Weak Instances of PLWE

Kirsten Eisenträger[1⋆], Sean Hallgren[2⋆⋆], and Kristin Lauter[3]

[1] Department of Mathematics, The Pennsylvania State University, University Park, PA 16802, USA, and Harvard University. eisentra@math.psu.edu
[2] Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA 16802, USA. hallgren@cse.psu.edu
[3] Microsoft Research, One Microsoft Way, Redmond, WA 98052. klauter@microsoft.com

Abstract. In this paper we present a new attack on the polynomial version of the Ring-LWE assumption, for certain carefully chosen number fields. This variant of RLWE, introduced in [BV11] and called the PLWE assumption, is known to be as hard as the RLWE assumption for 2-power cyclotomic number fields, and for cyclotomic number fields in general with a small cost in terms of error growth. For general number fields, we articulate the relevant properties and prove security reductions for number fields with those properties. We then present an attack on PLWE for number fields satisfying certain properties.

1 Introduction

Lattice-based cryptography has been an active area of study for at least two decades. The Ajtai-Dwork [AD99] public-key cryptosystem was based on the worst-case hardness of a variant of the Shortest Vector Problem (SVP). The NTRU family of cryptosystems [HPS98] were defined in particularly efficient lattices connected to number theory and were standardized in the IEEE P1363.1 Lattice-Based Public Key Cryptography standard [IEEE]. Recently, a new assumption has been introduced, Learning-With-Errors (LWE) [Reg09] and the Ring-Learning-With-Errors (RLWE) variant [LPR10], which is related via various security reductions from hard lattice problems such as (Gap-)SVP and Bounded Distance Decoding (BDD) [LPR10,Reg09,BLP+13]. NTRUEncrypt and NTRUSign can be slightly modified so that their security also based on the hardness of a variant of the RLWE problem [SS11], and applications to Homomorphic Encryption were proposed in [BV11] and extended in [BGV11] and [GHS12].

⋆ Partially supported by National Science Foundation grant DMS-1056703 and by the National Security Agency (NSA) under Army Research Office (ARO) contract number W911NF-12-1-0522. Part of this work was done while the first author was visiting Microsoft, Harvard University and MIT.

⋆⋆ Partially supported by National Science Foundation awards CCF-0747274 and CCF-1218721, and by the National Security Agency (NSA) under Army Research Office (ARO) contract number W911NF-12-1-0522. Part of this work was done while visiting Microsoft and MIT.

The advantage of RLWE over LWE is better efficiency and functionality for cryptosystems based on this hardness assumption, but with extra structure in the ring variant of the assumption comes the possibility of special attacks which take advantage of this structure. So far, we have not seen special attacks which take advantage of the extra structure.

The PLWE decisional hardness assumption was proposed in [BV11] as the basis for a fully homomorphic encryption scheme, and was introduced as a variant of the RLWE assumption. But the worst-case to average-case reduction from the shortest vector problem on ideal lattices to the PLWE problem was only proved for the special case of 2-power cyclotomic fields and the proof of the reduction was cited from [LPR10]. A clear explanation of why this reduction works in the 2-power cyclotomic case was given in [DD12], which identified the necessary properties of the 2-power cyclotomic ring, and extended the proof to work for general cyclotomic fields, with minimal loss in the growth of the error bounds.

On the other hand, we can ask about the hardness of the PLWE assumption for general number rings, and its relationship to the hardness of the RLWE assumption. The key point is the distortion in the error distribution which occurs when passing between the Gaussian error distribution in the continuous complex space, where the ring is embedded via the Minkowski (or "canonical") embedding, and the error distribution when sampling error vectors coefficient-wise.

In this paper, we investigate the extent to which the hardness of these problems holds in more general number rings, that is, when the number field is not necessarily a cyclotomic field generated by roots of unity. We present an attack on the PLWE problem in certain carefully constructed examples of number fields. We also give a sequence of reductions between the Search and Decision versions of RLWE and the PLWE assumptions, under various conditions on the number field. An intuitive way to explain our results is that, for number fields satisfying our conditions, our attack on PLWE works by "guessing" one of q possibilities for the value of the secret polynomial evaluated at 1, and distinguishing PLWE samples with non-negligible probability when the error vectors are sampled from Gaussian distributions coefficient-wise in the polynomial ring.

