Attacks on the Search-RLWE problem with small errors

Hao Chen[1], Kristin Lauter[2], and Katherine E. Stange[3]

[1] Microsoft Research, Redmond, USA
haoche@microsoft.com
[2] Microsoft Research, Redmond, USA
klauter@microsoft.com
[3] University of Colorado, Boulder, USA
kstange@math.colorado.edu

Abstract. The Ring Learning-With-Errors (RLWE) problem shows great promise for post-quantum cryptography and homomorphic encryption. We describe a new attack on the non-dual search RLWE problem with small error widths, using ring homomorphisms to finite fields and the chi-square statistical test. In particular, we identify a "subfield vulnerability" (Section 5.2) and give a new attack which finds this vulnerability by mapping to a finite field extension and detecting non-uniformity with resepect to the number of elements in the subfield. We use this attack to give examples of vulnerable RLWE instances in Galois number fields. We also extend the well-known search-to-decision reduction result to Galois fields with any unramified prime modulus q, regardless of the residue degree f of q, and we use this in our attacks. The time complexity of our attack is $O(nq^{2f})$, where n is the degree of K and f is the *residue degree* of q in K. We also show an attack on the non-dual (resp. dual) RLWE problem with narrow error distributions in prime cyclotomic rings when the modulus is a ramified prime (resp. any integer). We demonstrate the attacks in practice by finding many vulnerable instances and successfully attacking them. We include the code for all attacks.

Key words: attacks, RLWE, cryptanalysis.

1 Introduction

The Ring Learning-with-Errors (RLWE) problem, proposed in [18], is a variant of the Learning-with-Errors (LWE) problem, and is an active research area in lattice based cryptography, and a candidate for post-quantum cryptography. It has drawn increased attention because it can be used for homomorphic encryption ([2,3,4,5,13,17,24]). The problem, which comes in *search* and *decision* variants, is based on the geometry of ideal lattices, in particular the rings of integers of number fields, or their duals.

It is of great importance to understand the security of RLWE. The first piece of the puzzle is provided by proofs of security [18]. However, it is also important to mount direct attacks on the problem and its variants, in order to illuminate the protective properties of the provably secure instances, the dangers of deviating from established parameters, and the practical runtimes for certain parameter sizes. This work is part of that programme, which has seen much recent interest, e.g. [11,12]. An eprint version of the current paper [8] has already generated much follow-up work [6,7,9,22]. In this paper we provide a brief overview of past work, and then present attacks which are novel in their mathematical underpinnings (based on new homomorphisms to finite fields of higher degree which detect in particular "subfield vulnerabilities", see Section 5.2). We also discuss the underlying number theory and geometry of these attacks to provide a framework for future work.

An instance of the RLWE problem is determined by a choice of a number field K and a prime q called the *modulus*, along with an error distribution. The authors of [18] proved a reduction from certain hard lattice problems to an instance of search RLWE involving a continuous Gaussian error distribution modulo the dual ideal $R^\vee$ of the ring of integers R of K. Ducas and Durmus proposed a non-dual variant of RLWE in the cyclotomic setting and proved its hardness in [10]. Also in [18], a search-to-decision reduction was proved for RLWE problems in cyclotomic fields and modulus q which splits completely. This reduction was then generalized in [11] to hold for general Galois number fields where q splits. As an auxiliary result in this paper, we generalize this search-to-decision reduction to work for the case of unramified modulus q of arbitrary degree.

The non-dual variant of RLWE generates the error distribution as a discrete Gaussian on the ring of integers R under the canonical embedding, instead of in the image of the dual ideal. The dual and non-dual variants are equivalent up to a change in the error distribution (see Section 2). For the non-dual variant of RLWE, the authors of [12] proposed an attack on the *decision* RLWE problem. The attack makes use of a modulus q of residue degree 1, giving a ring homomorphism $\rho : R \to \mathbb{F}_q$ (so that it could be called a *mod q* attack, although it differs from [20,21]). The attack works when, with overwhelming probability, the image of the RLWE error distribution under the map ρ takes values only in a small subset of $\mathbb{F}_q$. The authors of [12] then gave an infinite family of examples vulnerable to the attack. Unfortunately, the vulnerable number fields in [12] are not Galois. Hence, the search-to-decision reduction theorem does not apply, and the attack can not be directly used to solve the search variant of RLWE for those instances.

In this paper, we generalize the attack of [12] to certain Galois number fields and moduli of higher degree. As a result, we have an attack on the *search* RLWE problem and an implementation of the attack on concrete RLWE instances, including the search-to-decision reduction. Our attack is new in two major ways: first, the attack considers ring homomorphisms from $R \to \mathbb{F}_{q^f}$, for $f > 1$, instead of just homomorphisms from $R \to \mathbb{F}_q$ (so it is no longer 'mod q'); second, the error distribution is distinguished from random (i.e. from the uniform distribution) using the statistical chi-squared test, instead of relying on the values of the error polynomial to be small or in a small subset. The attack aims at an intermediate problem used in the search-to-decision proof of [18], which is to recover the secret modulo a prime ideal (denoted $\mathrm{SRLWE}(\mathcal{R}, \mathfrak{q})$; see Definition 7). The time complexity of our attack is $O(nq^{2f})$, where n is the degree of K and f is the residue degree of q in K.

Importantly, we also show an attack on prime cyclotomic rings under certain assumptions on the modulus and error rate, which succeeds with high probability and with surprising efficiency. First, we give attacks for the *decision* version of the non-dual variant of RLWE considered in [12], when the modulus q is equal to the unique ramified prime p. For example, we show that in dimension $n = 808$, we can attack an RLWE instance in the cyclotomic ring $\mathbb{Q}(\zeta_{809})$ effectively in 35 seconds, where the modulus is 809. This opens up the question of whether general cyclotomic fields are safe for cryptography, depending on whether modulus switching can be used to transfer this attack from the ramified modulus to other larger moduli which are used in practice. In addition, we attack the *decision* version of the *dual* RLWE problem on the p-th cyclotomic field with *arbitrary* modulus q, assuming that the width r of the error distribution is around $\frac{1}{\sqrt{p}}$.

The error widths for which our attacks work are below those required by the security proof of [18], which requires $r = \omega(\sqrt{\log n})$. In particular, this work does not affect the hardness results of [18]. On the other hand, in practice, in implementations of homomorphic encryption systems based on the hardness of RLWE, it has been common practice to use small errors to improve efficiency for the systems. We show in this work that for errors in the width range below provably secure but above linear algebra vulnerable (errorless LWE), the security of RLWE depends in an interesting way on the choice of ring and modulus. To be more specific, the geometry of the ring of integers and the manner in which certain prime ideals exist as sublattices are important factors (see Section 5.2). Finally, it is important to note that most implementations of RLWE-based schemes use exclusively 2-power cyclotomic rings, on which our attacks are not effective. Hence the impact of our attacks on the security of existing practical implementations of RLWE-based homomorphic encryption schemes is limited.

Auxiliary results we present include several stand-alone items of possibly independent interest: we prove a search-to-decision reduction for Galois fields which applies for any unramified modulus q, regardless of the residue degree of q (this relies heavily on Galois theory and Galois fields are the largest class to which we expect this to apply). We also present some heuristic arguments as to whether modulus switching techniques are likely to be successfully combined with our attacks.

We end this section with a table summarizing what is known about the security of RLWE for certain choices of number fields. The first table deals with the continuous dual version. For comparison, we normalize the error width: let $\tilde{r} = r \cdot |d_K|^{1/2n}$, where d_K is the discriminant of the number field K.

Table 1. Security of dual RLWE

field	modulus	$\tilde{r}$	security
$\mathbb{Q}(\zeta_m)$	$q \equiv 1 \mod m$, $q = poly(m)$	$\omega(\sqrt{\log n}) \cdot \Theta(\sqrt{n})$	decision is secure [18]
Any	$q = poly(n)$	$\omega(\sqrt{\log n}) \cdot \lvert d_K \rvert^{1/2n}$	search is secure [18]
$\mathbb{Q}(\zeta_p)$	any	~ 1	decision is not secure (this paper)

The second table deals with the non-dual discrete version. Here we normalize by $\tilde{r} = r/|d_K|^{1/2n}$. The heuristic expectation is that when $\tilde{r} = \Omega(\sqrt{n})$ and $q = poly(n)$, decision RLWE problem should be hard.

Table 2. Security of non-dual RLWE

field	modulus	$\tilde{r}$	security
$\mathbb{Q}(\sqrt[n]{1-q})$	$poly(n)$	~ 1	decision is not secure [12]
$\mathbb{Q}(\sqrt[n]{1-q})$	$poly(n)$	~ 1	search is not secure [7]
certain $\mathbb{Q}(\zeta_m)^H$	$poly(n)$ w. properties	~ 1	search is not secure (this paper)
$\mathbb{Q}(\zeta_p)$	p	~ 1	decision is not secure (this paper)
certain $\mathbb{Q}(\zeta_p, \sqrt{d})$	$poly(n)$ w. properties	$o(\sqrt{d/p})$	search is not secure [9]

1.1 Organization

In Section 2, we review some definitions related to the RLWE problems. In Section 3, we prove a search-to-decision reduction for Galois extensions K and unramified moduli. In Section 4, we introduce an attack on non-dual RLWE problems based on the chi-square statistical test, which directly generalizes the attack in [12]. In Section 5, we give examples of subfields of cyclotomic fields vulnerable to our new attack, where the modulus q has residue degree two. In Section 6, we give attacks on the non-dual RLWE in prime cyclotomic fields when the modulus is the unique ramified prime and dual RLWE for any modulus, assuming the errors are sufficiently narrow. In Section 7, we consider the possibility of combining modulus switching with our attack.

All computations in this paper were performed in Sage [25]. All the relevant code is available and can be found at `https://github.com/haochenuw/GaloisRLWE`.

Acknowledgements

The authors thank Chris Peikert for helpful discussions, and the anonymous referees for helpful suggestions. The third author was supported by NSF grant DMS-1643552.

2 Background

Let K be a number field of degree n with ring of integers R, and let $\sigma_1, \cdots, \sigma_n$ be the embeddings of K into the field of complex numbers. We define the *adjusted canonical embedding* of K as follows: Let r_1, r_2 denote the number of real embeddings and conjugate pairs of complex embeddings of K. Without loss of generality, assume $\sigma_1, \cdots, \sigma_{r_1}$ are the real embeddings of K and $\sigma_{r_1+r_2+j} = \overline{\sigma_{r_1+j}}$ for $1 \leq j \leq r_2$. Then the adjusted canonical embedding is the following map:

$$\iota : K \to \mathbb{R}^n : \quad x \mapsto \begin{bmatrix} \sigma_1(x) \\ \vdots \\ \sigma_{r_1}(x) \\ \sqrt{2}\,\mathrm{Re}(\sigma_{r_1+1})(x) \\ \sqrt{2}\,\mathrm{Im}(\sigma_{r_1+1})(x) \\ \vdots \\ \sqrt{2}\,\mathrm{Re}(\sigma_{r_1+r_2})(x) \\ \sqrt{2}\,\mathrm{Im}(\sigma_{r_1+r_2})(x). \end{bmatrix} \tag{1}$$

It turns out that $\Lambda_R := \iota(R)$ is a lattice in $\mathbb{R}^n$. Let $w = (w_1, \cdots, w_n)$ be an integral basis for R. The embedding matrix of w, denoted by A_w, is the n-by-n matrix whose i-th column is $\iota(w_i)$. Note that the columns of A_w form a basis for the lattice Λ_R.

For $\sigma > 0$, define the Gaussian function $\rho_\sigma : \mathbb{R}^n \to (0, 1]$, depending on the usual inner product on $\mathbb{R}^n$, to be $\rho_\sigma(x) = e^{-||x||^2/2\sigma^2}$.

Definition 1. *For a lattice $\Lambda \subset \mathbb{R}^n$ and $\sigma > 0$, the discrete Gaussian distribution on Λ with parameter σ is:*

$$D_{\Lambda,\sigma}(x) = \frac{\rho_\sigma(x)}{\sum_{y \in \Lambda} \rho_\sigma(y)}, \forall x \in \Lambda.$$

Equivalently, the probability of sampling any lattice point x is proportional to $\rho_\sigma(x)$.

We follow [12] in setting up the non-dual RLWE problem for general number fields. In particular, the error distribution we use is a spherical discrete Gaussian distribution on Λ_R.

Definition 2. *A (non-dual) RLWE instance is a tuple $\mathcal{R} = (K, q, \sigma, s)$, where K is a number field with ring of integers R, q is a prime, $\sigma > 0$ is a positive real number, and $s \in R/qR$ is called the secret.*

Suppose $\mathcal{R} = (K, q, \sigma, s)$ is an RLWE instance and let R be the ring of integers of K. The *error distribution* of $\mathcal{R}$ is the discrete Gaussian distribution $D_{\Lambda_R,\sigma}$.

Let R_q denote the quotient ring R/qR; then a (non-dual) RLWE sample is a pair

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

where the first coordinate a is chosen uniformly at random in R_q, and e is sampled from the error distribution and considered modulo qR. The reader unfamiliar with this problem should consider this analogous to a discrete logarithm pair $(g, g^s) \in \mathbb{F}_q \times \mathbb{F}_q$, where s is a secret exponent.

Definition 3 (Search RLWE). *Let $\mathcal{R}$ be an RLWE instance. The* search RLWE problem*, denoted by* SRLWE($\mathcal{R}$)*, is to discover s given access to arbitrarily many independent samples (a, b).*

Definition 4 (Decision RLWE). *Let $\mathcal{R}$ be an RLWE instance. The* decision RLWE problem*, denoted by* DRLWE($\mathcal{R}$)*, is to distinguish between the same number of independent samples in two distributions on $R_q \times R_q$. The first is the RLWE distribution of $\mathcal{R}$, and the second consists of uniformly random and independent samples from $R_q \times R_q$.*

Remark 5. As pointed out in [12], when analyzing the error distribution, one needs to take into account the sparsity of the lattice Λ_R, which is measured by its covolume $\det(\Lambda_R) = \sqrt{|\operatorname{disc}(K)|}$. In light of this, we define a relative version of the standard deviation parameter: $\sigma_0 = \frac{\sigma}{|\operatorname{disc}(K)|^{\frac{1}{2n}}}$.

Remark 6. There are different approximate algorithms to sample from discrete Gaussian distributions on lattices. In this paper, we choose to use the sampling algorithm developed in [14].

