Vulnerable Galois RLWE Families and Improved Attacks

Hao Chen, Kristin E. Lauter, Katherine E. Stange

1. *Introduction*

Lattice-based cryptography was introduced in the mid 1990s in two different forms, independently by Ajtai-Dwork [**AD97**] and Hoffstein-Pipher-Silverman [**HPSS08**]. Thanks to the work of Stehlé-Steinfeld [**SS11**], we now understand the NTRU cryptosystem introduced by Hoffstein-Pipher-Silverman to be a variant of a cryptosystem which has security reductions to the Ring Learning With Errors (RLWE) problem. The RLWE problem was introduced in [**LPR13**] as a version of the LWE problem [**Reg09**]: both problems have reductions to hard lattice problems and thus are interesting for practical applications in cryptography. RLWE has more structure which allows for greater efficiency, but also in some cases additional attacks.

The hardness of RLWE is an important problem to study due to applications in cryptography, in particular as the basis of numerous homomorphic encryption schemes [**BV11**, **BV14**, **BGV12**, **Bra12**, **SS11**, **LATV12**, **BLLN13**]. Although so far in practical cryptographic applications only cyclotomic rings are used, it is interesting to study the hardness of RLWE for general number rings. Recently, new attacks on the RLWE problem for certain number rings and special moduli were introduced [**EHL14**, **ELOS15**, **CLS15**, **CIV**].

This paper is an extension of [**CLS15**], and here we explore further the hardness of the RLWE problem for various number rings, construct a new family of vulnerable Galois number fields, give improved attacks for certain rings satisfying some additional assumptions, and apply some number theoretic results on Gauss sums to deduce the likely failure of these attacks for cyclotomic rings and unramified moduli.

To be more specific, the RLWE problem is stated given a choice of number ring of degree n, R, modulus q, and error distribution. In cryptographic applications, it is most efficient to sample the error distribution coordinate-wise according to a polynomial basis for the ring. For 2-power cyclotomic rings which are monogenic with a well-behaved power basis, it is justified to sample the RLWE error distribution directly in the polynomial basis for the ring, according to results in [**BV11**, **LPR13**, **EHL14**], where the Polynomial Learning With Errors (PLWE) problem was introduced. Although the PLWE and RLWE problems are equivalent for 2-power cyclotomic fields, in general number rings the two problems are not at all equivalent, as was shown in [**ELOS15**]. For certain choices of ring, R, and modulus q, efficient attacks on PLWE were presented in [**EHL14**]. In [**ELOS15**], these attacks were extended to apply to the decision version of the RLWE problem in certain rings, and in [**CLS15**, **CIV**], attacks on the search version of the RLWE problem for certain choices of ring and modulus were presented. So it is important to study the hardness of the both PLWE and RLWE problems and the relationship between the two problems in general rings.

1.1. *Preliminaries*

We give a short introduction to some of the basics of RLWE. A more detailed introduction can be found in Section 2 of [**CLS15**].

2000 *Mathematics Subject Classification* 11T71.

First, to define an RLWE instance, we specify a number field K of degree n with ring of integers R, a prime q called the *modulus*, a positive real number r, and an element $s \in R/qR$ called the *secret*.

Let $\iota : K \to \mathbb{R}^n$ be the "adjusted canonical embedding" defined in [**CLS15**, Section 2]. We recall the definition of the embedding ι: suppose K is a number field of degree n and signature (r_1, r_2) with embeddings $\sigma_1, \ldots, \sigma_{r1}, \sigma_{r_1+1}, \ldots, \sigma_n$, such that $\sigma_1, \cdots, \sigma_{r_1}$ are the real embeddings and $\sigma_{r_1+r_2+j} = \overline{\sigma_{r_1+j}}$ for $1 \leq j \leq r_2$. Then we define

$$\iota : K \to \mathbb{R}^n : \quad x \mapsto (\sigma_1(x), \cdots, \sigma_{r_1}(x), \mathrm{Re}(\sigma_{r_1+1})(x), \mathrm{Im}(\sigma_{r_1+1})(x), \cdots, \mathrm{Re}(\sigma_{r_1+r_2})(x), \mathrm{Im}(\sigma_{r_1+r_2})(x)).$$

Then the (non-dual and discrete) *RLWE error distribution* is the discrete Gaussian distribution on the lattice $\iota(R)$ with width r (see [**LPR13**, Section 2], for example, for the definition of discrete Gaussian distributions over lattices).

Let R_q denote the quotient ring R/qR; then an RLWE sample is a pair

$$(a, b = as + e \bmod qR) \in R_q \times R_q,$$

where the first coordinate a is chosen uniformly at random in R_q, and e is a sampled from the error distribution.

The *search* RLWE problem is to discover the secret s given access to arbitrarily many independent RLWE samples (a, b). The *decision* RLWE problem is to distinguish RLWE samples from samples of the uniform distribution on $R_q \times R_q$. Let $\mathfrak{q}$ be a prime ideal of K lying above q; then the RLWE problem modulo $\mathfrak{q}$ means discovering $s \mod \mathfrak{q}$ from arbitrarily many RLWE samples. In [**LPR13**] the authors showed a polynomial time reduction from RLWE modulo $\mathfrak{q}$ to *decision* RLWE.

For an element $v \in K$, we let $||v||$ denote its 2-norm under the embedding, i.e., $||v|| := ||\iota(v)||_2$. We will call this the *embedding length* of v.