Practical encryption schemes based on PLWE all work based on the assumption that the error distribution is sampled coefficient-wise in the polynomial ring. For 2-power cyclotomic fields, this is equivalent to sampling from the usual Gaussian error distribution for the ring embedded in a real vector space, but our attack does not work for these fields. On the other hand, our attack shows that it is not safe to work directly with the PLWE assumption in arbitrary number fields. So for the purpose of constructing secure and efficient cryptosystems, it is a reasonable conclusion that one should stick to cyclotomic number fields, until the class of fields for which there exists a reduction to RLWE is enlarged.

More specifically, for a degree n number field $K = \mathbb{Q}[x]/(f(x))$ and an integer modulus q, if $f(1)$ is congruent to zero modulo q, then our attack runs in time $\widetilde{O}(q)$. For all current recommendations on parameter selection for RLWE, our attack runs much faster than the known distinguishing attacks based on solutions to the shortest vector problem, or decoding attacks based on computing a

reduced basis for a lattice, which run in time exponential in n. For example, recommended high security parameters for LWE and RLWE-based cryptosystems given in [LP11, Figure 4] specify $n = 320$ and $q = 4093$. While the distinguishing and decoding attacks, estimated to run in time 2^{122} and 2^{119} seconds, are impractical, an attack which runs in time $\widetilde{O}(q)$ is certainly feasible.

We emphasize that this does not constitute a practical attack on existing PLWE/RLWE-based cryptosystems. First of all, all practical systems known to us are based on the RLWE problem in a cyclotomic ring, normally a 2-power cyclotomic ring, $R = \mathbb{Z}[x]/(\Phi_m(x))$ where m is a power of 2. Our attack does not apply to $f = \Phi_m$ because $f(1) = \Phi_m(1)$ cannot be zero modulo q when q is much larger than m. Secondly, our attack runs in time proportional to q. While this is an improvement over algorithms which need to find a short vector or compute a reduced basis and run in time $O(2^n)$, it is still far from a practical attack when q is taken to be of size 2^{128}, (with $n = 2^{12}$), which is the minimum size required for homomorphic computations in [GLN12] and [BLN14], for example.

2 Background

2.1 Distances and distributions

For the definition of the RLWE and PLWE hardness assumptions and for the implementation of related cryptosystems, it is necessary to define certain distributions, which will be used in particular for error distributions.

Adopting the notation from [LPR10], we will let $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ denote the space

$$H = \{(x_1, \ldots, x_n) \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} \mid x_{s_1+s_2+j} = \overline{x_{s_1+j}}, j = 1, \ldots, s_2\},$$

where s_1 and s_2 are non-negative integers and $n = s_1 + 2s_2$. Since the last s_2 complex coordinates depend on the previous s_2 coordinates, as they are just the complex conjugates of them, H is isomorphic to $\mathbb{R}^n$. It inherits the usual inner product $\langle (x_i), (y_i) \rangle := \sum_{i=1}^n x_i \cdot y_i$.

There are several natural notions of distance on an inner product space, and we will primarily need the ℓ_2-norm, given by $\|x\|_2 := (\sum_{i=1}^n x_i^2)^{1/2} = \sqrt{\langle x, x \rangle}$ for $x \in H$, and the ℓ_∞-norm, given by $\|x\|_\infty := \max |x_i|$.

For a real number $\sigma > 0$, the Gaussian function $\rho_\sigma : H \to (0, 1]$ is given by

$$\rho_\sigma(x) := \exp(-\pi \langle x, x \rangle / \sigma^2).$$

The *continuous Gaussian probability distribution* D_σ is given by $D_\sigma(x) = \frac{\rho_\sigma(x)}{\sigma^n}$.

As in [LPR10, Definition 5] we now define the family of LWE error distributions to which the results apply.

Definition 1. *For a positive real $\alpha > 0$, the family $\Psi_{\leq \alpha}$ is the set of all elliptical Gaussian distributions $D_{\mathfrak{r}}$ (over $K \otimes_{\mathbb{Q}} \mathbb{R}$) where each parameter $r_i \leq \alpha$.*

2.2 Lattices

A lattice is a discrete subgroup of a continuous space. For example, in $\mathbb{R}^n$, the real vector space of dimension n, a lattice can be specified by a set of n linearly independent vectors, and the lattice is the integral span of those vectors. An orthogonal basis for a lattice, if one exists, is a basis such that the basis vectors are pairwise orthogonal, with respect to the given inner product.