We now discuss dual RLWE and its relation to non-dual RLWE. In dual RLWE, the secret s lies in $R_q^\vee := R^\vee/qR^\vee$, where $R^\vee$ is the dual ideal of R, and the error e is sampled from $R^\vee$ with discrete spherical Gaussian distribution with width $r = \sqrt{2\pi}\sigma$. Therefore the RLWE samples are of the form

$$(a, b = as + e) \in R_q \times R_q^\vee.$$

The dual and non-dual versions of the RLWE problem are very closely related when the dual ideal is principal: in this case, R and $R^\vee$ are related by a scaling factor (which may alter a spherical Gaussian to an ellipsoidal one). Even when $R^\vee$ is not principal, we have $R \subseteq k_1 R^\vee$ and $R^\vee \subseteq k_2 R$ for some constants k_1 and k_2, so that a problem in one formulation can be reduced to a problem in the other, with a different error distribution. Several of the non-dual examples of this paper are known to have principal dual ideal [22]. In particular, our attack on the ramified prime for cyclotomic fields is translated to the dual situation in Section 6.2. The full class of elliptic Gaussians (not just spherical) is also considered in the security reductions of [18].

Finally, there is a *continuous* version of RLWE which is more amenable to security reductions. Since one can always discretize, the continuous version reduces to the discrete one presented here, which is more practical for applications.

3 Search-to-Decision Reduction

In [11], the search-to-decision reduction of [18] is extended to RLWE for Galois number fields, where q is an unramified prime of degree one. The approach is via an intermediate problem, denoted $\mathfrak{q}_i$-LWE in [18]. In this section, we extend this result to primes $\mathfrak{q}$ of arbitrary residue degree. Our intermediate problem, which we denote by $\mathrm{SRLWE}(\mathcal{R}, \mathfrak{q})$, is the same as $\mathfrak{q}_i$-LWE, and it amounts to finding the secret modulo the prime $\mathfrak{q}$. The Galois group allows us to bootstrap this piece of information to discover the full secret.

The attack in Section 4 targets $\mathrm{SRLWE}(\mathcal{R}, \mathfrak{q})$ and hence, by the results of this section, will solve Search RLWE. In Section 5, we demonstrate the attack on Search RLWE in practice.

Definition 7. *Let $\mathcal{R} = (K, q, \sigma, s)$ be an RLWE instance and let $\mathfrak{q}$ be a prime of K lying above q. The problem* $\mathrm{SRLWE}(\mathcal{R}, \mathfrak{q})$ *is to determine $s \pmod{\mathfrak{q}}$, given access to arbitrarily many independent samples $(a, b) \leftarrow \mathcal{R}$.*

We recall some facts from algebraic number theory in the following lemma.

Lemma 8. *Let $K/\mathbb{Q}$ be a finite Galois extension of degree n with ring of integers R, and let q be a prime unramified in K. Then there exists a unique divisor g of n and a set of g distinct prime ideals $\mathfrak{q}_1, \cdots, \mathfrak{q}_g$ of R such that:*

1. *$qR = \prod_{i=1}^{g} \mathfrak{q}_i$,*
2. *the quotient $R/\mathfrak{q}_i$ is a finite field of cardinality q^f for each i, where $f = \frac{n}{g}$,*
3. *there is a canonical isomorphism of rings*

$$R_q \cong R/\mathfrak{q}_1 \times \cdots \times R/\mathfrak{q}_g, \tag{2}$$

4. *the Galois group acts transitively on the ideals $\mathfrak{q}_1, \ldots, \mathfrak{q}_g$ and this action descends to an action on R_q which permutes the corresponding factors in* (2) *in the same way.*

The number f in the above lemma is called the *residue degree* of q in K. Note that the prime q splits completely in K if and only if its residue degree is one.

Theorem 9. *Let $\mathcal{R} = (K, q, \sigma, s)$ be an RLWE instance such that $K/\mathbb{Q}$ is Galois of degree n and q is unramified in K with residue degree f. Let $\mathscr{A}$ be an oracle which solves* $\mathrm{SRLWE}(\mathcal{R}, \mathfrak{q})$ *using a list of m samples modulo $\mathfrak{q}$. Let S be a set of m RLWE samples in $R_q \times R_q$. Then the problem* $\mathrm{SRLWE}(\mathcal{R})$ *can be solved using S by n/f calls to the oracle $\mathscr{A}$, $2mn/f$ reductions $R_q \to R/\mathfrak{q}$, and $2mn/f$ evaluations of a Galois automorphism on R_q.*

Proof. The Galois group $G = \mathrm{Gal}(K/\mathbb{Q})$ acts on the set $\{\mathfrak{q}_1, \cdots, \mathfrak{q}_g\}$ transitively. Hence for each i, there exists $\sigma_i \in \mathrm{Gal}(K/\mathbb{Q})$, such that $\sigma_i(\mathfrak{q}) = \mathfrak{q}_i$. Next, remark that $\sigma_i^{-1}(S)$ is itself a set of RLWE samples in $R_q \times R_q$, since the action of Galois is an isometry of the Minkowski embedding. (Here it is essential that we consider only spherical Gaussians.) Furthermore, the secret for this set of samples is $\sigma_i^{-1}(s)$. Then we call the oracle $\mathscr{A}$ on the input $(\sigma_i^{-1}(S) \pmod{\mathfrak{q}}, \mathfrak{q})$. The algorithm will output $\sigma_i^{-1}(s) \pmod{\mathfrak{q}}$, from which we can recover $s \pmod{\mathfrak{q}_i}$ using σ_i. We do this for all $1 \leq i \leq g = n/f$ and use (2) of Lemma 8 to recover s. □

In particular, if the number of samples m is polynomial in n and the time taken to evaluate Galois automorphisms on a single sample is also polynomial in n, then Theorem 9 gives a polynomial time reduction from SRLWE($\mathcal{R}$) to SRLWE($\mathcal{R}, \mathfrak{q}$).

Remark 10. For a proper runtime analysis of the reduction, one must examine the implementation, in particular with regards to Galois automorphisms. The runtime for evaluating an automorphism depends rather strongly on the instance and on the way ring elements are represented. For example, for subfields of cyclotomic fields represented with respect to normal integral bases, the Galois automorphisms are simply permutations of the coordinates, so the time needed to apply these automorphisms is trivial.

The search-to-decision reduction will follow from the lemma below.

Lemma 11. *There is a probabilistic polynomial time reduction from* SRLWE($\mathcal{R}, \mathfrak{q}$) *to* DRLWE($\mathcal{R}$).

Proof. This is a rephrasing of [18, Lemma 5.9 and Lemma 5.12]. □

Corollary 12. *Suppose $\mathcal{R}$ is an RLWE instance where K is Galois and q is an unramified prime in K. Then there is a probabilistic polynomial-time reduction from* SRLWE($\mathcal{R}$) *to* DRLWE($\mathcal{R}$).

4 The Chi-square Attack

In this section, we extend the $f(1) \equiv 0 \pmod{q}$ attack of [11] and the root-of-small-order attack of [12]. These attacks can be viewed as examples of a more general principle, as follows. Suppose one has a ring homomorphism

$$\phi : R_q \to F$$

where F is a finite field, and where two properties hold:

1. F is small enough that its elements can be examined exhaustively; and
2. the error distribution on R_q, transported by ϕ to F, is detectably non-uniform.