As pointed out in [**ELOS15**], when analyzing the error distribution, one needs to take into account the sparsity of the lattice $\iota(R)$, measured by its covolume in $\mathbb{R}^n$. We know this covolume is equal to $|\operatorname{disc}(K)|^{1/2}$. In light of this, we define the scaled error width to be

$$r_0 = \frac{r}{|\operatorname{disc}(K)|^{\frac{1}{2n}}}.$$

1.2. *Summary of contributions*

- In Section 2, we present a new infinite family of Galois number fields vulnerable to our attack in [**CLS15**, Section 4], where the relative standard deviation parameter is allowed to grow to infinity, and we give a table of examples.
- In Section 3, we present an improvement to the attack in [**CLS15**, Section 4] and use it to dramatically cut down the runtime of the attacks on the weak instances found in [**CLS15**, Section 5].
- In Section 4, we analyze the security of cyclotomic fields with unramified moduli under our attack. We give some heuristics based on a modified discrete RLWE error distribution. Then we support the heuristics with Theorem 3, which gives an upper bound on the statistical distance between a modified reduced error distribution and the uniform distribution on $R/\mathfrak{q}$. We conclude that cyclotomic fields are very likely safe against our attack when the modulus q is unramified with small residue degree (1 or 2).

Acknowledgements. We thank Chris Peikert, Igor Shparlinski, Leo Ducas and Ronald Cramer for helpful discussions.

2. *Infinite family of vulnerable Galois RLWE instances*

In this section, describe Galois number fields which are vulnerable to the attack in Section 4 of [**CLS15**]. In contrast to the vulnerable instances found by computer search in Section 5 of [**CLS15**], in this section we explicitly construct infinite families of such fields with flexible parameters. Furthermore, the attacks of [**CLS15**] were successful only on instances where the size of the distribution (in the form of the scaled standard deviation) is a small constant, where as in this paper the scaled standard deviation parameter can be taken to be $o(|d|^{1/4})$, where d is an integer parameter and can go to infinity.

We briefly review the method of attack in Section 4 of [**CLS15**]. The basic principle of this family of attacks is to find a homomorphism

$$\rho : R_q \to F$$

to some small finite field F, such that the error distribution on R_q is transported by ρ to a non-uniform distribution on F. In this case, errors can be distinguished from elements uniformly drawn from R_q by a statistical test in F, for example, by a χ^2-test. The existence (or non-existence) of such a homomorphism depends on the parameters of the field, prime, and distribution in the setup of RLWE. In this section, we will describe parameters under which such a map exists.

Once such a map is known, the basic method of attack on Decision RLWE is as follows:

(1) Apply ρ to samples (a, b) in $R_q \times R_q$, to obtain samples in $F \times F$.
(2) Guess the image of the secret $\rho(s)$ in F, calling the guess g.
(3) Compute the distribution of $\rho(b) - \rho(a)g$ for all the samples. If $g = \rho(s)$, this is the image of the distribution of the errors. Otherwise it is the image of a uniform distribution.
(4) If the image looks uniform, try another guess g until all are exhausted. If any non-uniform distribution is found, the samples are RLWE samples. Otherwise they are not.

This is actually an attack on RLWE modulo $\mathfrak{q}$, for some prime $\mathfrak{q}$ lying above q. See [**CLS15**] for how this can be used to attack Search RLWE when K is Galois.

Notice that the attack requires looping through all guesses g in F. In the next section, we will improve this attack to avoid such a large loop.

To set up, let p be an odd prime and let $d > 1$ be a squarefree integer such that d is coprime to p and $d \equiv 2, 3 \mod 4$. We choose an odd prime q such that

(1) $q \equiv 1 \pmod{p}$.
(2) $\left(\frac{d}{q}\right) = -1$ (equivalently, the prime q is inert in $\mathbb{Q}(\sqrt{d})$).

Remark 1. Fix a pair (p, d) that satisfies the conditions described above. By quadratic reciprocity, condition (2) on q above is a congruence condition modulo $4d$. So by Dirichlet's theorem on primes in arithmetic progressions, there exists infinitely many primes q satisfying both (1) and (2).

Let $M = \mathbb{Q}(\zeta_p)$ be the p-th cyclotomic field and $L = \mathbb{Q}(\sqrt{d})$. Let $K = M \cdot L$ be the composite field and let $\mathcal{O}_K$ denote its ring of integers. Our goal in this section is to prove:

Theorem 1. *Let K and q be as above, and R_q defined as in the preliminaries in terms of K and q. Suppose $\mathfrak{q}$ is a prime ideal in K lying over q. We consider the reduction map $\rho : R/qR \to R/\mathfrak{q}R \cong \mathbb{F}_{q^f}$, where f is the residue degree. Suppose $\mathcal{D}$ is the RLWE error*

distribution with error width r such that $r < 2\sqrt{\pi d}$. Let

$$\beta = \min\left\{ \left(\frac{\sqrt{4\pi e d}}{r} e^{-\frac{2\pi d}{r^2}} \right)^n, 1 \right\}.$$

Then, for $x \in R_q$ drawn according to $\mathcal{D}$, we have $\rho(x) \in \mathbb{F}_q$ with probability at least $1 - \beta$.

EXAMPLE 1. As a sample application of the theorem, we take $d = 4871, r = 68.17$ and $p = 43$. Then we computed $\beta = 0.11\ldots$. So if $x \in R_q$ is drawn from the error distribution, then $\rho(x) \in \mathbb{F}_q$ with probability at least 0.88.

LEMMA 1. *Under the notation above, we have*
(1) $K/\mathbb{Q}$ is a Galois extension.
(2) $[K : \mathbb{Q}] = [M : \mathbb{Q}][L : \mathbb{Q}] = 2(p-1)$.
(3) The prime q has residue degree 2 in K.
(4) $\mathcal{O}_K = \mathcal{O}_M \cdot \mathcal{O}_L = \mathbb{Z}[\zeta_p, \sqrt{d}]$.
(5) $|\operatorname{disc}(\mathcal{O}_K)| = p^{2(p-2)}(4d)^{(p-1)}$.