For a lattice $\Lambda \subset H$, define the dual lattice as

$$\Lambda^\vee = \{y \in H \mid \forall x \in \Lambda, \langle x, y \rangle = \sum_{i=1}^{n} x_i \overline{y_i} \in \mathbb{Z}\}.$$

We will also need to refer to the *smoothing parameter* of a lattice, $\eta_\epsilon(\Lambda)$ introduced by Micciancio and Regev [MR07], which for a lattice Λ and a positive real number ϵ is defined to be the smallest s such that $\rho_{1/s}(\Lambda^\vee \setminus \{0\}) \leq \epsilon$.

2.3 Number fields

A *number field* is a finite algebraic extension of the field of rational numbers $\mathbb{Q}$. It is a field which contains $\mathbb{Q}$ and is a finite dimensional vector space over $\mathbb{Q}$. The *degree* of the number field is its dimension as a vector space. A number field K is *Galois* if $K/\mathbb{Q}$ is a Galois extension, which means it is both separable and normal. An extension is *separable* if every element in the extension is separable, which means that its minimal polynomial has distinct roots. An extension $K/\mathbb{Q}$ is *normal* if every irreducible polynomial with rational coefficients which has one root in K has all of its roots in K. In particular, this means that every isomorphism of the field K into its algebraic closure which fixes $\mathbb{Q}$ actually maps into K, and thus is an automorphism of K. For a Galois extension $K/\mathbb{Q}$, the set of automorphisms of K which fix $\mathbb{Q}$ forms a group, and is called the Galois group, $\mathrm{Gal}(K/\mathbb{Q})$, of $K/\mathbb{Q}$.

The *ring of integers* in a number field K is the set of all algebraic integers in the number field, which means the elements which satisfy a *monic* irreducible polynomial with integer coefficients. This set is a ring, called a number ring, and is usually denoted by $\mathcal{O}_K$. If the ring $R = \mathcal{O}_K$ is generated over $\mathbb{Z}$ by (sums and powers of multiples of) a single element, $R = \mathbb{Z}[\beta]$, then we say that R is *monogenic*.

The m^{th} cyclotomic field is the number field generated by the m^{th} roots of unity. Let ζ_m be a primitive m^{th} root of 1, i.e. $\zeta_m^m = 1$ but no smaller power is 1. Then $K = \mathbb{Q}(\zeta_m) = \mathbb{Q}[x]/(\Phi_m(x))$, where $\Phi_m(x)$ is the m-th cyclotomic polynomial $\Phi_m(x) = \prod_{k \in (\mathbb{Z}/m\mathbb{Z})^\times} (x - \zeta_m^k)$ with degree equal to $n = \varphi(m)$. When m is an odd prime we have $\Phi_m(x) = 1 + x + x^2 + \cdots + x^{m-1}$, and when m is a power of 2, $\Phi_m(x) = x^{m/2} + 1$.

For a finite algebraic extension $K/\mathbb{Q}$, the Trace and Norm maps from K to $\mathbb{Q}$ are defined as the sum (*resp.* the product) of all the algebraic conjugates of an element. The Trace map induces a non-degenerate bilinear form $\mathrm{Tr}(xy)$ on

K. The *dual* or the *codifferent* of the ring of integers, R, with respect to this bilinear form is the collection of elements

$$R^\vee = \{y \in K \mid \mathrm{Tr}(xy) \in \mathbb{Z}, \forall x \in R\}.$$

This is often denoted $\mathcal{D}_K^{-1}$ in algebraic number theory.

It is known that if $R = \mathcal{O}_K = \mathbb{Z}[\beta] = \mathbb{Z}[x]/(f(x))$ is monogenic, then the codifferent is generated by the single element $(1/f'(\beta))$ [Ser79, p. 56, Cor 2]. In that case, there is a simple isomorphism between $R^\vee$ and R which scales elements by multiplication by $f'(\beta)$. For $K = \mathbb{Q}(\zeta_m)$, we have $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$.

The ring of integers R is embedded in H via the Minkowski embedding (called the *canonical embedding* in [LPR10]) which sends any $x \in K$ to $(\sigma_1(x), \ldots \sigma_n(x))$, where σ_i are the real and complex embeddings of K, ordered to coincide with the definition of H. Under this embedding, the notions of duality and codifferent coincide. In particular, we can define an *ideal lattice* to be the image under this embedding of any fractional ideal of R by taking the lattice generated by the image of the n basis elements of the $\mathbb{Z}$-basis for the ideal.