Then the attack on decision RLWE is as follows:

1. Transport the samples (a, b) in $R_q \times R_q$ to $F \times F$ via ϕ.
2. Loop through possible guesses for the image of the secret, $\phi(s)$, in F.
3. For each guess g, compute the distribution of $\phi(b) - \phi(a)g$ on the available samples. Note that if we let $g^* = \phi(s)$ denote the true value,

$$\phi(b) - \phi(a)g = (\phi(b) - \phi(a)g^*) + \phi(a)(g - g^*) = \phi(e) + \phi(a)(g - g^*),$$

   which equals $\phi(e)$ if the guess is correct, and looks uniform otherwise.
4. If the samples are RLWE samples with secret s and $g = \phi(s)$, then this distribution will follow the error distribution, which will look non-uniform.
5. If all such distributions look uniform, then the samples were uniform, not RLWE, samples.

The fact that ϕ is a ring homomorphism is essential in guaranteeing that for the correct guess, the distribution in question is the image of the error distribution. The only ring homomorphisms from R_q to a finite field are given by reduction modulo a prime ideal $\mathfrak{q}$ lying above q in R.

4.1 Chi-square Test for Uniform Distribution

We briefly review the properties and usage of the chi-square test for uniform distributions over a finite set S. We partition S into r subsets $S = \bigsqcup_{j=1}^{r} S_j$, called *bins*. Suppose there are M samples $y_1, \ldots, y_M \in S$. For each $1 \leq j \leq r$, we compute the expected number of samples in the j-th subset: $c_j := \frac{|S_j| M}{|S|}$. Then we compute the actual number of samples in S_j, i.e., $t_j := |\{1 \leq i \leq r : y_i \in S_j\}|$. Finally, the χ^2 value is computed as

$$\chi^2(S, y) = \sum_{j=1}^{r} \frac{(t_j - c_j)^2}{c_j}.$$

Suppose the samples are drawn from the uniform distribution on S. Then the χ^2 value follows the chi-square distribution with $(r-1)$ degrees of freedom, which we denote by χ^2_{r-1}. Let $\mathcal{F}_{r-1}(x)$ denote its cumulative distribution function. For the chi-square test, we choose a confidence level parameter $\alpha \in (0,1)$ and compute $\delta = \mathcal{F}_{r-1}^{-1}(\alpha)$. Then we reject the hypothesis that the samples are drawn from the uniform distribution if $\chi^2(S, y) > \delta$.

If P, Q are two probability distributions on the set S, then their *statistical distance* is defined as $d(P,Q) = \frac{1}{2}\sum_{t \in S} |P(t) - Q(t)|$. For convenience, we also define the l_2 *distance* between P and Q as $d_2(P,Q) = (\sum_{t \in S} |P(t) - Q(t)|^2)^{\frac{1}{2}}$. We have the inequality $d(P,Q) \leq \frac{\sqrt{|S|}}{2} d_2(P,Q)$.

Remark 13. We chose to use the chi-square test for our attack since we are distinguishing a known distribution (uniform on $R/\mathfrak{q}$) from an unkown distribution (discrete Gaussians mod $\mathfrak{q}$). If the latter distribution is efficiently computable, then one might switch to other statistical tests, e.g., the Neyman-Pearson test, for better results.

4.2 The Chi-square Attack on SRLWE($\mathcal{R}, \mathfrak{q}$)

Let $\mathcal{R}$ be an RLWE instance with error distribution $D_{\Lambda_R,\sigma}$ and $\mathfrak{q}$ be a prime ideal above q. The basic idea of our attack relies on the assumption that the distribution $D_{\Lambda_R,\sigma} \mod \mathfrak{q}$ is distinguishable from the uniform distribution on the finite field $F = R/\mathfrak{q}$. More precisely, the attack loops through all q^f possible values $\bar{s} = s \pmod{\mathfrak{q}}$, and for each guess s', it computes the values $\bar{e}' = \bar{b} - \bar{a}s' \pmod{\mathfrak{q}}$ for every sample $(a,b) \in S$. If the guess is wrong, or if the samples are taken from the uniform distribution in $(R_q)^2$, the values $\bar{e}'$ would be uniformly distributed in F and it is likely to pass the chi-square test. On the other hand, if the guess is correct, then we expect the test on the errors $\bar{e}'$ to reject the null hypothesis. Let $N := q^f$ denote the cardinality of F. We remark that as a general rule of thumb for the chi-square test, we need to generate at least $5N$ samples.

For the attack to be successful, we need the $(N-1)$ tests corresponding to wrong guesses of $s \pmod{\mathfrak{q}}$ to pass, and the one test corresponding to the correct guess to be rejected. For this purpose, we need to choose the confidence level α to be close enough to one (a reasonable choice is $\alpha = 1 - \frac{1}{10N}$). The detailed attack is described in Algorithm 1. Let $\mathcal{F}_{N-1}(x)$ denote the cumulative distribution function of χ^2_{N-1}.

Remark 14. For simplicity of exposition, we use N bins in Algorithm 1, that is one element per bin. In some situations, it might be advantageous to choose the bins differently.

The time complexity of the attack is $O(nq^{2f})$ since there are q^f possible values for $s \pmod{\mathfrak{q}}$, each reduction modulo $\mathfrak{q}$ takes $O(n)$ to compute, and the number of samples needed is $O(q^f)$. The correctness of the attack is captured in Theorem 15 below. For $\lambda \in \mathbb{R}$ and $d \in \mathbb{Z}$, we use $\mathcal{F}_{d,\lambda}(x)$ to denote the cumulative distribution function of the noncentral chi-square distribution with degree of freedom d and parameter λ.

Theorem 15. *Let $\mathcal{R} = (K, q, s, \sigma)$ be an RLWE instance. Suppose $\mathfrak{q}$ be a prime ideal in K above q, and let Δ denote the statistical distance between the distribution $D_{\Lambda_R,\sigma} \mod \mathfrak{q}$ and the uniform distribution on $R/\mathfrak{q}$. Let M be the number of samples used in Algorithm 1, and let $\lambda = 4M\Delta^2$. Let $0 < \alpha < 1$ and let $\delta = \mathcal{F}_{N-1}^{-1}(\alpha)$. If p is the probability of success of the attack in Algorithm 1, then*

$$p \geq \alpha^{N-1}(1 - \mathcal{F}_{N-1;\lambda}(\delta)).$$

Algorithm 1 chi-square attack on SRLWE($\mathcal{R}, \mathfrak{q}$)

Input: $\mathcal{R} = (K, q, \sigma, s)$ – an RLWE instance; R – the ring of integers of K; $\mathfrak{q}$ – a prime ideal in K above q; $F = R/\mathfrak{q}$ – the residue field of $\mathfrak{q}$; $N = q^f$ – the cardinality of F; $\mathcal{S}$ – a collection of M ($M = \Omega(N)$) RLWE samples from $\mathcal{R}$; $0 < \alpha < 1$ – the confidence level.

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

```
δ ← F_{N−1}^{−1}(α), G ← ∅.
for s in F do
    E ← ∅.
    for a, b in S do
        ā, b̄ ← a (mod q), b (mod q).
        ē ← b̄ − ās.
        add ē to E.
    end for
    χ²(E) ← Σ_{j=1}^{N} (|{c∈E : c=j}| − M/N)² / (M/N).
    if χ²(E) > δ then
        add s to G.
    end if
end for
if G = ∅ then
    return NOT-RLWE
else if G = {g} then
    return g
else
    return INSUFFICIENT-SAMPLES
end if
```

Proof. It is a standard fact (see [23], for example) that the chi-square value on samples from $D_{\Lambda_R,\sigma} \bmod \mathfrak{q}$ follows the noncentral chi-square distribution with $(N-1)$ degrees of freedom and parameter λ_0 given by

$$\lambda_0 = d_2(D_{\Lambda_R,\sigma} \pmod{\mathfrak{q}}, U(R/\mathfrak{q}))^2 \cdot MN.$$

Note that we have $\lambda_0 \geq (2d(D_{\Lambda_R,\sigma} \bmod \mathfrak{q}, U(R/\mathfrak{q}))/\sqrt{N})^2 MN = 4M\Delta^2 = \lambda$. Recall that our attack succeeds if the "error" set $\mathcal{E}$ from each of the $(N-1)$ wrong guesses of $s \pmod{\mathfrak{q}}$ passes the test, and the true reduced errors fail the test. We assume that the results of these tests are independent of each other. Then the first event happens with probability α^{N-1}, whereas the second event has probability $1 - \mathcal{F}_{N-1;\lambda_0}(\delta)$. Since this is an increasing function in λ_0, we can replace λ_0 by λ, and the theorem follows. □

Remark 16. One could choose the value of α in Theorem 15 to suit the specific instance. The probability of success will change accordingly. When we expect the statistical distance Δ to be large, it is preferable to choose a larger α to increase the probability of success. For example, if we choose $\alpha = 1 - \frac{1}{10N}$, then $\alpha^{N-1} \geq e^{-1/10} = 0.904\cdots$.

Figure 1 shows a plot of p versus Δ for various choices of N, made according to Theorem 15, where we fix the number of samples to be $M = 5N$ and fix $\alpha = 1 - \frac{1}{10N}$.

Remark 17. For linear equations with small errors, there is the attack on the search RLWE problem proposed by Arora and Ge [1]. However, the attack requires solving a linear system in $\approx n^d/d!$ variables. Here d is the number of possible values for the error: e.g., if the error can take values $0, 1, 2, -1, -2$, then $d = 5$. Since it requires $\approx n^d/d!$ samples, the attack of Arora and Ge requires $\geq 10^8$ samples when $n \geq 100$ and $d \geq 5$, for example. 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).

5 Vulnerable Instances among Subfields of Cyclotomic Fields

We searched for instances of RLWE vulnerable to the chi-square attack. For this purpose, we restricted attention to subfields of cyclotomic fields $\mathbb{Q}(\zeta_m)$. Throughout this section, we assume m is a positive integer

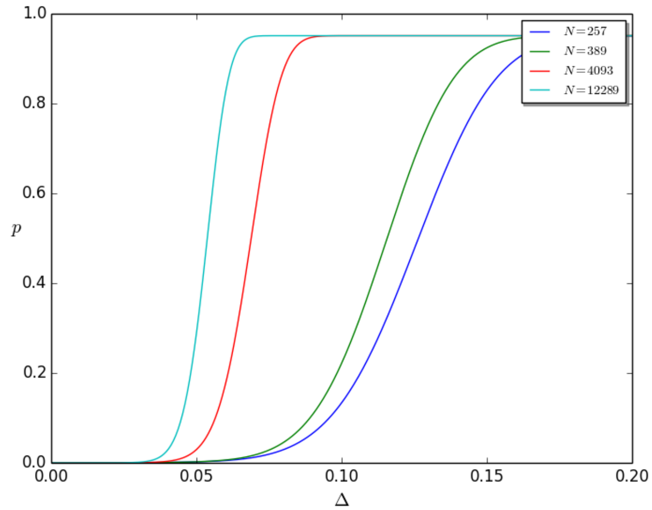

Fig. 1. Success probability versus statistical distance

that is *odd and squarefree*. The Galois group $\mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ is canonically isomorphic to $G = (\mathbb{Z}/m\mathbb{Z})^*$. For a subgroup H of G, let $K_{m,H} = \mathbb{Q}(\zeta_m)^H$ be the subfield of elements fixed by H. Then the extension $K_{m,H}/\mathbb{Q}$ is Galois with degree $n = \frac{\varphi(m)}{|H|}$. Also, the residue degree of a prime q in $K_{m,H}$ is equal to the order of q in the quotient group G/H. Moreover, $K_{m,H}$ has a canonical *normal integral basis*, as follows. For each integer i coprime to m, set $w_i = \sum_{h\in H} \zeta_m^{hi}$. Then $w := (w_i)_{i\in G/H}$ is a $\mathbb{Z}$-basis of R. (For a proof of this fact, see [15, Proposition 6.1]). Thus we have $A_w = TA'_w$, where

$$T = \begin{bmatrix} I_{r_1} & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} I_{r_2} & \frac{1}{\sqrt{2}} I_{r_2} \\ 0 & \frac{-i}{\sqrt{2}} I_{r_2} & \frac{i}{\sqrt{2}} I_{r_2} \end{bmatrix}, \qquad (A'_w)_{ij} = \sum_{h\in H} \zeta_m^{hij}, \text{ for } i, j \in G/H.$$

Remark 18. The field $K_{m,H}$ is totally real if and only if $-1 \in H$, in which case $(r_1, r_2) = (n, 0)$. Otherwise, it is totally complex, and $(r_1, r_2) = (0, n/2)$.

Lemma 19. *Suppose $\mathcal{R}$ is an RLWE instance such that the underlying field K is a Galois number field and that q is unramified in K. Then the reduced error distribution $D_{\Lambda_R,\sigma} \mod \mathfrak{q}$ is independent of the choice of prime ideal $\mathfrak{q}$ above q.*

Proof. From Lemma 8, we may switch from a prime $\mathfrak{q}$ to $\mathfrak{q}'$ via $\mathrm{Gal}(K/\mathbb{Q})$. On the other hand, the Galois group acts on the lattice Λ_R by permuting the coordinates. Hence we have a group homomorphism

$$\phi : \mathrm{Gal}(K/\mathbb{Q}) \to \mathrm{Aut}(\Lambda).$$

Since permutation matrices are orthogonal, the Galois group action on Λ_R given by ϕ is distance-preserving. In particular, it preserves any spherical discrete Gaussian distribution on Λ_R. □

5.1 Searching for Vulnerable Instances

Algorithm 1 allows us to search for vulnerable instances among fields of the form $K_{m,H}$ by generating actual RLWE samples and running the attack. Success of the attack will indicate vulnerability of the instance. Note that our field searching requires sampling efficiently from a discrete Gaussian $D_{\Lambda,\sigma}$, for which we use the efficient algorithm developed in [14].

In Table 3, we list some instances on which the attack has succeeded. The columns of Table 3 are as follows. The first two columns specify m and the generators of H, where H is represented as a subgroup of $(\mathbb{Z}/m\mathbb{Z})^*$; the column labeled f is the residue degree of q. The last column consists of either the runtime for an actual attack which succeeded, or an estimation of the runtime. Note that we omitted our choice of prime ideal $\mathfrak{q}$, since due to Lemma 19 the choice of $\mathfrak{q}$ is irrelevant to our attack. The parameters σ_0 in Table 3 represent the boundary of the power of our attack, i.e., we tried higher σ_0 and the attack failed. Note that although σ_0 is relatively small, in practice it still provides exponentially many error vectors. Intuitively, when $\sigma_0 = 1$, our σ is equal to the geometric mean of the lengths of a Gram-Schmidt basis of Λ_R. In practice, the lengths of these basis vectors do not differ by a lot, so we still expect to get at least $\Omega(2^n)$ error vectors.