Proof. (1) follows from the fact that K is a composition of Galois extensions M and L; (2) is equivalent to $M \cap L = \mathbb{Q}$, which holds because $L/\mathbb{Q}$ is unramified away from primes dividing $2d$ and $M/\mathbb{Q}$ is unramified away from p; for (3), note that our assumptions imply that q splits completely in M and is inert in L, hence the claim. The claims (4) and (5) follow directly from [**Mar77**, II. Theorem 12], and the fact that $\operatorname{disc}(\mathcal{O}_M) = p^{p-2}$ and $\operatorname{disc}(\mathcal{O}_L) = 4d$ are coprime. □

The following lemma is a standard bound on Euclidean length of samples from discrete Gaussians over lattices. It can be deduced from [**MR07**, Lemma 2.10], for example.

LEMMA 2. *Suppose $\Lambda \subseteq \mathbb{R}^n$ is a lattice. Let $D_{\Lambda,r}$ denote the discrete Gaussian over Λ of width r. Suppose c is a positive constant such that $c > \frac{r}{\sqrt{2\pi}}$. Let v be a sample from $D_{\Lambda,r}$. Then*

$$\operatorname{Prob}(||v||_2 > c\sqrt{n}) \leq C_{c/r}^n,$$

where $C_s = s\sqrt{2\pi e} \cdot e^{-\pi s^2}$.

Proof of Theorem. Part (3) of Lemma 1 implies that

$$1, \zeta_p, \ldots, \zeta_p^{p-2}; \sqrt{d}, \ldots, \zeta_p^{p-2}\sqrt{d} \qquad (*)$$

is an integral basis of $R = \mathcal{O}_K$. By our assumptions, we have $R/\mathfrak{q}R \cong \mathbb{F}_{q^2}$, the finite field of q^2 elements. Under the map ρ, the first $(p-1)$ elements of the basis reduce to $\mathbb{F}_q$, and the rest reduce to the complement $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$, because d is not a square modulo q.

Let $n = p - 1$ be the degree of M over $\mathbb{Q}$. Then the extension $K/\mathbb{Q}$ has degree $2n$. Moreover, the first n elements in the basis (*) have embedding length $\sqrt{2n}$, while the last n have embedding length $\sqrt{2nd}$. We denote the elements in (*) by $v_1, \ldots, v_n$ and $w_1, \ldots, w_n$.

We compute the root volume $c := (vol(R))^{1/n}$. It is a general fact that $vol(R) = |\operatorname{disc}(R)|^{\frac{1}{2}}$, so we have

$$c = |\operatorname{disc}(R)|^{\frac{1}{2n}} = \sqrt{2} p^{\frac{p-2}{2(p-1)}} d^{\frac{1}{4}}.$$

So when $d \gg p$, we have $|v_i| \ll c \ll |w_i|$. We have a decomposition $R = V \oplus W$, where V and W are the vector spaces with bases $v_1, \ldots, v_n$ and $w_1, \ldots, w_n$ respectively. The embeddings

of V and W are orthogonal subspaces, because $\mathrm{Tr}(v_i\bar{w_j}) = 0$ for all i, j. For any element $e \in R$, we can write $e = e_1 + e_2\sqrt{d}$ where e_1, e_2 are elements of $\mathbb{Z}[\zeta_p]$, and it follows that $||e||^2 = ||e_1||^2 + d||e_2||^2$. In particular, if $e_2 \neq 0$, then $||e|| \geq \sqrt{2nd}$.

By applying Lemma 2 with $c = \sqrt{2d}$, the assumptions in the statement of our theorem imply that the probability that the discrete Gaussian $D_{\iota(R),r}$ will output a sample with $e_2 \neq 0$ is $\leq 1 - \beta$. So the statement of theorem follows, since $e_2 = 0$ implies $\rho(e) \in \mathbb{F}_q$, i.e., the image of e lies in the prime subfield. □

Therefore, we can specialize the general attack in this situation as follows. Given a set S of samples $(a, b) \in (R/qR)^2$, we loop through all q^2 possible guesses g of the value $s \mod \mathfrak{q}$ and compute $e_g = \rho(b) - g\rho(a)$. We then perform a chi-square test on the set $\{e_g : (a, b) \in S\}$, using two bins $\mathbb{F}_q$ and $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$. If the samples are not taken from the RLWE distribution, or if the guess is incorrect, we expect to obtain uniform distributions; for the correct guess, we have $e_g = \rho(e)$, and by the above analysis, if the error parameter r_0 is sufficiently small, then the chi-square test might detect non-uniformness, since the portion of elements that lie in $\mathbb{F}_q$ might be larger than $1/q$.

The theoretical time complexity of our attack is $O(nq^3)$: the loop runs through q^2 possible guesses. In each passing of the loop, the number of samples we need for the chi-square test is $O(q)$, and the complexity of computing the map ρ on one sample is $O(n)$. Note that using the techniques in Section 3 of this paper, we could reduce the complexity to $O(nq^2)$.

Remark 2. It is easy to verify that if a triple (p, q, d) satisfies our assumptions, then so does $(p, q, d + 4kq)$ for any integer k, as long as $d + 4kq$ is square free. This shows one infinite family of Galois fields vulnerable to our attack.

2.1. *Examples*

Table 1 records some of the successful attacks we performed on the instances described previously. In each row of Table 1, the degree of the number field is $2(p - 1)$. Note that the runtimes are computed based on the improved version of the attack described in Section 3 of this paper. Also, by varying the parameters p and d, we can find vulnerable instances with $r_0 \to \infty$. For example, any $r_0 = o(d^{1/4}/\sqrt{p})$ will suffice.

Remark 3. From Table 1, we see that the the attack in practice seems to work better (i.e., we can attack larger width r) than what is predicted in Theorem 1. As a possible explanation, we remark that in proving the theorem we bounded the probability of $e_2 = 0$ from below. However, the condition $e_2 = 0$ is sufficient but not necessary for $\rho(e)$ to lie in $\mathbb{F}_q$, so our estimation may be a very loose one.