An *ideal* I in a commutative ring R is an additive subgroup which is closed under multiplication by elements of R. A *prime ideal* $I \neq R$ is an ideal with the property that if the product of two elements $a, b \in R$ is such that $ab \in I$, then either a or b is in I. Ideals in the ring of integers of number fields have unique factorization into products of prime ideals. We say that a prime ideal $(p) = p\mathbb{Z}$ *splits completely* in an extension of number fields $K/\mathbb{Q}$ if the ideal $p\mathcal{O}_K$ factors into the product of n distinct ideals of degree 1, where $n = [K : \mathbb{Q}]$ is the degree of the extension $K/\mathbb{Q}$.

2.4 Definition of the Ring-LWE distribution and problem

The Ring-LWE distribution and hardness assumptions were introduced in [LPR10, Section 3] using the notation $K_\mathbb{R} = K \otimes \mathbb{R}$ and $\mathbb{T} = K_\mathbb{R}/R^\vee$. For an integer q, let R_q denote R/qR.

Definition 2. *(Ring-LWE Distribution) For $s \in R_q^\vee$ a secret, and an error distribution ψ over $K_\mathbb{R}$, the Ring-LWE distribution $A_{s,\psi}$ over $R_q \times \mathbb{T}$ consists of samples generated as follows: choose a uniformly at random from R_q and choose the error vector e from the error distribution ψ, then the samples are pairs of the form $(a, (a \cdot s)/q + e)$.*

Definition 3. *(Ring-LWE Search Problem) Let Ψ be a family of distributions over $K_\mathbb{R}$. The Ring-LWE Search problem (*R-LWE$_{q,\Psi}$*), for some $s \in R_q^\vee$ and $\psi \in \Psi$, is to find s, given arbitrarily many independent samples from $A_{s,\psi}$.*

Definition 4. *(Ring-LWE Average-Case Decision Problem) Let Υ be a family of error distributions over $K_\mathbb{R}$. The Ring-LWE Average-Case Decision problem (*R-DLWE$_{q,\Upsilon}$*) is to distinguish with non-negligible advantage between arbitrarily many independent samples from $A_{s,\psi}$, for a random choice of $s \in R_q^\vee$ and $\psi \in \Upsilon$, and the same number of samples chosen independently and uniformly at random from $R_q \times \mathbb{T}$.*

2.5 Worst-case hardness of search version of ring-LWE

Theorem 1. *([LPR10]) Let K be an arbitrary number field of degree n, with $R = \mathcal{O}_K$, $\alpha = \sigma/q \in (0,1)$, and $q \geq 2 \in N$ such that $\alpha \cdot q \geq \omega(\sqrt{\log n})$. Then there is a probabilistic polynomial-time quantum reduction from the $\tilde{O}(\sqrt{n}/\alpha)$-approximate SIVP problem on ideal lattices in K to* R-LWE$_{q,\Psi_{\leq\alpha}}$*. For K a cyclotomic number field, this gives a reduction from the $\tilde{O}(\sqrt{n}/\alpha)$-approximate SVP problem.*

2.6 Known attacks

When selecting secure parameters for cryptographic applications of the hardness of RLWE, the following known attacks are currently taken into account. The *distinguishing attack* considered in [MR09,RS10] for LWE requires the adversary to find a short vector in the scaled dual of the LWE lattice. The distinguishing advantage is then given in terms of the length of the vector found. According to [LP11], the vector should be of length less than $q/(2\sigma)$. Writing q in terms of n, this amounts to solving a short-vector problem in an n-dimensional lattice, and if q is too large with respect to n, this problem will be easy. This gives some insight as to why q cannot be too large with respect to n.

Concrete security estimates given in [LP11, Figure 4] against this attack lead to suggested parameters, for example at the "high security" level, of $n = 320$, $q \approx 2^{12}$, and $\sigma = 8$ (however recall that for 2-power cyclotomic fields, n should be a power of 2). For those parameter choices, the distinguishing attack is estimated to run in time 2^{122} (seconds) to obtain a distinguishing advantage of 2^{-64}.

The *decoding attack* presented in [LP11] is an attack which actually recovers the secret error vector in the ciphertext. To run the attack requires a reduced basis, and the estimated time to compute the reduced basis when $n = 320$ and $q \approx 2^{12}$ is 2^{119} seconds for decoding probability 2^{-64}.

3 Overview of results

We work with the ring of integers $R = \mathcal{O}_K$ in a number field K of degree n and a prime number q and consider the following properties:

1. (q) splits completely in K, and $q \nmid [R : \mathbb{Z}[\beta]]$;
2. K is Galois over $\mathbb{Q}$;
3. the ring of integers of K is generated over $\mathbb{Z}$ by β, $\mathcal{O}_K = \mathbb{Z}[\beta] = \mathbb{Z}[x]/(f(x))$ with $f'(\beta) \mod q$ "small" ;
4. the transformation between the Minkowski embedding of K and the power basis representation of K is given by a scaled orthogonal matrix.
5. let $f \in \mathbb{Z}[x]$ be the minimal polynomial for β. Then $f(1) \equiv 0 \pmod{q}$;
6. q can be chosen suitably large.

Note that by the Chebotarev Density Theorem, there are infinitely many choices for q satisfying this and only finitely many exclusions.

In Section 4 we will show that for pairs (K, q) satisfying conditions (1) and (2) we have a search-to-decision reduction from R-LWE$_q$ to R-DLWE$_q$.

In Section 5 we will consider a second reduction, from R-DLWE$_q$ to PLWE, which is essentially a slightly more general version of the reduction given in [LPR10] and [DD12]. For that step we require that K satisfies conditions (3) and (4).

In Section 6 we will then give an attack which breaks instances of the PLWE decision problem whenever (K, q) satisfy conditions (5) and (6), and we will consider possible extensions of our attack.

4 Search to decision reduction for the ring-LWE problem

In this section we will prove the following theorem.

Theorem 2. *Let K be a number field such that $K/\mathbb{Q}$ is Galois of degree n and let $R = \mathcal{O}_K$ be its ring of integers. Let $R^\vee$ be the dual (the codifferent ideal) of R. Let β be an algebraic integer such that $K = \mathbb{Q}(\beta)$, and let $f(x) \in \mathbb{Z}[x]$ be the minimal polynomial of β over $\mathbb{Q}$. Let q be a prime such that (q) splits completely in K and such that $q \nmid [R : \mathbb{Z}[\beta]]$. Let α be such that $\alpha \cdot q \geq \eta_\varepsilon(R^\vee)$ for some negligible $\varepsilon = \varepsilon(n)$. Then there is a randomized polynomial-time reduction from* R-LWE$_{q,\Psi_{\leq\alpha}}$ *to* R-DLWE$_{q,\Upsilon_\alpha}$.

Proof. Let n denote the degree of K over $\mathbb{Q}$. Since $K/\mathbb{Q}$ is Galois, f factors completely in K as $f(x) = (x - \beta) \cdot (x - \beta_2) \cdots (x - \beta_m)$. Let q be a prime as in the theorem statement, which factors as $(q) = \mathfrak{q}_1 \dots \mathfrak{q}_n$.

The field K has n embeddings, and since $K/\mathbb{Q}$ is Galois either all of these embeddings are real or they are all complex.

Let $\beta_1 := \beta$, and for $i = 1, \dots, n$, let $\sigma_i : K \hookrightarrow \mathbb{C}$ $(i = 1, \dots, n)$ be the embedding which sends β_1 to β_i. Let $\sigma : K \hookrightarrow \mathbb{C} \times \cdots \times \mathbb{C}$ be the Minkowski embedding sending $x \in K$ to $(\sigma_1(x) = x, \sigma_2(x), \dots, \sigma_n(x))$.

Before we can finish the proof we need two lemmas:

Lemma 1. *Let $K = \mathbb{Q}(\beta_1)$ be as above. For any $\alpha > 0$, the family $\Psi_{\leq\alpha}$ is closed under every automorphism τ of K, i.e. $\psi \in \Psi_{\leq\alpha}$ implies that $\tau(\psi) \in \Psi_{\leq\alpha}$.*

Proof. Let τ be an automorphism of K. Then $\tau(\beta_1) = \beta_j$ for some $1 \leq j \leq n$. Let $x \in K$. Then $x = \sum_{i=0}^{n-1} k_i \beta_1^i$ for $k_i \in \mathbb{Q}$ and

$$\sigma(x) = \left(\sum_{i=0}^{n-1} k_i \beta_1^i, \sum_{i=0}^{n-1} k_i \beta_2^i, \dots, \sum_{i=0}^{n-1} k_i \beta_n^i \right).$$

On the other hand, $\sigma(\tau(\beta_1))$ is a vector whose entries are simply a permutation of $\beta_1, \dots, \beta_n$ and whose first entry is β_j, and so for any $x \in K$, the coordinates of $\sigma(x)$ and $\sigma(\tau(x))$ are simply a rearrangement of each other.

Hence for any $\psi = D_\mathfrak{r} \in \Psi_{\leq\alpha}$, we have $\tau(D_\mathfrak{r}) = D_{\mathfrak{r}'} \in \Psi_{\leq\alpha}$, where the entries of $\mathfrak{r}'$ are simply a rearrangement of the entries of $\mathfrak{r}$ and hence are all at most α. □

Worst-case search to worst-case decision

Definition 5. *The* $\mathfrak{q}_i$-LWE$_{q,\Psi}$ *problem is: given access to* $A_{s,\psi}$ *for some arbitrary* $s \in R_q^\vee$ *and* $\psi \in \Psi$*, find* $s \bmod \mathfrak{q}_i R^\vee$.

Lemma 2. *(LWE to* $\mathfrak{q}_i$*-LWE) Suppose that the family* Ψ *is closed under all automorphisms of* K*. Then for every* $1 \leq i \leq n$ *there is a deterministic polynomial-time reduction from* LWE$_{q,\psi}$ *to* $\mathfrak{q}_i$-LWE$_{q,\psi}$.

Proof. The proof proceeds almost word for word as the proof in [LPR10]. Given two prime ideals $\mathfrak{q}_i$ and $\mathfrak{q}_j$ above q, [LPR10] uses the explicit automorphism τ_k with $\tau_k(\zeta) = \zeta^k$ where $k = j/i \in \mathbb{Z}_m^*$ that maps $\mathfrak{q}_j$ to $\mathfrak{q}_i$. Instead we use the fact that the Galois group of K over $\mathbb{Q}$ acts transitively on the prime ideals above $\mathfrak{q}$. Hence in our situation, given $\mathfrak{q}_i, \mathfrak{q}_j$ there is also an automorphism τ of K such that $\tau(\mathfrak{q}_i) = \mathfrak{q}_j$. The rest of the argument is identical to the argument in [LPR10]. □

Conclusion of the proof of Theorem 2: To finish the proof based on these Lemmas, we argue as in [LPR13a, Lemma 5.9] and the proof given there goes through for Galois fields exactly as stated. □

5 Reduction from R-DLWE$_q$ to PLWE

This section essentially summarizes and slightly generalizes one of the main results from [DD12]. The reduction for general cyclotomic fields is also covered in [LPR13b].

5.1 The PLWE problem

The PLWE problem was first defined in [LPR10] and [BV11].

Definition 6. *(The PLWE assumption). For all* $\kappa \in \mathbb{N}$*, let* $f(x) = f_\kappa(x)$ *be a polynomial of degree* $n = n(\kappa)$*, and let* $q = q(\kappa)$ *be a prime integer. Let* $R = \mathbb{Z}[x]/(f)$*, let* $R_q = R/qR$ *and let* χ *denote a distribution over* R.

The PLWE assumption PLWE$_{f,q,\chi}$ *states that for any* $\ell = poly(\kappa)$ *it holds that*

$$\{(a_i, a_i \cdot s + e_i)\}_{i \in [\ell]} \text{ is computationally indistinguishable from } \{a_i, u_i\}_{i \in [\ell]},$$

where s *is sampled from the noise distribution* χ*, the* a_i *are uniform in* R_q*, the error polynomials* e_i *are sampled from the error distribution* χ *and the ring elements* u_i *are uniformly random over* R_q.

The PLWE assumption is a decisional assumption.

5.2 Reduction

In [DD12, p. 39], the authors explain the reduction in the 2-power cyclotomic case in terms of the two key properties of the ring $R = \mathcal{O}_K$ which are used:

1. When $R = \mathbb{Z}[\zeta_m]$ with m a power of 2, then $nR^\vee = R$, for $n = m/2$.
2. The transformation between the embedding of R in the continuous real vector space H and the representation of R as a $\mathbb{Z}$-vector space with the power basis consisting of powers of ζ_m is an orthogonal linear map.

Their argument shows that one can slightly generalize those conditions to our Properties (3) and (4) and obtain the reduction for general number fields with those properties. Note that the claim is that these conditions are sufficient to obtain the reduction, not that they are necessary. There may be a reduction which works for an even more general class of number fields.

Step 1 of the reduction uses the property that $R = \mathbb{Z}[\beta] = \mathbb{Z}[x]/(f(x))$ is monogenic to transform the ring-LWE samples between distributions on $R^\vee$ and R, at the cost of a scaling by $f'(\beta)$, where $f'(\beta)$ is "small" modulo q.

When reducing RLWE to PLWE we take samples from the Minkowski embedding and consider them in the coefficient embedding. The main point is whether vectors that are short in the Minkowski embedding have small coefficients in the coefficient embedding.

Step 2 uses the fact that the matrix which transforms between the embedding of R in H and the power basis representation of R is a scaled orthogonal matrix, so it transforms the spherical Gaussian distribution in H into a spherical Gaussian distribution in the power basis representation. Thus the error distribuition can be sampled directly from small values coefficient-wise in the polynomial ring.

Note that a different reduction is given in [DD12] for general cyclotomic fields because of the fact that ζ_m potentially does not satisfy the requirement that $\Phi'_m(\zeta_m)$ is small modulo q compared to n. As noted there, according to a result of Erdős [Erd46], the coefficients of Φ_m can be superpolynomial in size. In any case, even for m prime, Φ'_m has coefficients of size up to roughly m.

6 Breaking certain instances of PLWE

6.1 The Attack

Let K be a number field such that $f(1) \equiv 0 \pmod{q}$, and such that q can be chosen large enough. Let $R := \mathcal{O}_K$, and let $R_q := R/qR$.

Now, given samples, $(a_i, b_i) \in R_q \times R_q$, we have to decide whether the samples are uniform or come from a PLWE distribution. To do this we take the representatives of a_i and b_i in R, call them a_i and b_i again, and evaluate them at 1. This gives us elements $a_i(1), b_i(1) \in \mathbb{F}_q$. If (a_i, b_i) are PLWE samples, then by definition,

$$b_i = a_i \cdot s + e_i,$$

and so

$$b_i(1) \equiv (a_i \cdot s)(1) + e_i(1) \pmod{q}.$$

Since $f(1) \equiv 0 \pmod{q}$, the Chinese Remainder Theorem gives us that

$$b_i(1) \equiv a_i(1) \cdot s(1) + e_i(1) \pmod{q}.$$

Now we can guess $s(1)$, and we have q choices. For each of our guesses we compute $b_i(1) - a_i(1) \cdot s(1)$. If (a_i, b_i) are PLWE samples and our guess for $s(1)$ is correct, then $b_i(1) - a_i(1) \cdot s(1) = e_i(1)$, and we will detect that it is non-uniform, because e_i is taken from χ. (For example, if e_i is taken from a Gaussian with small radius, then $e_i(1)$ will be "small" for all i and hence not uniform.) If (a_i, b_i) are uniform samples, then $b_i(1) - a_i(1) \cdot s(1)$ for any fixed choice of $s(1)$ will still be uniform, since $a_i(1)$, $b_i(1)$ are both uniform modulo q.

6.2 A family of examples

Let $f(x) = X^n + (k-1)pX + p$, where p is a prime less than n, and k is chosen such that $1 + kp = q$ with q prime and $q > n$. This polynomial is Eisenstein at p and hence irreducible. By Dirichlet's theorem about primes in arithmetic progressions, there are infinitely many values of k that give a prime q.

Also, by construction

$$f(1) = 1 + (k-1)p + p = 1 + kp = 1 \equiv 0 \pmod{q}.$$

Moreover, $f'(1)$ is not zero modulo q since

$$f'(1) = n + (k-1)p = (1 + kp) + (n - 1 - p) = q + a,$$

with a a number which, by construction is $< n$. Hence f has 1 as a simple root modulo q, and by the Chinese Remainder Theorem

$$\mathbb{Z}[X]/(f(X)) \cong Z[X]/(X-1) \times Z[X]/(h(X))$$

with $h(X)$ coprime to $(X-1)$. As explained in the previous section, this allows us to guess $s(1)$, since $(a_i \cdot s)(1) = a_i(1) \cdot s(1)$. Hence for this choice of polynomials and choice of q, we can distinguish uniform samples from PLWE samples and break PLWE.

6.3 Extension of the attack on PLWE

The attack we presented in Section 6.1 above on PLWE for number fields satisfying property (5) can be extended to a more general class of number fields as follows:

Suppose that $f(x)$ has a root β modulo q which has small order in $(\mathbb{Z}/q\mathbb{Z})^*$. If q is a prime, then this is equivalent to β having small order modulo $q-1$. If $f(\beta) \equiv 0 \mod q$, then the same attack above will work by evaluating samples at β, instead of at 1. Now unfortunately, the value of the error polynomials $e_i(\beta)$ are harder to distinguish from random ones than in the case $\beta = 1$: although the $e_i(x)$ have small coefficients modulo q, the powers of β may grow large and also may wrap around modulo q. However, if β has small order in $(\mathbb{Z}/q\mathbb{Z})^*$, then the set $\{\beta^i\}_{i=0,\ldots,n-1}$ takes on only a small number values, and this can be used to distinguish samples arising from $e_i(\beta)$ from random ones with non-negligible advantage.

6.4 Security implications for RLWE and PLWE-based cryptosystems

Putting all the results of this paper together, if there exist number fields satisfying all 6 properties, then for those number fields we would also have an attack on RLWE. A toy example of a field satisfying the first five conditions listed above is $K = \mathbb{Q}(\sqrt{11})$, $\beta = \sqrt{11}$, $f(x) = x^2 - 11$, and $q = 5$.

In general, for a given degree n, it is not hard to generate irreducible polynomials $f(x)$ of degree n, such that, letting $q = f(1)$, q is sufficiently large. Each such polynomial $f(x)$ gives rise to a weak instance of PLWE, according to our attack. However, to obtain an attack on RLWE, we would need to check that the first 4 properties are also satisfied. The first two properties are easy to check, but not necessarily easy to assure by construction. Properties (3) and (4) are not as easy to check, and harder to assure by construction.

The security of RLWE in general and its reduction to hard lattice problems is an interesting theoretical question and thus the construction of a number field satisfying all 6 properties would be a significant result. But from the point of view of practical applications to cryptography and homomorphic encryption, the security of the proposed cryptosystems is based on the hardness of the PLWE assumption. Thus the attack presented here and the results of this section are of interest in themselves.

References

[IEEE] P1363.1: Standard Specifications for Public-Key Cryptographic Techniques Based on Hard Problems over Lattices, December 2008. `http://grouper.ieee.org/groups/1363/`

[AD99] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *STOC '97: Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 284–293. ACM, New York, 1999.

[BLN14] Joppe W. Bos, Kristin Lauter, Michael Naehrig. Private Predictive Analysis on Encrypted Medical Data, Journal of Biomedical Informatics (2014) DOI 10.1016/j.jbi.2014.04.003

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC'13: Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 575–584. ACM, New York, 2013.

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

[BGV11] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science—ITCS 2012*, pages 309–325. ACM 2012.

[DD12] Léo Ducas and Alain Durmus. Ring-LWE in polynomial rings. In *Public key cryptography—PKC 2012*, volume 7293 of *Lecture Notes in Comput. Sci.*, pages 34–51. Springer, Heidelberg, 2012.

[Erd46] Paul Erdős. On the coefficients of the cyclotomic polynomial. *Bull. Amer. Math. Soc.*, 52:179–184, 1946.

[GHS12] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully homomorphic encryption with polylog overhead. In *Advances in cryptology—EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Comput. Sci.*, pages 465–482. Springer, Heidelberg, 2012.

[GLN12] Thore Graepel, Kristin Lauter, and Michael Naehrig. ML Confidential: Machine Learning on Encrypted Data. In *International Conference on Information Security and Cryptology—ICISC 2012*, volume 7839 of *Lecture Notes in Comput. Sci.*, pages 1–21. Springer, Heidelberg, 2013.

[HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: a ring-based public key cryptosystem. In *Algorithmic number theory (Portland, OR, 1998)*, volume 1423 of *Lecture Notes in Comput. Sci.*, pages 267–288. Springer, Berlin, 1998.

[LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In *Topics in cryptology—CT-RSA 2011*, volume 6558 of *Lecture Notes in Comput. Sci.*, pages 319–339. Springer, Heidelberg, 2011.

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

[LPR13a] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):Art. 43, 35, 2013.

[LPR13b] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In *Advances in cryptology—EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Comput. Sci.*, pages 35–54. Springer, Heidelberg, 2013.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302 (electronic), 2007.

[MR09] Daniele Micciancio and Oded Regev. Lattice-based cryptography. In *Post-quantum cryptography*, pages 147–191. Springer, Berlin, 2009.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009.

[RS10] M. Rückert and M. Schneider. Selecting secure parameters for lattice-based cryptography. Cryptology ePrint Archive, Report 2010/137, 2010.

[Ser79] Jean-Pierre Serre. *Local fields*, volume 67 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1979.

[SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Advances in cryptology—EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Comput. Sci.*, pages 27–47. Springer, Heidelberg, 2011.