Also, in terms of the normal integral basis of R, the coefficients of the error e are of the same size. In particular, none of the coefficients will be zero with overwhelming probability. Thus a standard linear algebra attack does not apply to this case.

The rows of Table 3 with "estimated" runtime mean the following. First, we ran the chi-square test on the correct reduced errors to obtain an estimate $\hat{\Delta}$ of the statistical distance Δ. We then chose α according to $\hat{\Delta}$ and obtained an estimation $\hat{p}$ of the success probability of our attack, using the formula in Theorem 15. The corresponding rows in the table all have $\hat{p} > 1 - 2^{-10}$, suggesting that the attack is very likely to succeed. Finally, we ran a few chi-square tests on samples obtained from a few randomly chosen incorrect guesses to compute the average time t for running one chi-square test. We set the estimated runtime for the attack to be tN.

Table 3. Attacked sub-cyclotomic RLWE instances

m	generators of H	n	q	f	σ_0	no. samples	runtime (in hours)
2805	[1684, 1618]	40	67	2	1	22445	3.49
15015	[12286, 2003, 11936]	60	43	2	1	11094	1.05
15015	[12286, 2003, 11936]	60	617	2	1.25	8000	228.41 (estimated)
90321	[90320, 18514, 43405]	80	67	2	1	26934	4.81
255255	[97943, 162436, 253826, 248711, 44318]	90	2003	2	1.25	15000	1114.44 (estimated)
285285	[181156, 210926, 87361]	96	521	2	1.1	5000	75.41 (estimated)
1468005	[312016, 978671, 956572, 400366]	100	683	2	1.1	5000	276.01 (estimated)
1468005	[198892, 978671, 431521, 1083139]	144	139	2	1	4000	5.72

Remark 20. Castryk et al [7] show that there are even more weaknesses in the weak instances found in [12] so that one can attack the corresponding search RLWE problem using a standard linear algebra attack using only a few samples. The approach from [7] using linear algebra will not work on the examples in this paper. Although each coordinate of the error vector only takes on small integers, it is unlikely that any fixed coordinate of the error vector will equal to zero. Hence one can not hope to extract exact linear equations from the samples.

Nonetheless, in [6] Castryck et al. performed an analysis of the instances in our Table 3, and showed that one can recover a certain number of approximate linear equations from each RLWE sample. One can certainly run the Arora-Ge attack using these approximate equations. However, a careful analysis shows that for instances in our table, our attack is more efficient than the Arora-Ge attack.

For example, we take the first instance from Table 3. In Section 5 of [6], it is shown that out of each RLWE sample, one can recover 20 noisy linear equations in the secret key, with each noise sampled from a Gaussian of mean zero and standard deviation 0.5381. (In [6], the standard deviation was incorrectly claimed to be $0.5381/\sqrt{2\pi}$ due to a misunderstanding: they used $r = 1$ in their analysis but our instance has $r = \sqrt{2\pi}$.)

First, we try $d = 7$. In order to run the Arora-Ge attack, we need in the best case $(\binom{40+7-1}{7}/20) \approx 2^{21}$ RLWE samples, assuming all errors after rounding to integers lies in [-3,3]. If we choose $d = 5$ instead, then in the best case we need $(\binom{40+5-1}{5}/20) \approx 54300$ RLWE samples. However, to achieve this, we need the rounded errors in each equation to lie in [-2,2], which happens with probability $\mathrm{erf}(\frac{2.5}{\sqrt{2}\cdot 0.5381})^{\binom{40+5-1}{5}} \approx 0.025$. Our attack on the other hand requires 22445 samples and succeeds with probability greater than 1/2. Moreover, the computational complexity of Arora-Ge attack is cubic in the number of samples, while our attack is linear in the number of samples. Hence we conclude that our attack is more efficient. A similar analysis can be done for other instances in Table 3.

5.2 Discussion of the Reason for Vulnerability

We searched for vulnerable instances where the modulus has residue degree one or two. It turns out that all vulnerable instances we found and listed in Table 3 have a modulus of degree two. In this section we give a heuristic explanation for the existence of examples of higher degree. Let K be a Galois number field and suppose q is a prime of residue degree f in K. We will give a scenario under which a vulnerability to our attack may appear.

For the purposes of the thought experiment, we will suppose there exists a "good" integral basis $w_1, \cdots, w_n$ of the ring of integers R, by which we mean that the vectors $\iota(w_i)$ and $\iota(w_j)$ are almost orthogonal and short for $i \neq j$; this is only for convenience in the discussion. Fix a prime ideal $\mathfrak{q}$ above q. Then the images of the basis under the reduction modulo $\mathfrak{q}$ map are elements of $F := R/\mathfrak{q}$. Now if for some index i, the element w_i lies inside some proper subfield K' of K, and if q has residue degree $f' < f$ in K', then $w_i \pmod{\mathfrak{q}}$ will lie in a proper subfield of F. *If this occurs for a large number of the basis elements* w_i, then we could expect the distribution $D_{\Lambda_R,\sigma} \mod \mathfrak{q}$ to take values in a proper subfield of F more frequently than the uniform distribution. This would allow us to distinguish it from the uniform distribution on F.

In practice, we found that the above scenario is particularly likely when the field K has a subfield K' of index 2 such that q splits completely in K' and has residue degree 2 in K. Since the ring of integers of K' is a subring of the ring of integers of K, one has at least $n/2$ linearly independent vectors in Λ_R with the desired property, i.e., their reduction modulo some prime $\mathfrak{q}$ above q lie inside $\mathbb{F}_q$ instead of $\mathbb{F}_{q^2}$.

5.3 A Detailed Example

In order to illustrate our discussion above together with the search-to-decision reduction, we present a vulnerable Galois instance in detail, where we generated RLWE samples, performed the attack, and used the search-to-decision reduction to recover the entire secret s.

Example 21. Let $m = 3003$ and H be the subgroup of $(\mathbb{Z}/m\mathbb{Z})^*$ generated by 2276, 2729 and 1123. Then $K = K_{m,H}$ is a Galois number field of degree $n = 30$. We take the modulus to be $q = 131$, a prime of residue degree 2, and take $\sigma_0 = 1$. We generate the secret s from the discrete Gaussian $D_{\Lambda_R,\sigma}$. There are 15 prime ideals in K lying above q, which we denote by $\mathfrak{q}_1, \cdots, \mathfrak{q}_{15}$. We then generate 1000 RLWE samples and use Algorithm 1 and Theorem 9 to recover $s \pmod{\mathfrak{q}_i}$ for each $1 \leq j \leq 15$. Then we use the Chinese remainder theorem to recover s. The attack succeeded in 32.8 hours. The code for this attack is in the appendix.

6 Attacks on the Prime Cyclotomic Fields

6.1 Attacking non-dual RLWE when $q = p$

Let p be an odd prime and let $K = \mathbb{Q}(\zeta_p)$ be the p-th cyclotomic field. Then K has degree $(p-1)$ and discriminant p^{p-2}. The prime p is totally ramified in K, so there is a unique prime ideal $\mathfrak{p} = (1 - \zeta_p)$ above p, and the reduction from R/pR to $R/\mathfrak{p}R \cong \mathbb{F}_p$ takes all powers of ζ_p to 1.

We give a heuristic argument that the attack could work: writing the error e as $\sum e_i \zeta_m^i$, we have $e \pmod{\mathfrak{p}} = \sum_i e_i$. Since the coefficients e_i tend to be small, it may be that $e \pmod{\mathfrak{p}}$ takes on small values with higher probability, making the instance vulnerable to our chi-square attack. Table 4 contains data of some actual attacks we have done. Note that the parameters σ_0 represent the boundary of the power of our attack in practice, i.e., we tried higher σ_0 and the attack failed.

Table 4. Attacked instances of DRLWE for $K = \mathbb{Q}(\zeta_p)$

q $(= p)$	n	σ_0	runtime (in seconds)
251	250	0.5	2.62
503	503	0.575	12.02
809	808	0.61	34.38

6.2 Attacking dual RLWE

We adopt our attack to the decision version of dual RLWE for the field $K = \mathbb{Q}(\zeta_p)$, with no assumptions on the modulus q. Keep the notations as above, and let $R^\vee$ be the dual ideal of R. Let $r > 0$ be the width parameter. Then the error e is sampled from the continuous spherical Gaussian distribution of width r, which is denoted D_r in [18]. Recall that the secret $s \in R^\vee/qR^\vee$, and an RLWE sample is $(a, b = as + e) \in R_q \times K_\mathbb{R}/qR^\vee$.

We start by scaling the second coordinate by p. Then $b' = bp = a(ps) + pe$. Using the fact that $pR^\vee = \mathfrak{p}$, we see that $s' = ps \in \mathfrak{p}/q\mathfrak{p}$, and $e' = pe \in K_\mathbb{R}/q\mathfrak{p}$. Thus we can regard s' as the new secret, and e' as the new error.

Note that the scaled error $e' = pe$ is sampled from the continuous spherical Gaussian D_{pr}. Equivalently, by [10], we may assume e' is sampled as

$$e' = p \cdot e = \sum_{i=0}^{p-1} e_i \zeta_p^i.$$

where the coefficients e_i are i.i.d. one-dimensional Gaussians with width $\sqrt{p}r$.

Recall that our goal is to tell the difference between the above samples and samples chosen uniformly from $R_q \times K_\mathbb{R}/q\mathfrak{p}$. Let $\beta = \zeta_p - 1$. Note that $K = \mathbb{Q}(\beta)$, hence every element in $K_\mathbb{R}$ can be uniquely written as $\sum_{i=0}^{p-2} a_i\beta^i$ $(a_i \in \mathbb{R})$. Consider the map

$$\rho : K_\mathbb{R} \to \mathbb{R} : \sum a_i\beta^i \mapsto a_0. \qquad (*)$$

It is clear that ρ is additive. We examine the image of e' under the map ρ. We write

$$e' = \sum_{i=0}^{p-2} \epsilon_i\beta^i, (\epsilon_i \in \mathbb{R}).$$

Then one verfies that $\epsilon_0 = e_0 + \cdots + e_{p-2} - (p-1)e_{p-1}$ and we have $\rho(e) = \epsilon_0$.

Now we make two observations: first, since the ideal $\mathfrak{p}$ is generated by β, we have $\rho(\mathfrak{p}) = p\mathbb{Z}$; second, we have $as' = a(p \cdot s) \in \mathfrak{p}/q\mathfrak{p}$. Combining these observations, we see that

$$\rho(b') = \rho(as) + \rho(e') \equiv \rho(e') \equiv \epsilon_0 \mod p\mathbb{Z}.$$

We could describe our attack on the decision RLWE as follows: for each scaled sample (a, b'), we compute $\rho(b')$. Then we perform a statistical test on the set $\{\rho(b') \mod p\mathbb{Z}\} \subseteq \mathbb{R}/p\mathbb{Z}$ to distinguish it from the uniform distribution on the circle $\mathbb{R}/p\mathbb{Z}$.

Note that this attack did not involve the modulus q, thus it can be applied to *any* modulus. This is in contrast to the previous attack on the non-dual case, where the attack was only performed under the assumption that $q = p$ is the unique ramified prime.

Remark 22. The search-to-decision reduction for dual RLWE in cyclotomic fields and completely split modulus is proved in [18]. However, the theorem requires that the error width $r \geq \eta_\epsilon(R^\vee)$ for some negligible $\epsilon = \epsilon(n)$. (Here $\eta_\epsilon(R^\vee)$ is the smoothing parameter defined in [19]. For $R = \mathbb{Z}[\zeta_p]$, if we take $\epsilon = 2^{-p+1}$, then one sees that $\eta_\epsilon(R^\vee) \leq 1$, and $\eta_\epsilon(R^\vee)$ tends to 1 in the limit as $p \to \infty$). Hence the search-to-decision reduction of [18] essentially requires $r \geq 1$, which is above the parameters we can attack ($r \sim 1/\sqrt{p}$). So in this particular case, our attack on the decision problem cannot be transferred to an attack on the search problem using this search-to-decision reduction.

Table 5 records some successful attacks. Note that we have omitted the modulus q since it is irrelevant to the attack. We used 50∼400 bins for the chi-square tests. We observe from the table that the error width we can attack is about a constant times $1/\sqrt{p}$, and that the constant is growing (if slowly) as p grows.

Table 5. Attacking the dual RLWE in $\mathbb{Q}(\zeta_p)$

p	$r\sqrt{p}$	no. samples	average run time	success rate
307	0.82	1535	0.048 second	6 out of 10
507	0.83	2515	0.076 second	8 out of 10
809	0.85	4045	0.134 second	6 out of 10
997	0.86	4985	0.154 second	5 out of 10
1103	0.87	5515	0.192 second	5 out of 10
1201	0.88	6005	0.202 second	2 out of 10

7 Can Modulus Switching be Used?

The modulus switching procedure is a technique to reduce noise in RLWE samples, and has been discussed extensively in [3] and [16]. We recap the basic ideas of modulus switching. Let $\mathcal{R} = (K, q, \sigma, s)$ be an RLWE instance. Choose another prime p less than q as the new modulus and consider the instance $\mathcal{R}' = (K, p, \sigma', s)$ for some $\sigma' > \sigma$. We can "switch modulus" if there exists a map

$$\pi_{q,p} : R_q \to R_p,$$

which takes RLWE samples with respect to $\mathcal{R}$ to RLWE samples with respect to $\mathcal{R}'$. In what follows, we give a heuristic argument that our attack will not work in combination with modulus switching under a naïve implementation, and isolate the key characteristics a successful implementation of the attack would require.

One example of a map $\pi_{q,p}$ being used in practice is as follows. Let $\alpha = \frac{p}{q}$ and fix a small positive number τ. For an equivalence class $[a]$ in R_q, we sample a vector a' from the "shifted discrete Gaussian" $D_{\Lambda_R,\tau,\alpha a}$, defined as follows. For a lattice Λ and a vector $c \in \mathbb{R}^n$,

$$D_{\Lambda,\tau,c}(x) = \frac{\rho_\tau(x-c)}{\sum_{y\in\Lambda} \rho_\tau(y-c)}, \ \forall x \in \Lambda.$$

Finally, we set $\pi_{q,p}([a]) = a' \pmod{pR}$. Note that the definition of $\pi_{q,p}([a])$ is independent of the choice of representative a, as follows. Suppose we choose another representative a_1, then $a_1 = a + \lambda q$ for some $\lambda \in R$, hence $\alpha a_1 = \alpha a + \lambda p$. Finally, observe that the shifted discrete Gaussian behaves well under translating by a lattice point, i.e., we have $D_{\Lambda,\tau,c+u} = D_{\Lambda,\tau,c} + u$ for any $u \in \Lambda$.

Put loosely, the map $\pi_{q,p}$ scales a by p/q and then rounds back into the lattice. It is a natural question then to ask whether modulus switching can be combined with our attack, to switch from a "strong" modulus to a "weak" modulus. However, a heuristic argument shows that the naive combination of our attack with modulus switching will not work.

Let $a'' = \alpha a - a'$. By construction, we expect a'' to be a short vector in $\mathbb{R}^n$, and the point a' can be viewed as a "rounding" of the point αa to the lattice Λ_R.

We will make two heuristic assumptions:

1. That $\pi_{q,p}$ takes the uniform distribution on R_q to an almost uniform distribution on R_p.
2. The distribution of b'' and $(sa)''$ is independent modulo $\mathfrak{q}$, for $s \neq \pm 1$.

Proposition 23. *Under the assumption that $\pi_{q,p}$ takes the uniform distribution on R_q to an almost uniform distribution on R_p, the reduction of a'' modulo $\mathfrak{p}$ will be almost uniformly distributed in $R/\mathfrak{p}R$.*

Proof. The reduction map $R \to R/\mathfrak{p}$ is a ring homomorphism that can be extended to a homomorphism of additive groups $\phi : \frac{1}{q}R \to R/\mathfrak{p}$ by the following chain of maps:

$$\frac{1}{q}R \xrightarrow{(\text{mod } \mathfrak{p}\frac{1}{q}R)} \frac{1}{q}R \Big/ \mathfrak{p}\frac{1}{q}R \xrightarrow{\times q} R/\mathfrak{p}R \xrightarrow{\times [q]^{-1}} R/\mathfrak{p}R.$$

Then the relation $a'' + a' = \alpha a$ is preserved by this map. However, $\phi(\alpha a) = 0 \pmod{\mathfrak{p}}$, so that $\phi(a'') \equiv -\phi(a')$. ☐

Suppose we have a sample (a, b) and the switched sample $(a', b') = (\pi_{q,p}(a), \pi_{q,p}(b))$. Consider the error $e' := b' - a's$. Suppose $b = as + e + \lambda q$ for some $\lambda \in R$. Then

$$\begin{aligned} e' &= b' - a's \\ &= \alpha(b - as) - b'' + a''s. \\ &= \alpha e + \lambda p - b'' + a''s. \end{aligned}$$

and therefore, considering this as an additive relation in $\frac{1}{q}R$ and applying the map of the proof above,

$$e' \equiv -b'' + a''s \pmod{\mathfrak{p}}.$$

By the Proposition above, a'' and b'' are uniformly distributed modulo $\mathfrak{p}$. Hence, if we assume the a'' and b'' are independent, then the reduced rounding errors $a'' \pmod{\mathfrak{p}}$ and $b'' \pmod{\mathfrak{p}}$ are also independent, and the new reduced errors $e' \pmod{\mathfrak{p}}$ would follow the uniform distribution. So our chi-square attack will fail on these modulus-switched samples, even though p might be a "weak" modulus.

Therefore, the best hope of attack is if one of our two assumptions is violated by a map $\pi_{q,p}$. The second is the most likely target. Note that a'' and b'' are the rounding errors when we try to round αa and αb to the lattice Λ_R. However, Λ_R is a n-dimensional lattice, so there are $\Omega(2^n)$ options of rounding a vector in $\mathbb{R}^n$ to a moderately close lattice point. Even in the scenario with zero error, i.e., $e = 0$, an attacker will face the task of finding a "nice" rounding algorithm, so that the roundings of the two vectors αa and $\alpha b = \alpha a s$ are somehow related.

So far, we are not aware of any such algorithm, unless the secret s is trivial, e.g., $s = 1$, in which case αa is almost equal to αb, and one expects that a'' is close to b''.

References

1. Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: Automata, Languages and Programming, pp. 403–415. Springer (2011)
2. Bos, J.W., Lauter, K., Loftus, J., Naehrig, M.: Improved security for a ring-based fully homomorphic encryption scheme. In: Cryptography and Coding, pp. 45–64. Springer (2013)
3. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. pp. 309–325. ACM (2012)
4. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from Ring-LWE and security for key dependent messages. In: Advances in Cryptology–CRYPTO 2011, pp. 505–524. Springer (2011)
5. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. SIAM Journal on Computing 43(2), 831–871 (2014)
6. Castryck, W., Iliashenko, I., Vercauteren, F.: On error distributions in ring-based LWE. LMS Journal of Computation and Mathematics 19(A), 130–145 (2016)
7. Castryck, W., Iliashenko, I., Vercauteren, F.: Provably weak instances of Ring-LWE revisited. In: 35th Annual international conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt 2016). vol. 9665, pp. 147–167. Springer (2016)
8. Chen, H., Lauter, K.E., Stange, K.E.: Attacks on search RLWE. Cryptology ePrint Archive, Report 2015/971, 2015.
9. Chen, H., Lauter, K.E., Stange, K.E.: Vulnerable Galois RLWE families and improved attacks. In: Selected Areas in Cryptography – SAC 2016 (2016)
10. Ducas, L., Durmus, A.: Ring-LWE in polynomial rings. In: Public Key Cryptography–PKC 2012, pp. 34–51. Springer (2012)
11. Eisenträger, K., Hallgren, S., Lauter, K.E.: Weak instances of PLWE. In: Selected Areas in Cryptography–SAC 2014, pp. 183–194. Springer (2014)
12. Elias, Y., Lauter, K., Ozman, E., Stange, K.: Provably weak instances of Ring-LWE. In: Advances in Cryptology – CRYPTO 2015, Lecture Notes in Comput. Sci., vol. 9215, pp. 63–92. Springer, Heidelberg (2015)
13. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: Advances in Cryptology–EUROCRYPT 2012, pp. 465–482. Springer (2012)
14. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the 40th annual ACM symposium on Theory of computing. pp. 197–206. ACM (2008)
15. Johnston, H.: Notes on Galois modules. Notes accompanying the course Galois Modules, Cambridge (2011)

16. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. Designs, Codes and Cryptography 75(3), 565–599 (2014)
17. López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: Proceedings of the 44th annual ACM Symposium on Theory Of Computing. pp. 1219–1234. ACM (2012)
18. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. Journal of the ACM (JACM) 60(6), 43:1–43:35 (2013)
19. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. SIAM Journal on Computing 37(1), 267–302 (2007)
20. Micciancio, D., Regev, O.: Lattice-based cryptography. In: Post-quantum cryptography, pp. 147–191. Springer (2009)
21. Nguyen, P.: Cryptanalysis of the Goldreich-Goldwasser-Halevi cryptosystem from Crypto 97. In: Advances in Cryptology – CRYPTO 1999. pp. 288–304. Springer (1999)
22. Peikert, C.: How (not) to instantiate Ring-LWE. In: SCN. pp. 411–430 (2016)
23. Ryabko, B.Y., Stognienko, V., Shokin, Y.I.: A new test for randomness and its application to some cryptographic problems. Journal of statistical planning and inference 123(2), 365–376 (2004)
24. Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: Advances in Cryptology–EUROCRYPT 2011, pp. 27–47. Springer (2011)
25. Stein, W., et al.: Sage Mathematics Software (Version 6.4). The Sage Development Team (2014), `http://www.sagemath.org`