Table 1. *New vulnerable Galois RLWE instances*

p	d	q	r_0	r	no. samples	runtime (in seconds)
31	4967	311	8.94	592.94	3110	144.92
43	4871	173	8.97	694.94	1730	6.44
61	4643	367	8.84	815.11	3670	205.28
83	4903	167	8.94	963.84	1670	5.74
103	4951	619	8.94	1076.32	6190	579.77
109	4919	1091	8.94	1105.44	10910	1818.82
151	100447	907	14.08	4356.02	9070	1394.18
181	100267	1087	14.11	4777.17	10870	1973.47

2.2. *Remarks on other possible attacks*

First, we note that the instances we found in this section are not directly attackable using linear algebra, as in the recent paper [**CIV**]. The reason is that although the last $n/2$-coordinates of the error e under the basis (*) are small integers, they are nonzero most of the time, so it is not clear how one can extract exact linear equations from the samples. On the other hand, note that for linear equations with small errors, there is the attack on the search RLWE problem proposed by Arora and Ge. However, the attack requires $O(n^{d-1})$ samples and solving a linear system in $O(n^d)$ variables. Here d is the width of the discrete error: (e.g. if the error can take values $0, 1, 2, -1, -2$, then $d = 5$). Thus the attack of Arora and Ge becomes impractical when n is larger than 10^2 and $d \geq 5$, say. In contrast, the complexity of our attack depends linearly on n and quadratically on q. In particular, it does not depend on the error size (although the success rate does depend on the error size).

3. *An improved attack using cosets*

In this section, we describe an improvement to our chi-square attack on RLWE mod $\mathfrak{q}$ in [**CLS15**] for a special case. As a result, we have an updated version of [**CLS15**, Table 1], where we attacked each instance in the table in much shorter time. Note that the complexity of the previous attack in this special case is $O(nq^3)$. In contrast, our new attack has complexity $O(nq^2)$. Hence we have saved a factor of q.

To clarify, the special case we consider in this section is characterised by the following assumptions (we need not be in the special family of the previous section):

- The modulus q is a prime of residue degree 2 in the number field K.
- There exists a prime ideal $\mathfrak{q}$ above q such that the map $\rho : R_q \to R_{\mathfrak{q}}$ satisfies the following property: Let $e \in R_q$ be taken from the discrete RLWE error distribution. The probability that $\rho(e)$ lies in the prime subfield $\mathbb{F}_q$ of $\mathbb{F}_{q^2}$ is computationally distinguishable from $1/q$.

Granting these assumptions, we can distinguish the distribution of the "reduced error" $\rho(e)$ from the uniform distribution on $\mathbb{F}_{q^2}$. More precisely, the attack in [**CLS15**] works exactly as we described in Section 2: with access to $\Omega(q)$ samples, one loops over all q^2 possible values of $\rho(s)$. It marks the correct guess $\rho(s)$ based on chi-square test with two bins $\mathbb{F}_q$ and $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$.

The distinguishing feature of the improved attack is to loop over the cosets of $\mathbb{F}_q$ of $\mathbb{F}_{q^2}$ instead of the whole space. Fix $t_1, \cdots t_q$ to be a set of coset representatives for the additive group $\mathbb{F}_{q^2}/\mathbb{F}_q$. Recall that s denotes the secret and $\rho : R_q \to R_{\mathfrak{q}} \cong \mathbb{F}_{q^2}$ is a reduction map modulo some fixed prime ideal $\mathfrak{q}$ lying above q. Then there exists a unique index i such that $\rho(s) = s_0 + t_i$ for some $s_0 \in \mathbb{F}_q$. Our improved attack will recover s_0 and t_i separately.

We start with an identity $b = as + e$, where $a, b, s, e \in \mathbb{F}_{q^2}$. We will regard s as fixed and a, b, e as random variables, such that a is uniformly distributed in $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$ and b is uniformly distributed in $\mathbb{F}_{q^2}$. The reason why a is not taken to be uniform will become clear later in this section. We use a bar to denote the Frobenius automorphism, i.e.,

$$\bar{a} \stackrel{def}{=} a^q, \ \forall a \in \mathbb{F}_{q^2}.$$

Then $\bar{b} = \bar{a}\bar{s} + \bar{e}$. Using the identity $s = s_0 + t_i$ and subtracting, we obtain $\bar{b} - b - \overline{at_i} + at_i = s_0(\bar{a} - a) + \bar{e} - e$. Since $a \neq \bar{a}$, we can divide through by $\bar{a} - a$ and get

$$\frac{\bar{b} - b - \overline{at_i} + at_i}{\bar{a} - a} = s_0 + \frac{\bar{e} - e}{\bar{a} - a}. \qquad (**)$$

Now for each $1 \leq j \leq q$, we can compute

$$m_j(a, b) := \frac{\bar{b} - b - \overline{at_j} + at_j}{\bar{a} - a}$$

with access to a and b, but without knowledge of s or s_0. Note that m_j is in the prime field $\mathbb{F}_q$ by construction.

PROPOSITION 1. *For each $1 \le j \le q$,*
(1) If $j \neq i$, then $m_j(a,b)$ is uniformly distributed in $\mathbb{F}_q$, for RLWE samples (a,b).
(2) If $j = i$, then $m_j(a,b) = s_0 + \frac{\bar{e}-e}{\bar{a}-a}$.

We postpone the proof of Proposition 1 until the end of this section. Assuming the proposition, our improved attack works as follows: for $1 \le j \le q$, we compute a set of m_j from the samples. To avoid dividing by zero, we ignore the samples with $\rho(a) \in \mathbb{F}_q$ (which happens with probability $1/q$ since $\rho(a)$ is uniformly distributed). We then run a chi-square test on the m_j values. If $j \neq i$, then the distribution should be uniform; if $j = i$, then $P(m_i = s_0) = P(e \in \mathbb{F}_q)$, which by our assumption is larger than $1/q$. Hence if we plot the histogram of the m_i computed from the samples, we will see a spike at s_0. So we could recover s_0 as the element with the highest frequency, and output $\rho(s) = s_0 + t_i$. We give the pseudocode of the attack below.

Algorithm 1 Improved chi-square attack on RLWE modulo $\mathfrak{q}$)

Input: K – a number field R – the ring of integers of K; $\mathfrak{q}$ – a prime ideal in K above q with residue degree 2; $\mathcal{S}$ – a collection of M RLWE samples; $\beta > 0$ – the parameter used for comparing χ^2 values.

Output: a guess of the value $s \pmod{\mathfrak{q}}$, or **NOT-RLWE**, or **INSUFFICIENT-SAMPLES**

```
Let G ← ∅.
for j in 1, ..., q do
    E_j ← ∅.
    for a, b in S do
        ā, b̄ ← a (mod q), b (mod q).
        m_j ← (b̄ − b − ā t_j + a t_j) / (ā − a).
        add m_j to E_j.
    end for
    Run a chi-square test for uniform distribution on E_j.
    if χ²(E_j) > β then
        s_0 := the element(s) in E_j with highest frequency.
        s ← s_0 + t_j, add s to E_j.
    end if
end for
if G = ∅ then
    return NOT-RLWE
else if G = {s} is a singleton then
    return s
else
    return INSUFFICIENT-SAMPLES
end if
```

We analyze the complexity of our improved attack. There are q iterations, each operating on $O(q)$ samples, and reduction of each sample is $O(n)$. So our new attack has complexity $O(nq^2)$.

3.1. *Examples of successful attacks*

To illustrate the idea, we apply our improved attack to the instances in Table 1 of [**CLS15**]. Comparing the last column with the current Table 2, we see that the runtime has been improved significantly.

TABLE 2. *RLWE instances under our improved attack*

n	q	f	r_0	no. samples	old runtime (in hours)	new runtime (in seconds)
40	67	2	2.51	22445	3.49	210.42115
60	197	2	2.76	3940	1.05	142.68
60	617	2	2.76	12340	228.41 (est.)	1280.44
80	67	2	2.51	3350	4.81	31.77
90	2003	2	3.13	60090	1114.11 (est.)	18349.36
96	521	2	2.76	15630	75.41 (est.)	1301.93
100	683	2	2.76	20490	276.01 (est.)	2191.26
144	953	2	2.51	38120	5.72	6871.66

3.2. *Proof of Proposition 1*

For notational convenience, we let A_q denote the set $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$.

LEMMA 3. *Let the random variable a be uniformly distributed in A_q. Suppose e is a random variable with value in $\mathbb{F}_{q^2}$ independent of a. Fix $\delta \in A_q$ and $s_0 \in \mathbb{F}_q$. Then*

$$m_\delta = g_\delta + s_0 + \frac{\bar{e} - e}{\bar{a} - a}$$

is uniformly distributed in $\mathbb{F}_q$. Here

$$g_\delta = \frac{\overline{a\delta} - a\delta}{\bar{a} - a}.$$

Proof. Since the uniform distribution is invariant under translation, we may assume $s_0 = 0$. We introduce a new set $V = \{x \in \mathbb{F}_{q^2} : \bar{x} = -x\}$. We claim that for any $c, d \in V$ with $c \neq 0$, we have $P(\bar{a} - a = c, \overline{a\delta} - a\delta = d) = \frac{1}{q(q-1)}$. To prove the claim, note that V is a $\mathbb{F}_q$-vector space of dimension one, and we have the following $\mathbb{F}_q$-linear map $f_\delta : \mathbb{F}_{q^2} \to V^2$.

$$f_\delta : a \mapsto (\bar{a} - a, \overline{a\delta} - a\delta).$$

First we show f_δ is injective: if $f_\delta(a) = 0$, then $a \in \mathbb{F}_q$ and thus $a(\bar{\delta} - \delta) = 0$, so $a = 0$. By dimension counting, f_δ is an isomorphism. Restricting to A_q, we see that $f_\delta|_{A_q}$ gives an isomorphism between A_q and $(V \setminus \{0\}) \times V$. This proves the claim.

Let $e' = \frac{\bar{e}-e}{\bar{a}-a}$. For any $z \in \mathbb{F}_q$, we have

$$\begin{aligned}
&P(g_\delta + e' = z)\\
&= \sum_{x+y=z} P(g_\delta = x, e' = y)\\
&= \sum_{x+y=z} \sum_{c \in V\setminus\{0\}} P(\bar{a\delta} - a\delta = xc, \bar{e} - e = yc, \bar{a} - a = c)\\
&= \sum_{x+y=z, c\in V\setminus\{0\}} P(\bar{a\delta} - a\delta = xc, \bar{a} - a = c) P(\bar{e} - e = yc)\\
&= \frac{1}{q(q-1)} \sum_{y\in\mathbb{F}_q, c\in V\setminus\{0\}} P(\bar{e} - e = yc)\\
&= \frac{1}{q(q-1)} \cdot (q-1) \sum_{c'\in V} P(\bar{e} - e = c')\\
&= \frac{1}{q}.
\end{aligned}$$

□

Proof of Proposition 1. The second claim follows directly from (1). For the first claim, let $\delta = t_i - t_j$. Then $m_j \sim g_\delta + s_0 + \frac{\bar{e}-e}{\bar{a}-a}$, where $g_\delta = \frac{\overline{a\delta}-a\delta}{\bar{a}-a}$. Now the first claim is precisely Lemma 3.

4. *Security of cyclotomic rings with unramified moduli*

In this section we provide some numerical evidence that for cyclotomic fields, the image of a fairly narrow RLWE error distribution modulo an unramified prime ideal $\mathfrak{q}$ of residue degree one or two is practically indistinguishable from uniform, implying that the cyclotomic fields are protected against the family of attacks in this paper.

For simplicity of analysis, we define a helper error distribution (Definition 2) on R with the goal of mimicking the RLWE error distribution. The advantage of this helper distribution is that it admits a closed form formula for a bound on the statistical distance between its reduction modulo $\mathfrak{q}$ and the uniform distribution. This allows for a rigorous bound on the statistical distance. We also generate the actual RLWE samples, run our chi-square attack, and confirm that the errors modulo $\mathfrak{q}$ are indeed uniform.

Let $m \geq 1$ be an integer and let $K = \mathbb{Q}(\zeta_m)$ be the m-th cyclotomic field. Let q be a prime such that $q \equiv 1 \pmod{m}$, so q is unramified in K. Finally, let $\mathfrak{q}$ be a prime ideal above q.

Next, with the aim of simplifying our analysis, we introduce a class of "shifted binomial distributions" indexed by even integers $k \geq 2$, which is then used to generate our modified error distribution.

Definition 1. For an even integer $k \geq 2$, let $\mathcal{V}_k$ denote the distribution over $\mathbb{Z}$ such that for every $t \in \mathbb{Z}$,

$$\mathrm{Prob}(\mathcal{V}_k = t) = \begin{cases} \frac{1}{2^k}\binom{k}{t+\frac{k}{2}} & \text{if } |t| \leq \frac{k}{2} \\ 0 & \text{otherwise} \end{cases}$$

We will abuse notation and also use $\mathcal{V}_k$ to denote the reduced distribution $\mathcal{V}_k \pmod{q}$ over $\mathbb{F}_q$, and let ν_k denote its probability density function. Figure 1 shows a plot of ν_8.

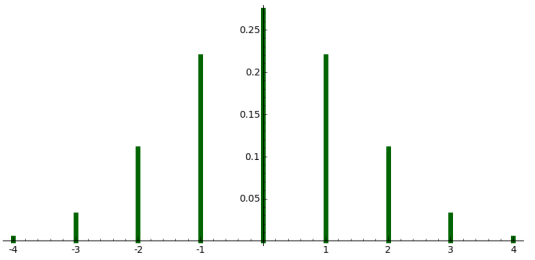

FIGURE 1. *Probability density function of* $\mathcal{V}_8$

DEFINITION 2. Let $k \geq 2$ be an even integer. Then a sample from the helper error distribution $P_{m,k}$ is

$$e = \sum_{i=0}^{n-1} e_i \zeta_m^i,$$

where the coefficients e_i are sampled independently from $\mathcal{V}_k$.

4.1. *Bounding the Distance from Uniform*

We recall the definition and key properties of Fourier transform over finite fields. Suppose f is a real-valued function on $\mathbb{F}_q$. The *Fourier transform* of f is defined as

$$\widehat{f}(y) = \sum_{a \in \mathbb{F}_q} f(a)\overline{\chi_y(a)},$$

where $\chi_y(a) := e^{2\pi i a y/q}$.

Let u denote the probability density function of the uniform distribution over $\mathbb{F}_q$, that is $u(a) = \frac{1}{q}$ for all $a \in \mathbb{F}_q$. Let δ denote the characteristic function of the one-point set $\{0\} \subseteq \mathbb{F}_q$. Recall that the convolution of two functions $f, g : \mathbb{F}_q \to \mathbb{R}$ is defined as $(f * g)(a) = \sum_{b \in \mathbb{F}_q} f(a-b)g(b)$. We list without proof some basic properties of the Fourier transform.

(1) $\widehat{\delta} = qu$; $\widehat{u} = \delta$.
(2) $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$.
(3) $f(a) = \frac{1}{q} \sum_{y \in \mathbb{F}_q} \widehat{f}(y)\chi_y(a)$ (the Fourier inversion formula).

The following is a standard result.

LEMMA 4. *Suppose the random variables F, G are independent random variables with values in $\mathbb{F}_q$, having probability density functions f and g. Then $h = f * g$. In general, suppose $F_1, \ldots, F_n$ are mutually independent random variables in $\mathbb{F}_q$, with probability density functions $f_1, \ldots, f_n$. Let f denote the density function of the sum $F = \sum F_i$, then $f = f_1 * \cdots * f_n$.*

The Fourier transform of ν_k has a nice closed-form formula, as below.

LEMMA 5. *For all even integers $k \geq 2$, $\widehat{\nu_k}(y) = \cos\left(\frac{\pi y}{q}\right)^k$.*

Proof. We have

$$\begin{aligned}
2^k \cdot \widehat{\nu_k}(y) &= \sum_{m=-\frac{k}{2}}^{\frac{k}{2}} \binom{k}{m+\frac{k}{2}} e^{2\pi i y m/q} \\
&= e^{-\pi i y k/q} \sum_{m=-\frac{k}{2}}^{\frac{k}{2}} \binom{k}{m+\frac{k}{2}} e^{2\pi i y (m+k/2)/q} \\
&= e^{-\pi i y k/q} \sum_{m'=0}^{k} \binom{k}{m'} e^{2\pi i y m'/q} \\
&= e^{-\pi i y k/q} (1 + e^{2\pi i y/q})^k \\
&= (e^{-\pi i y/q} + e^{\pi i y/q})^k \\
&= (2\cos(\pi y/q))^k.
\end{aligned}$$

Dividing both sides by 2^k gives the result. □

Next, we concentrate on the reduced distributions $P_{m,k} \pmod{\mathfrak{q}}$. Note that there is a one-to-one correspondence between primitive m-th roots of unity in $\mathbb{F}_q$ and the prime ideals above q in $\mathbb{Q}(\zeta_m)$. Let α be the root corresponding to our choice of $\mathfrak{q}$. Then a sample from $P_{m,k} \pmod{\mathfrak{q}}$ is of the form

$$e_\alpha = \sum_{i=0}^{n-1} \alpha^i e_i \pmod{q},$$

where e_i are independently sampled from $\mathcal{V}_k$. We abuse notations and use e_α to denote its own probability density function.

Lemma 6.

$$\widehat{e_\alpha}(y) = \prod_{i=1}^{n} \cos\left(\frac{\alpha^i \pi y}{q}\right)^k.$$

Proof. This follows directly from Lemma 5 and the independence of the coordinates e_i. □

Proposition 2. *Let $f : \mathbb{F}_q \to \mathbb{R}$ be a function such that $\sum_{a\in\mathbb{F}_q} f(a) = 1$. Then for all $a \in \mathbb{F}_q$,*

$$|f(a) - 1/q| \le \frac{1}{q} \sum_{y\in\mathbb{F}_q, y\neq 0} |\hat{f}(y)|. \tag{4.1}$$

Proof. For all $a \in \mathbb{F}_q$,

$$\begin{aligned} f(a) - 1/q &= f - u(a) \\ &= \frac{1}{q} \sum_{y \in \mathbb{F}_q} (\hat{f}(y) - \widehat{u}(y)) \chi_y(a) \\ &= \frac{1}{q} \sum_{y \in \mathbb{F}_q} (\hat{f}(y) - \delta(y)) \chi_y(a) \\ &= \frac{1}{q} \sum_{y \in \mathbb{F}_q, y \neq 0} \hat{f}(y) \chi_y(a). \qquad (\text{since } \hat{f}(0) = 1) \end{aligned}$$

Now the result follows from taking absolute values on both sides, and noting that $|\chi_y(a)| \leq 1$ for all a and all y. □

Taking $f = e_\alpha$ in Proposition 2, we immediately obtain

THEOREM 2. *The statistical distance between e_α and u satisfies*

$$d(e_\alpha, u) \leq \frac{1}{2} \sum_{y \in \mathbb{F}_q, y \neq 0} |\widehat{e_\alpha}(y)|. \tag{4.2}$$

Now let $\epsilon(m, q, k, \alpha)$ denote the right hand side of (4.2), i.e.,

$$\epsilon(m, q, k, \alpha) = \frac{1}{2} \sum_{y \in \mathbb{F}_q, y \neq 0} \prod_{i=0}^{n-1} \cos\left(\frac{\alpha^i \pi y}{q}\right)^k.$$

To take into account all prime ideals above q, we let α run through all primitive m-th roots of unity in $\mathbb{F}_q$ and define

$$\epsilon(m, q, k) := \max\{\epsilon'(m, q, k, \alpha) : \alpha \text{ has order } m \text{ in } \mathbb{F}_q\}.$$

If $\epsilon(m, q, k)$ is negligibly small, the distribution $P_{m,k} \pmod{\mathfrak{q}}$ will be computationally indistinguishable from uniform.

There is a heuristic argument as to why one expects $\epsilon(m, q, k, \alpha)$ to be small. Each term in the summand is a product of form $\prod_{i=0}^{n-1} \cos\left(\frac{\alpha^i \pi y}{q}\right)^k$. For each $0 \neq y \in \mathbb{F}_q$, if one assumes the elements α^i are distinct and uniformly distributed in $\mathbb{F}_q$, it is very likely that $\alpha^i y$ is close to $q/2$ for at least some values of i, making the product of cosines small. Our goal now is to give a rigorous proof to this heuristic. Specifically, we want to give an upper bound for

$$\epsilon(m, q, k) = \frac{1}{2} \sum_{y \in \mathbb{F}_q, y \neq 0} \prod_{i=0}^{m-1} \cos\left(\frac{\alpha^i \pi y}{q}\right)^k.$$

Here m is the index of the cyclotomic field $K = \mathbb{Q}(\zeta_m)$, q is a prime congruent to 1 modulo m, k is a positive even integer, and α is a primitive m-th root of unity in $\mathbb{F}_q$. Note that the product is independent of the choice of α.

We will prove the following theorem.

THEOREM 3. *Let q, m be positive integers such that q is a prime, $q \equiv 1 \mod m$ and $q < m^2$. Let $\beta = \frac{1 + \frac{\sqrt{q}}{m}}{2}$; then $0 < \beta < 1$ and*

$$\epsilon(m, q, k) \leq \frac{q-1}{2} \beta^{\frac{km}{2}}.$$

In particular, if $\beta^{k/4} < 0.5$, then the theorem says that $\epsilon(m, q, k) = O(q2^{-m})$ as $m \to \infty$.

COROLLARY 1. *The statistical distance between $P_{m,k}$ modulo $\mathfrak{q}$ and a uniform distribution is bounded above, independently of the choice of $\mathfrak{q}$ above q, by*

$$\frac{q-1}{2}\left(\frac{1+\frac{\sqrt{q}}{m}}{2}\right)^{\frac{km}{2}}.$$

To prepare proving the theorem, we set up some notations of Shparlinski in [**Shp95**]. Let $\Omega = (\omega_k)_{k=1}^{\infty}$ be a sequence of real numbers and let m be a positive integer. We define the following quantities:

– $L_\Omega(m) = \prod_{k=1}^{m}(1 - \exp(2\pi i\omega_k))$

– $S_\Omega(m) = \sum_{k=1}^{m}\exp(2\pi i\omega_k)$.

The following lemma is a special case of [**Shp95**, Theorem 1.3].

LEMMA 7.

$$L_\Omega(m) \le 2^{m/2}(1 + S_\Omega(m)/m)^{m/2}.$$

Now we specialize the above discussion to our situation. We take the sequence to be $\omega_k = \frac{\alpha^k y}{q} + 1/2$, where we abuse notations and let α also denote a lift of $\alpha \in \mathbb{F}_q$ to $\mathbb{Z}$. We compute that

$$L_\Omega(m) = 2^m\left|\prod_{i=0}^{m-1}\cos\left(\frac{\alpha^i\pi y}{q}\right)\right|$$

and

$$S_\Omega(m) = -\sum_{i=1}^{m}\exp\left(\frac{\alpha^i\pi y}{q}\right).$$

The following fact is a standard bound on exponential sums.

FACT 1. Let α be an primitive m-th root of unity in $\mathbb{F}_q$ and let $y \in \mathbb{F}_q$ be nonzero. Then

$$\left|\sum_{i=1}^{m}\exp\left(\frac{\alpha^i\pi y}{q}\right)\right| \le q^{1/2}.$$

Proof of Theorem 3. Fact 1 implies that $|S_\Omega(m)| \le q^{1/2}$. Then using Lemma 7, we get

$$\left|\prod_{i=0}^{m-1}\cos\left(\frac{\alpha^i\pi y}{q}\right)\right| \le \beta^{m/2}$$

for β in the theorem and for any nonzero $y \in \mathbb{F}_q$. Our result in the theorem now follows from taking both sides to k-th power and summing over y.

4.2. *Numerical Distance from Uniform*

We have computed $\epsilon(m, q, k)$ for various choices of parameters. Smaller values of ϵ imply that the error distribution looks more uniform when transferred to $R/\mathfrak{q}$, rendering the instance of RLWE invulnerable to the attacks in [**CLS15**].

The following is a table of data. For each instance in the table, we also generated the actual RLWE samples (where we fixed $r_0 = \sqrt{2\pi}$) and ran the chi-square attack of [**CLS15**] using confidence level $\alpha = 0.99$. The column labeled "χ^2" contains the χ^2 values we obtained, and the column labeled "uniform?" indicates whether the reduced errors are uniform. We can see from data how the practical situation agrees with our analysis on the approximated distributions.

TABLE 3. *Values of $\epsilon(m, q, 2)$ and the χ^2 values*

m	n	q	$-[\log_2(\epsilon(m,q,2))]$	χ^2	uniform?
96	32	193	35	231.6	yes
55	40	331	44	308.8	yes
160	64	641	55	658.0	yes
101	100	1213	177	1254.4	yes
244	120	1709	230	1721.2	yes
256	128	3329	194	3350.0	yes
197	196	3547	337	3475.2	yes
512	256	10753	431	10732.8	yes

The data in Table 3 shows that when $n \geq 100$ and the size of the modulus q is polynomial in n, the statistical distances between $P_{m,k} \pmod{\mathfrak{q}}$ and the uniform distribution are negligibly small. Also, note that we fixed $k = 2$, and the epsilon values becomes even smaller when k increases.

It is possible to generalize our discussion in this section to primes of arbitrary residue degree f, in which case the Fourier analysis will be performed over the field $\mathbb{F}_{q^f}$. The only change in the definitions would be $\chi_y(a) = e^{\frac{2\pi i \operatorname{Tr}(ay)}{q}}$. Here $\operatorname{Tr} : \mathbb{F}_{q^f} \to \mathbb{F}_q$ is the trace function. Similarly, we have

$$\widehat{e'_\alpha}(y) = \prod_{i=1}^{n} \cos\left(\frac{\pi \operatorname{Tr}(\alpha^i y)}{q}\right)^k.$$

Table 4 contains some data for primes of degree two.

TABLE 4. *Values of $\epsilon(m, q, 2)$ for primes of degree two*

m	n	q	$-[\log_2(\epsilon(m,q,2))]$
64	32	383	31
63	36	881	33
55	40	109	48
53	52	211	61
512	256	257	263

References

AD97 Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 284–293. ACM, 1997.

BGV12 Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 309–325. ACM, 2012.

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

Bra12 Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical gapsvp. In *Advances in Cryptology–CRYPTO 2012*, pages 868–886. Springer, 2012.

BV11 Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from Ring-LWE and security for key dependent messages. In *Advances in Cryptology–CRYPTO 2011*, pages 505–524. Springer, 2011.

BV14 Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM Journal on Computing*, 43(2):831–871, 2014.

CIV Wouter Castryck, Ilia Iliashenko, and Frederik Vercauteren. Provably weak instances of Ring-LWE revisited. To appear in Eurocrypt 2016.

CLS15 Hao Chen, Kristin Lauter, and Katherine E. Stange. Attacks on search-RLWE. Cryptology ePrint Archive, Report 2015/971, 2015. `http://eprint.iacr.org/`.

EHL14 Kirsten Eisenträger, Sean Hallgren, and Kristin Lauter. Weak instances of PLWE. In *Selected Areas in Cryptography–SAC 2014*, pages 183–194. Springer, 2014.

ELOS15 Yara Elias, Kristin Lauter, Ekin Ozman, and Katherine Stange. Provably weak instances of ring-lwe. In *Advances in Cryptology – CRYPTO 2015*, volume 9215 of *Lecture Notes in Comput. Sci.*, pages 63–92. Springer, Heidelberg, 2015.

HPSS08 Jeffrey Hoffstein, Jill Pipher, Joseph H Silverman, and Joseph H Silverman. *An introduction to mathematical cryptography*, volume 1. Springer, 2008.

LATV12 Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1219–1234. ACM, 2012.

LPR13 Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM (JACM)*, 60(6):43, 2013.

Mar77 Daniel A Marcus. *Number fields*, volume 18. Springer, 1977.

MR07 Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.

Reg09 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):34, 2009.

Shp95 Igor E Shparlinski. On some characteristics of uniformity of distribution and their applications. In *Computational Algebra and Number Theory*, pages 227–241. Springer, 1995.

SS11 Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Advances in Cryptology–EUROCRYPT 2011*, pages 27–47. Springer, 2011.