the Number Field Sieve needs at least 2^κ operations to factor N . A possible choice is $\log N = 1130$ (See full version). Since factoring N must also be hard for McKee-Pinch's algorithm, which works well when $(p - 1)$ and $(q - 1)$ share common factor, k must be chosen such that $N^{1/4}2^{-k} \geq 2^\kappa$, i.e. $k \leq \frac{\log(N)}{4} - \kappa$. This leads to $k = 203$. Given N and k , we can calculate the sizes of the primes $\log(p') \approx \log(q') \approx 362$ and $\log(p) \approx \log(q) \approx 565$ and the LWE modulus $\log(M) \approx 927$. Taking $n = 67000$ and $\sigma = 2^{97}$, Lemma 1 shows that the algorithm decrypts correctly with high probability.

References

1. Shor, P.: Algorithms for quantum computation: discrete logarithms and factoring. In: 35th Annual Symposium on Foundations of Computer Science, 1994 Proceedings, pp. 124–134 (1994)
2. Shor, P.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**, 1484–1509 (1997)
3. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 84–93, 22–24 May 2005. ACM, Baltimore, MD, USA (2005)
4. Regev, O.: The learning with errors problem (invited survey). In: IEEE Conference on Computational Complexity, pp. 191–204. IEEE Computer Society (2010)
5. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, pp. 333–342. ACM, New York (2009)
6. Lyubashevsky, V., Micciancio, D.: On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 577–594. Springer, Heidelberg (2009)
7. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the 45th Annual ACM Symposium on Symposium on Theory of Computing, pp. 575–584. ACM (2013)
8. Katz, J., Sahai, A., Waters, B.: Predicate encryption supporting disjunctions, polynomial equations, and inner products. *J. Cryptology* **26**, 191–224 (2013)
9. Boneh, D., Goh, E.-J., Nissim, K.: Evaluating 2-DNF formulas on ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Heidelberg (2005)

10. Brands, S.A.: An efficient off-line electronic cash system based on the representation problem. Technical report, Amsterdam, The Netherlands (1993)
11. Shacham, H.: A Cramer-Shoup encryption scheme from the linear assumption and from progressively weaker linear variants. Cryptology ePrint Archive, Report 2007/074 (2007). <http://eprint.iacr.org/>
12. Dagdelen, O., Gajek, S., Gopfert, F.: Learning with errors in the exponent. Cryptology ePrint Archive, Report 2014/826 (2014). <http://eprint.iacr.org/>
13. Regev, O.: Quantum computation and lattice problems. *SIAM J. Comput.* **33**, 738–760 (2004)
14. Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009)
15. Agrawal, S., Dodis, Y., Vaikuntanathan, V., Wichs, D.: On continual leakage of discrete log representations. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 401–420. Springer, Heidelberg (2013)
16. Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision diffie-hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)
17. Katz, J., Vaikuntanathan, V.: Signature schemes with bounded leakage resilience. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 703–720. Springer, Heidelberg (2009)
18. Dagdelen, Ö., Venturi, D.: A second look at Fischlin’s transformation. In: Pointcheval, D., Vergnaud, D. (eds.) AFRICACRYPT. LNCS, vol. 8469, pp. 356–376. Springer, Heidelberg (2014)
19. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (2013)
20. Micciancio, D.: Improving lattice based cryptosystems using the Hermite normal form. In: Silverman, J.H. (ed.) CaLC 2001. LNCS, vol. 2146, pp. 126–145. Springer, Heidelberg (2001)
21. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)
22. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-LWE and security for key dependent messages. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 505–524. Springer, Heidelberg (2011)
23. Goldwasser, S., Kalai, Y., Peikert, C., Vaikuntanathan, V.: Robustness of the learning with errors assumption. In: ICS, 2010, [GPV08] [GRS08] (2008)
24. Kuhn, F., Struik, R.: Random walks revisited: extensions of Pollard’s Rho algorithm for computing multiple discrete logarithms. In: Vaudenay, S., Youssef, A.M. (eds.) SAC 2001. LNCS, vol. 2259, pp. 212–229. Springer, Heidelberg (2001)
25. Hitchcock, Y., Montague, P., Carter, G., Dawson, E.: The efficiency of solving multiple discrete logarithm problems and the implications for the security of fixed elliptic curves. *Int. J. Inf. Secur.* **3**, 86–98 (2004)
26. Yun, A.: Generic hardness of the multiple discrete logarithm problem. Cryptology ePrint Archive, Report 2014/637 (2014). <http://eprint.iacr.org/>
27. Lyubashevsky, V., Palacio, A., Segev, G.: Public-key cryptographic primitives provably as secure as subset sum. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 382–400. Springer, Heidelberg (2010)

28. Lindner, R., Peikert, C.: Better key sizes (and attacks) for LWE-based encryption. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 319–339. Springer, Heidelberg (2011)
29. Joye, M., Libert, B.: Efficient cryptosystems from 2^k -th power residue symbols. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 76–92. Springer, Heidelberg (2013)
30. Pohlig, S., Hellman, M.: An improved algorithm for computing logarithms over $\text{GF}(p)$ and its cryptographic significance (corresp.). IEEE Trans. Inf. Theory **24**, 106–110 (1978)