Provably Weak Instances of Ring-LWE Revisited

Wouter Castryck[1,2], Ilia Iliashenko[1], and Frederik Vercauteren[1]

[1]KU Leuven ESAT/COSIC and iMinds
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
firstname.lastname@esat.kuleuven.be

[2]Vakgroep Wiskunde, Universiteit Gent
Krijgslaan 281/S22, B-9000 Gent, Belgium

Abstract. In CRYPTO 2015, Elias, Lauter, Ozman and Stange described an attack on the non-dual *decision* version of the ring learning with errors problem (RLWE) for two special families of defining polynomials, whose construction *depends on the modulus q that is being used*. For particularly chosen error parameters, they managed to solve non-dual decision RLWE given 20 samples, with a success rate ranging from 10% to 80%. In this paper we show how to solve the *search* version for the same families and error parameters, using only 7 samples with a success rate of 100%. Moreover our attack works for *every modulus q'* instead of the q that was used to construct the defining polynomial. The attack is based on the observation that the RLWE error distribution for these families of polynomials is very skewed in the directions of the polynomial basis. For the parameters chosen by Elias et al. the smallest errors are negligible and simple linear algebra suffices to recover the secret. But enlarging the error paremeters makes the largest errors wrap around, thereby turning the RLWE problem unsuitable for cryptographic applications. These observations also apply to dual RLWE, but do not contradict the seminal work by Lyubashevsky, Peikert and Regev.

1 Introduction

Hard problems on lattices have become popular building blocks for cryptographic primitives mainly because of two reasons: firstly, lattice based cryptography appears to remain secure even in the presence of quantum computers, and secondly, the security of the primitives can be based on worst-case hardness assumptions. