ERROR ANALYSIS OF WEAK POLY-LWE INSTANCES

YAO CHEN[1], BENJAMIN CASE[2], SHUHONG GAO[2], AND GUANG GONG[1]

Abstract. Error distribution plays a central role in the security of encryption based on the Learning with Errors (LWE) problem and its variants. In this paper, we investigate the error distribution of weak Poly-LWE instances. For this purpose, we derive a closed-form formula to compute the mapped error distribution. With this algebraic approach to evaluate the error, we examine the recently proposed attacks on Poly-LWE and Ring-LWE and reassess their parameters in order to include more instances. Notably, our method can also be applied to non-Gaussian error. We conduct experiments to investigate the shape of the mapped error distribution and confirm that in many cases it is no longer Gaussian nor uniform; our experimental results from distinguishers also validate our theoretical analysis.

Keywords: error distributuion, weak instances, Poly-LWE

1. Introduction

Since Regev proposed Learning with Errors (LWE) [Reg05, Reg09], it has found many applications in cryptography. It is conceptually simple but also enjoys worst-case hardness like some other lattice problems [Reg09, BLP+13]. Unfortunately, with the benefits of LWE usually come large keys and ciphertexts that add to communication and storage overhead. In the search of a better alternative that is more competitive in terms of key size, Ring-LWE (RLWE) emerged [LPR10] as a promising candidate. RLWE accomplishes its higher efficiency by exploiting additional algebraic structure of polynomial rings, but this approach is followed by attacks that can find weak instances, for the same reason. Similar to LWE, RLWE enjoys worst-case hardness, but only on ideal lattices.

Due to its efficiency, RLWE has been applied to many cryptographic constructions [BV11, FV12, BLLN13]; therefore, it has become more urgent to understand its strengths and weaknesses.

1.1. Related Work.

Since RLWE can be reduced to lattice and LWE problems, all lattice reduction attacks such as LLL and general decoding attacks such as BKW apply to RLWE. However, people generally believe and the weak-instance attacks described below show that the RLWE problem is potentially easier to tackle than the LWE problem because of the additional algebraic structure.

Weak instances of RLWE are analogous to weak primes in factorization: their special properties significantly reduce the difficulty of the problem so that specialized algorithms can be designed to launch an attack. Weak RLWE instances usually involve some ring homomorphism, under which the image of the error distribution can be distinguished from a uniform distribution.

Generally, weak instance attacks involve three steps:

(1) Exploiting the algebraic property to reduce the search space for Step (2);
(2) Exhausting the secret s;

[1]University of Waterloo, Waterloo, ON N2L 3G1, Canada
[2]Clemson University, Clemson, SC 29634, USA
Email address: y449chen@uwaterloo.ca (Yao Chen).

(3) Testing if the samples agree with a certain distribution generated with the guessed secret.

Considering different rings with special properties, leads to various weak instance attacks.

1.1.1. *Algebraic Structures.* The first such attack was developed by Eisenträger et al. [EHL14] on Polynomial-LWE (Poly-LWE), which is a special case of RLWE. They consider polynomial rings of the form $\mathbb{Z}_p[x]/\langle f(x)\rangle$, where p is prime and $f(x)$ is irreducible over $\mathbb{Q}$ but has a low-order root α modulo p. Their attack exploits a ring homomorphism induced by α into the finite field $\mathbb{F}_p$. When p is small enough, it becomes feasible to search for the image of secret s in $\mathbb{F}_p$.

Similar attacks were soon carried out on more general RLWE cases [ELOS15, CIV16]. The conditions for a possible attack are similar, but the attacker is faced with bigger distortions introduced by the conversion from RLWE instances to Poly-LWE ones in order for the original attack to work.

A more delicate attack was delivered by Chen et al. [CLS15, CLS16] on RLWE instances based on families of Galois number fields whose ring of integers can be decomposed into orthogonal subspaces, where their homomorphism will likely nullify a component of the error drawn from a discrete Gaussian distribution.

1.1.2. *Distinguishing Distributions.* In Step (3) of the attack, we need to decide which distribution fits the one computed from the guessed secret and samples better: uniform or the error distribution (under the homomorphism). All previous works achieve this by comparing the sampled distribution with uniform under the assumption that the mapped error distribution will be far enough apart from uniform.

For this particular task, there are general purpose distinguishers available such as the Chi-square test. Another distinguisher with dual lattices is considered by Peikert [Pei16] and can be used in conjunction with all existing weak instance attacks.

1.2. **Our Contribution.** In this paper, we first review the conditions for the attack in [ELOS15] to be launched. We show that the mapped error distribution can be precisely computed and therefore the restriction in their work can be relaxed to allow a broader range of instances to be attacked the same way. Then we show the mapped discrete Gaussian distribution, with different widths, according to calculation based on our method. We believe this work is the first to reveal the shape of exact mapped distributions, although estimations have been made repeatedly.

Having established the mapped distribution, we demonstrate how effective distinguishers can be built for weak instances recognized by our method, with or without information of the mapped distribution.

1.3. **Organization.** This paper is organized as follows: Section 2 summarizes the Poly-LWE problem and other related background. Section 3 shows the image of error distribution mapped under the homomorphism and introduces our method to compute it. Section 4 discusses different methods to distinguish distributions with samples and demonstrates our simulation results.

2. Background

Let $f(x)$ be a monic irreducible polynomial in $\mathbb{Z}[x]$ of degree n (not necessarily cyclotomic). Let $q \in \mathbb{Z}$ (not necessarily prime). If we let $p \in \mathbb{Z}$ it will always denote a prime. For the integers modulo q we use the notation $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z}$. We will be working in the following polynomial ring $\mathbb{Z}_q[x]/\langle f(x)\rangle$ which we denote as

$$R_q := \mathbb{Z}_q[x]/\langle f(x)\rangle.$$

Observe that there are q^n elements in R_q that are of the form

$$d_0 + d_1 x + d_2 x^2 + \cdots + d_{n-1} x^{n-1} + \langle f(x) \rangle$$

where $d_i \in \mathbb{Z}_q$ $0 \leq i \leq n-1$. Roots of $f(x) \bmod q$ will be denoted as $\alpha_1, ..., \alpha_i$ for as many roots as there are. If only one root is being considered, it may be denoted as just α. Since a root α is in $\mathbb{Z}_q$, we can talk about its order with respect to the multiplication in $\mathbb{Z}_q$,

$$\text{ord}(\alpha) := \min\{m \in \mathbb{Z}^+ | \alpha^m = 1\}.$$

In general we will not assume that $f(x)$ factors completely mod q; however, we will assume that $f(x)$ has at least one root α mod q.

On the ring R_q, we consider the following discrete probability distributions. Let ${}_{\mathcal{U}}R_q$ denote the discrete uniform distribution on R_q, i.e. coefficients of the polynomials coming from the discrete uniform distribution on $\mathbb{Z}_q$. Let ${}_{D}R_q$ denote a discrete Gaussian distribution on R_q, i.e. one that is the preimage of a discrete Gaussian distribution over a lattice in the canonical embedding of a number field K when considering R_q is considered as isomorphic to an ideal of K. Note this distribution ${}_{D}R_q$ is the one considered in [ELOS15]. Let ${}_{\chi}R_q$ denote the discrete Gaussian distribution on R_q, i.e. coefficients of the polynomials coming from a discrete Gaussian distribution on $\mathbb{Z}_q$. The precise formulation of a discrete Gaussian on $\mathbb{Z}_q$ is as follows where we try to keep the definitions consistent with [Reg09]. Let ${}_{\bar{\mathcal{U}}}R_q$ denote a bounded uniform distribution on R_q, i.e. the coefficients of the polynomials coming from a bounded uniform distribution on $\mathbb{Z}_q$. Let ${}_{E}R_q$ denote any distribution on R_q where the coefficients are sampled from a given distribution E on $\mathbb{Z}_q$.

Definition 2.1. For $\beta \in \mathbb{R}^+$ the continuous distribution Ψ_β on $[0, 1)$ is obtained by sampling from a normal distribution with mean 0 and standard deviation $\frac{\beta}{\sqrt{2\pi}}$ and reducing the result mod 1. This probability distribution is given as

$$\forall r \in [0, 1),\ \Psi_\beta(r) := \sum_{k=-\infty}^{\infty} \frac{1}{\beta} \cdot \exp\left(-\pi \left(\frac{r-k}{\beta}\right)^2\right).$$

Now using Ψ_β, the discrete Gaussian distribution on $\mathbb{Z}_q$ is defined as follows.

Definition 2.2. The discretization of a Gaussian distribution on $\mathbb{Z}_q$ (we denote as $G_{q,\beta}$) is obtained by sampling from Ψ_β and multiplying by q. This probability distribution is give by

$$G_{q,\beta}(i) = \int_{(i-\frac{1}{2})/q}^{(i+\frac{1}{2})/q} \Psi_\beta(x) dx.$$

Note that one can easily generate random values from $G_{q,\beta}$, and one can numerically approximate the probability distribution of $G_{q,\beta}$ by using approximations for the infinite sum and the integral. We now introduce the two main problems of interest.

Problem 2.1 (Search Poly-LWE Problem)**.** Let $s(x) \in {}_{\mathcal{U}}R_q$ be secret. The *Search Poly-LWE Problem* is that of finding $s(x)$ given a poly(n) number of samples of the form

$$(a_j(x), b_j(x) := a_j(x) \cdot s(x) + e_j(x)) \in R_q \times R_q$$

where $a_j(x) \in {}_{\mathcal{U}}R_q$ and $e_j(x)$ is sampled from an error distribution on R_q.

A related problem is to distinguish samples coming from a Search LWE Problem from uniform samples on $R_q \times R_q$, and our attack can be extended to work against this variant as well.

Problem 2.2 (Decision Poly-LWE Problem)**.** Given poly(n) samples from one of the following two distributions on $R_q \times R_q$ the *Decision Poly-LWE Problem* is to decide which

distribution the samples are coming from:
(1) samples from a Search Poly-LWE Problem, i.e. of the form

$$(a_j(x), b_j(x) := a_j(x) \cdot s(x) + e_j(x))$$

where $a_j(x) \in {}_{\mathcal{U}}R_q$ and $e_j(x)$ is sampled from an error distribution, or
(2) samples that are uniform, i.e. of the form

$$(a_j(x), b_j(x))$$

where $a_j(x), b_j(x) \in {}_{\mathcal{U}}R_q$.

Attack. In [ELOS15] they develop the following attack against the Decision Poly-LWE Problem which we describe here with a couple of adjustments. For a polynomial $f(x)$ that is irreducible over $\mathbb{Z}$ let $p \in \mathbb{Z}$ be a prime such that $f(x)$ has at least one root α mod p. Note that in [ELOS15] the authors assume that $f(x)$ factors completely mod p which is often the case in practice so that fast multiplication can be done in the ring using the Chinese Remainder Theorem, but we will not need this assumption for our attack.

Given access to $L :=\text{poly}(n)$ samples from a Decision Poly-LWE Problem

$$(a_j(x), b_j(x)) \in R_p \times R_p$$

we want to transfer the problem to $\mathbb{Z}_p$. To do this, we will build a well defined ring homomorphism

$$\bar{\phi} : \frac{Z_p[x]}{\langle f(x) \rangle} \to \mathbb{Z}_p.$$

Since we are assuming we have a root α of $f(x)$ mod p, we can consider the ring homomorphism

$$\phi : \mathbb{Z}_p[x] \to Z_p$$
$$y(x) \mapsto y(\alpha).$$

It is clearly well defined, and one can check it is a ring homomorphism. Moreover $\langle f(x) \rangle \subseteq \text{Ker}(\phi)$ which gives that $\bar{\phi}$ is a well defined ring homomorphism defined as

$$\bar{\phi} : \frac{Z_p[x]}{\langle f(x) \rangle} \to \mathbb{Z}_p$$
$$z(x) + \langle f(x) \rangle \mapsto \phi(z(x)) = z(\alpha).$$

Now we take the samples and map them according to $\bar{\phi}$

$$(a_j(x), b_j(x)) \mapsto (a_j(\alpha), b_j(\alpha)).$$

If the samples are coming from the Search Poly-LWE distribution, $s(\alpha)$ will be some element in $\mathbb{Z}_p$. For the attack we will guess $s(\alpha)$. For each $g \in \mathbb{Z}_p$ we assume g is the correct guess for $s(\alpha)$ and compute

$$b_j(\alpha) - a_j(\alpha)g.$$

Since the multiplication and addition is preserved by $\bar{\phi}$ this gives us

$$e_j(\alpha) = b_j(\alpha) - a_j(\alpha)g.$$

We can now analyze the distribution of $e_j(\alpha)$ to decide which distribution the samples came from. The error polynomial $e_j(x)$ can be written as

$$e_j(x) = \sum_{i=0}^{n-1} e_{ij} x^i,$$

which, when evaluated at α, is

$$e_j(\alpha) = \sum_{i=0}^{n-1} e_{ij} \alpha^i,$$

From now on we will be considering $e_j(\alpha)$ for a particular guess g and will drop the j subscript and just write

$$e(\alpha) = \sum_{i=0}^{n-1} e_i \alpha^i,$$

This can be simplified by considering the order of α, which we denote as r. For simplicity of notation assume r divides n. This gives

$$\begin{aligned} e(\alpha) =& (e_0 + e_r + \cdots + e_{\frac{n}{r}}) + (e_1 + e_{r+1} + \cdots + e_{\frac{n}{r}+1})\alpha + \cdots \\ &+ (e_{r-1} + e_{2r+(r-1)} + \cdots + e_{\frac{n}{r}+(r-1)})\alpha^{r-1}. \end{aligned}$$

This can be further simplified and written as

$$e(\alpha) = v_0 + v_1\alpha + \cdots + v_{r-1}\alpha^{r-1},$$

where

$$\begin{aligned} v_0 &= e_0 + e_r + \cdots + e_{\frac{n}{r}} \\ v_1 &= e_1 + e_{r+1} + \cdots + e_{\frac{n}{r}+1} \\ &\vdots \\ v_{r-1} &= e_{r-1} + e_{2r+(r-1)} + \cdots + e_{\frac{n}{r}+(r-1)}. \end{aligned}$$

If the errors are from ${}_DR_q$ as in [ELOS15], then depending on the value of α and the order of α, this distribution $e(\alpha)$ may be either a discrete Gaussian or a periodic distribution, and depending on the parameters it may be very close to a uniform one. In [ELOS15, Section 3.2 Case 2] the authors state that when α has small order ≥ 3, $e(\alpha)$ will be Gaussian; however this is not quite right as the distribution may be highly periodic as we will demonstrate. Since the algorithms in [ELOS15] are stronly dependent on $e(\alpha)$ being Gaussian to distinguish we will need to introduce some other methods for algorithms to distinguish in this case. We shall see that Chi-square test will work much better.

3. Error Distribution

3.1. **Computing probability distributions.** In the above attacks, we need to understand the probability distribution of $e(\alpha)$. More generally, we consider

$$e = e_0 + e_1a_1 + \cdots + e_{n-1}a_{n-1} \in \mathbb{Z}_p$$

where $a_i \in \mathbb{Z}_p$ are fixed and e_i's are chosen according to a given probability distribution E on $\mathbb{Z}_p$. In the following, we derive the probability distribution of e. We first show how to compute the exact distribution of e efficiently, and then demonstrate several possible distributions of e on $\mathbb{Z}_p$, including distributions that are neither Gaussian nor uniform.

Lemma 3.1. *Let u and v be independent random variables on $\mathbb{Z}_p$ with probability distributions $(a_0, a_1, \ldots, a_{p-1})$ and $(b_0, b_1, \ldots, b_{p-1})$, respectively. Let*

$$a(x) = \sum_{i\in\mathbb{Z}_p} a_i x^i, \qquad b(x) = \sum_{i\in\mathbb{Z}_p} b_i x^i.$$

Then the probability distribution of $u+v$ can be computed as the coefficients of the polynomial $a(x)b(x) \pmod{x^p - 1}$.

The proof is simple since, for any $k \in \mathbb{Z}_p$, the probability

$$P(u + v = k) = \sum_{i\in\mathbb{Z}_p} P(u = i)P(v = k - i \bmod p) = \sum_{i\in\mathbb{Z}_p} a_i b_{k-i},$$

where the subscript $k - i$ of b is computed modulo p.

Theorem 3.2. *Suppose $e_0, e_1, \ldots, e_{n-1}$ are independent random variables on $\mathbb{Z}_p$ with the same probability distribution $(c_0, c_1, \ldots, c_{p-1})$. Let $c(x) = \sum_{i\in\mathbb{Z}_p} c_i x^i$. Then, for any $a_1, \ldots, a_{n-1} \in \mathbb{Z}_p$, the probability distribution of $e = e_0 + e_1a_1 + \cdots + e_{n-1}a_{n-1} \bmod p$ can be computed as the coefficients of the polynomial*

$$c(x)c(x^{a_1})\cdots c(x^{a_{n-1}}) \pmod{x^p - 1}.$$

The theorem follows from the above lemma, since the random variables $e_0, e_1a_1, \ldots, e_{n-1}a_{n-1}$ are independent and $c(x^{a_i})$ represents the probability distribution of e_ia_i. Also, the product can be computed by a simple loop: $t := c(x)$; and for i from 1 to $n-1$ do $t := t \cdot c(x^{a_i}) \bmod x^p - 1$. This takes at most $O(np^2)$ operations in $\mathbb{Z}_p$.

3.2. **Intuition for the error distribution.** In this subsection we give several examples of the mapped error distribution

$$a_0e_0 + a_1e_1 + \cdots + a_{n-1}e_{n-1} \mod p$$

where e_i, $0 \le i < n$, are independent identically distributed $G_{p,\beta}$ distributions on $\mathbb{Z}_p$ and $a_i \in \mathbb{Z}_p$, $0 \le i < n$, are fixed constants. The number of terms n and the sizes of the a_i's greatly affects the shape of the above distribution. To get an idea of what one should expect, we look at three general cases.

3.2.1. ***Case 1: small coefficients.*** The first case we consider is when all the $a_0, ..., a_{n-1}$ coefficients are 1 and the standard deviation is fixed; in this case we consider how varying n affects the shape of the distribution. As n grows large, the distribution remains Gaussian but approaches uniform.

Example 3.3. Let $\beta = 0.01$, $p = 331$. The distribution $e_1+e_2+\cdots+e_n$ for $n = 1, 20, 40, 100$ with e_j iid $G_{p,\beta}$ is shown in Figure 1.

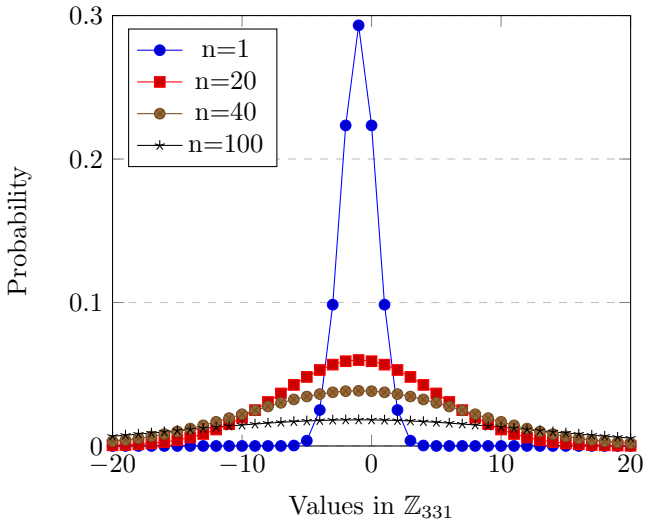

FIGURE 1. Sum of n iid discrete Gaussian distributions on $\mathbb{Z}_{331}$.

3.2.2. ***Case 2: large coefficients.*** The second case we consider is when all the $a_0, ..., a_{n-1}$ coefficients are large and the standard deviation is fixed; in this case we consider how varying n affects the shape of the distribution. We want to answer the question of how large n needs to be for the distribution to be almost uniform. In this case n can be quite small and the distribution already be very close to uniform. However, note that the distributions in the following example are neither Gaussian nor uniform.

Example 3.4. We consider the distributions of $23e_0+45e_1$ and $23e_0+45e_1+43e_2$ and $23e_0+45e_1+43e_2+95e_3$ where $\beta = 0.01$ and $p = 331$ and e_j iid $G_{p,\beta}$. For each additional term in the sum the distributions gets considerably closer to uniform while remaining periodic. The three graphs are shown in Figure 2.

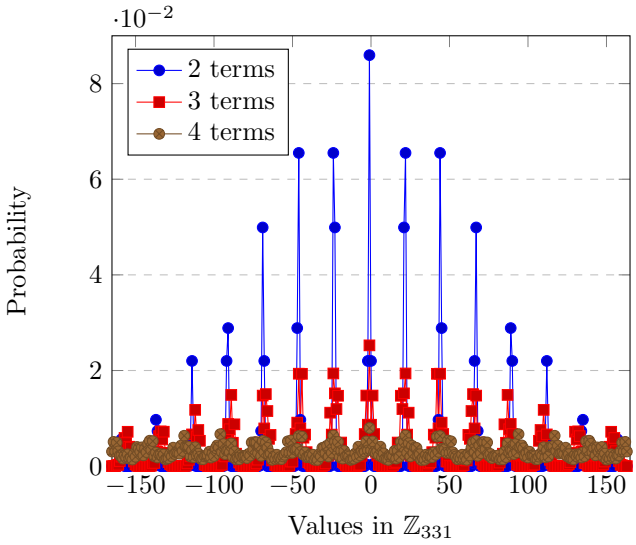

FIGURE 2. The distributions of $23e_0 + 45e_1$ and $23e_0 + 45e_1 + 43e_2$ and $23e_0 + 45e_1 + 43e_2 + 95e_3$ where e_j are iid discrete Gaussians on $\mathbb{Z}_{331}$

3.2.3. ***Case 3: root of small order.*** In the third case we consider a situation that may arise in the attack where the coefficients are all powers of a root α

$$e_0 + \alpha e_1 + \alpha^2 e_2 + \alpha^3 e_3 + \cdots + \alpha^{n-1} e_{n-1}$$

where e_j are iid discrete Gaussian distributions. We want to specifically consider what effect the order of α has on the distribution. We choose $f = x^n + ax + b$ to be an irreducible polynomial over $\mathbb{Z}$ that has a root α mod p.

We want to specifically consider the case when α has small order mod p and n is not too large. When α has low order, the distribution will be considerably farther from uniform compared to when α has large order. Consider the following example where f has one root of low order and another of high order.

Example 3.5. The polynomial $f = x^9 + 11x - 11$ is irreducible over $\mathbb{Z}$ but has a two roots $\alpha_1 = 31$, $\alpha_2 = 82$ mod 331 with α_1 having order 3 and α_2 having order 165. Consider the following distribution which arises if α_1 is used to define the homomorphism in the attack

$$e_0 + \alpha_1 e_1 + \alpha_1^2 e_2 + \alpha_1^3 e_3 + \cdots + \alpha_1^8 e_8.$$

Here we are assuming e_j are iid $G_{p,\beta}$ for $\beta = 0.01$. The graph of the distribution is show in Figure 3. It is neither Gaussian nor uniform.

When using the root α_2 with larger order, we see that the distribution $e_0+\alpha_2 e_1+\alpha_2^2 e_2+\alpha_2^3 e_3+\cdots+\alpha_2^8 e_8$ is much closer to uniform as seen in Figure 4.

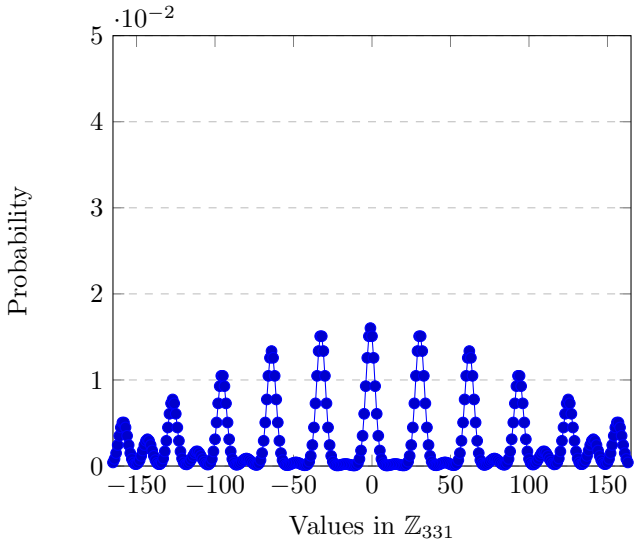

FIGURE 3. The distribution of $e_0 + \alpha_1 e_1 + \alpha_1^2 e_2 + \alpha_1^3 e_3 + \cdots + \alpha_1^8 e_8$ for $\alpha = 31$ of order 3 and e_j iid discrete Gaussians on $\mathbb{Z}_{331}$.

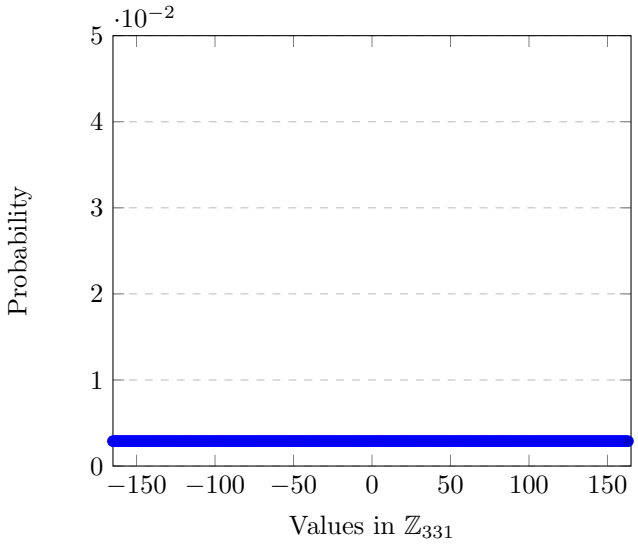

FIGURE 4. The distribution of $e_0 + \alpha_2 e_1 + \alpha_2^2 e_2 + \alpha_2^3 e_3 + \cdots + \alpha_2^8 e_8$ for $\alpha_2 = 84$ of order 165 and e_j iid discrete Gaussians on $\mathbb{Z}_{331}$.

3.2.4. ***Summarize Cases***. Considering these three cases, we note that if the distribution $e(\alpha)$ appears like Case 1 it can be fairly easily distinguished from the uniform for small n. For Case 2 distinguishing from the uniform gets much harder because when adding a large number terms with large coefficients, the distribution rapidly approaches uniform. For Case

3 when α has small order this is similar to Case 2 with a small number of terms, but if the order of α is small enough and n not too large, we can hope to be able to distinguish this from uniform; we will consider this case further in the next section. In Case 3 when α has large order, this is similar to Case 2 when there are a large number of terms, and it is unlikely one would be able to distinguish this from uniform.

4. Statistical Tests and simulation

4.1. **Distinguishing statistical tests.** In this section we will discuss how one can decide if a guess g in the attack is correct. If the guess g is correct, we have shown in the previous section how we can compute what the distribution of the error will be; we denote this computed error distribution as $\mathcal{E}$. If the guess g is not correct then the samples will be uniform, which we denote as $\mathcal{U}$.

Assume that we have L samples of RLWE public keys:

$$(a_i(x), b_i(x)), \quad b_i(x) = a_i(x)s(x) + e_i(x), \quad i = 1, \ldots, L$$

where $a_i(x) \in {}_{\mathcal{U}}R_p, e_i(x) \in {}_{\chi}R_p$. Let A and B be two random variables taking samples

$$A \in \{a_i(\alpha) \mid i = 1, \ldots, L\}$$
$$B \in \{b_i(\alpha) \mid i = 1, \ldots, L\}.$$

In general, (A, B) cannot be distinguished from (A, B'), where B' is uniform. But the homomorphism attack lets us the consider the distribution

$$\mathcal{S}(g) := B - Ag, g \in \mathbb{Z}_p.$$

Property 1. *$\mathcal{S}(g) \sim \mathcal{E}$ if $g = s(\alpha)$; otherwise $\mathcal{S}(g) \sim U$ if $g \neq s(\alpha)$.*

Now if $\mathcal{E}$ is not too close to $\mathcal{U}$, we will be able to decide which one $\mathcal{S}(g)$ matches by considering a reasonable number of samples. There are several methods one might use to decide which distribution $\mathcal{S}(g)$ fits.

4.1.1. ***Method 1: Chi-square Tests.*** First we will use a Chi-square goodness-of-fit test to test the $\mathcal{S}(g)$ against a uniform distribution. For a description of this test see [K$^+$99, pg 37]. Let $\tilde{e}_k$ denote the number of $e_j(\alpha)$'s equal to $k \bmod p$. Let L denote the total number of samples. We set up our null hypothesis to be that $\mathcal{S}(g)$ is distributed according to a uniform distribution

$$H_0 : \mathcal{S}(g) \sim \mathcal{U}$$
$$H_1 : \mathcal{S}(g) \not\sim \mathcal{U}.$$

Then we compute the Chi-square statistic as

$$V = \sum_{k=1}^{p-1} \frac{(\tilde{e}_k - L/p)^2}{L/p}$$

where there are $p - 1$ degrees of freedom. If V is too large or too small based on our choice of Type 1 error rate, we reject H_0. For this test to be considered reliable we need $L/p \geq 5$ samples, i.e. $L \geq 5p$.

Second, we test $S(g)$ against our computed distribution $\mathcal{E}$ using a Chi-squared test in a non-traditional way.

$$H_0 : S(g) \sim \mathcal{E}$$
$$H_1 : S(g) \not\sim \mathcal{E}.$$

Let $\mathcal{E}_k := Pr[\mathcal{E} = k \in \mathbb{Z}_p]$. Then we compute the Chi-square statistic as

$$V = \sum_{k=1}^{p-1} \frac{(\tilde{e}_k - L\mathcal{E}_k)^2}{L\mathcal{E}_k}$$

where there are $p-1$ degrees of freedom. We then reject or accept based on V and our desired Type 1 error rate. In this setup, we are inverting the usual setup for a Chi-square test. Note that the distribution in the null hypothesis has been switched. But this problem is also unusual in that it is a promise problem, giving us that the $S(g)$ must be one of two known distributions, which is not an assumption generally considered for most statistical tests like the Chi-square test.

4.1.2. ***Method 2: Statistical Distance***. For another test making use of the computed error distribution, we consider the statistical distance from $S(g)$ to $\mathcal{U}$ and $\mathcal{E}$ for each guess g and take the guess for which $S(g)$ is closest to $\mathcal{E}$. Let the distribution of $\mathcal{S}(g)$ be $\{t_i(g) \mid i \in \mathbb{F}_p\}$, i.e.,

$$P[\mathcal{S}(g) = i] = t_i(g), \quad i \in \mathbb{F}_p.$$

Let

$$\varepsilon(g) = \sum_i \left| t_i(g) - \frac{1}{p} \right|,$$
$$\delta(g) = \sum_i \left| t_i(g) - \mathcal{E}_k \right|.$$

Decode $s(\alpha) = g$, where

$$g = \arg\max_{g \in \mathbb{F}_p} \varepsilon(g) \tag{1}$$

or

$$g = \arg\min_{g \in \mathbb{F}_p} \delta(g). \tag{2}$$

This gives us a good idea of what the correct guess probably is if the samples are from Poly-LWE. If no one guess differed from the rest by more than some threshold, we would decide the samples were uniform. However, how this threshold should be chosen is unclear at this moment. We will leave that as a future work.

4.1.3. ***Type 1 Errors***. One further thing that must be considered when using any tests like Methods 1 with a fixed Type 1 error probability is that using the test repeated for each of the p guesses will result in a much higher overall Type 1 error probability. To see this in detail, if γ is set to be the Type 1 error probability for a single test, then $1-\gamma$ is the probability of not having a Type 1 error on that test. If one runs p such tests, the probability of no Type 1 errors in all p tests is $(1-\gamma)^p$. Thus the probability of at least one Type 1 error is $1-(1-\gamma)^p$. To be clear, this is an upper bound on probability of a Type 1 error and in practice a test may have a true Type 1 error rate much lower, but we may no longer have a very good upper bound. To see how much the Type 1 error bound can grow, consider that if $\gamma = 0.02$ and $p = 331$, then $1-(1-0.02)^{331} \approx 0.9987$.

It is possible to keep the overall Type 1 error bound to an desired rate. One way is to use a Bonferroni correction, which is a way of setting the Type 1 error rates on the individual tests to guarantee a particular overall Type 1 error. In particular if we set the new Type 1 error for each test at $\beta_\gamma := \frac{\gamma}{p}$, the overall Type 1 error will still be bounded by γ. However, this may result in an impractically high Type 2 error rate, so in practice we would recommend using the tests at multiple levels and using a variety of tests as described above.

4.2. **Simulations.** We simulate these Chi-square methods and look at the there true Type 1 and Type 2 error rates, the expected value of cosine test does not perform well on either of these two examples.

Example 4.1. Continuing with the earlier Example 3.5, we show how we can use Method 1 of the Chi-square tests to determine for which guess of g the distribution $e(\alpha)$ follows the computed error distribution rather than a uniform distribution. With 2000 samples and the individual Chi-Square tests' Type 1 error rates set at 2%, the overall test is successful at rejecting the null hypothesis that the error are from the uniform each time out of ten independent simulations while only every giving two false rejections.

For those same ten simulations the Chi-square test against the uniform with Type 1 error rates set at 1×10^{-9} (corresponding to this tests Type 2 error rates because of the inverted setup) it rejects all guesses correctly without giving any false rejections.

Example 4.2. Next we consider a new example. Let $f = x^{15} + 125x - 334$ which is irreducible over $\mathbb{Z}$ but has a root $\alpha = 396$ of order 3 mod 607. The distribution $\mathcal{E}$ is shown in Figure 5. Clearly it is neither Gaussian nor uniform.

Using the Chi-square test against uniform with 5000 samples and having the individual tests set at a Type 1 error rate of 2%, we are able to reject the correct the guess correctly for every one of ten independent simulations while only every giving one false rejection.

For those same ten simulations the Chi-square test against the uniform with Type 1 error rate set at 1×10^{-9} (corresponding to this test's Type 2 error rate because of the inverted setup) is not helpful on this example as it never rejects anything.

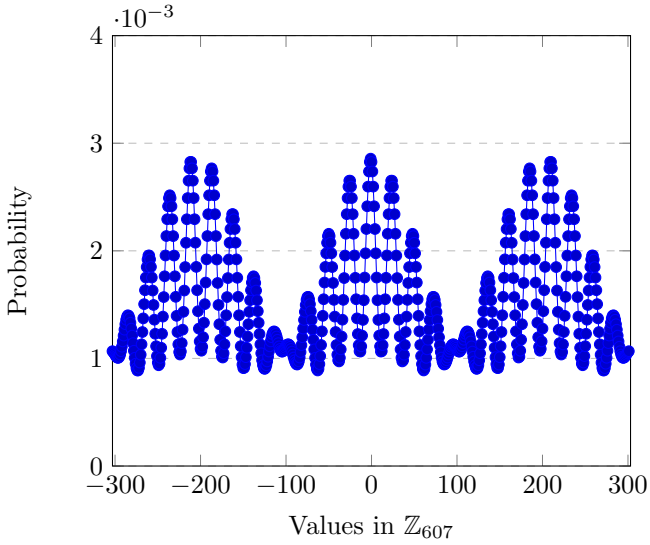

FIGURE 5. The distribution of $e_0 + \alpha e_1 + \alpha^2 e_2 + \alpha^3 e_3 + \cdots + \alpha^{12} e_{12}$ for $\alpha = 396$ of order 3 and e_j iid discrete Gaussians on $\mathbb{Z}_{607}$.

Example 4.3. We used Method 2 (statistical distance) on the following instance: $f(x) = x^n + p - 1$, where $n = 8$, $p = 257$, and $\beta = 0.2$. In this setup, $f(x)$ has root $\alpha = 1 \bmod p$. With 1200 samples, the test is successful at finding the image of s in all of the ten independent simulations, when either of (1) and (2) is used.

This instance has been attacked successfully by [EHL14], because of its simplicity. However, our method may be able to improve the efficiency of their attack by computing $s(\alpha)$ directly from the error distribution under the homomorphism. We leave that as a future work.

5. Concluding Remarks

In this work, we provide a general algebraic method to derive the mapped error distribution with a formula. Analyzing Poly-LWE instances with our method will allow such attacks to be applied to a broader range of parameters; in particular, those with non-Gaussian error distribution can be included. As an extension to this work, we may consider adapting our method to Ring-LWE or further investigating the characteristics of different statistic tests in order to more accurately recover the secret.

References

[BLLN13] Joppe W Bos, Kristin E Lauter, Jake Loftus, and Michael Naehrig. Improved security for a ring-based fully homomorphic encryption scheme. In *IMA Int. Conf.*, pages 45–64. Springer, 2013.

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 575–584. ACM, 2013.

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *Advances in cryptology—CRYPTO 2011*, volume 6841 of *Lecture Notes in Comput. Sci.*, pages 505–524. Springer, Heidelberg, 2011.

[CIV16] Wouter Castryck, Ilia Iliashenko, and Frederik Vercauteren. Provably weak instances of ring-LWE revisited. In *Advances in cryptology—EUROCRYPT 2016. Part I*, volume 9665 of *Lecture Notes in Comput. Sci.*, pages 147–167. Springer, Berlin, 2016.

[CLS15] Hao Chen, Kristin Lauter, and Katherine E Stange. Attacks on search rlwe. 2015.

[CLS16] Hao Chen, Kristin E Lauter, and Katherine E Stange. Vulnerable galois rlwe families and improved attacks. *IACR Cryptology ePrint Archive*, 2016:193, 2016.

[EHL14] Kirsten Eisenträger, Sean Hallgren, and Kristin Lauter. Weak instances of PLWE. In *Selected areas in cryptography—SAC 2014*, volume 8781 of *Lecture Notes in Comput. Sci.*, pages 183–194. Springer, Cham, 2014.

[ELOS15] Yara Elias, Kristin E. Lauter, Ekin Ozman, and Katherine E. Stange. Provably weak instances of ring-LWE. In *Advances in cryptology—CRYPTO 2015. Part I*, volume 9215 of *Lecture Notes in Comput. Sci.*, pages 63–92. Springer, Heidelberg, 2015.

[FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.

[K+99] Donald E Knuth et al. The art of computer programming. *Sorting and searching*, 3:426–458, 1999.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *Advances in cryptology—EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Comput. Sci.*, pages 1–23. Springer, Berlin, 2010.

[Pei16] Chris Peikert. How (not) to instantiate ring-LWE. In *Security and cryptography for networks*, volume 9841 of *Lecture Notes in Comput. Sci.*, pages 411–430. Springer, [Cham], 2016.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 84–93. ACM, New York, 2005.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):Art. 34, 40, 2009.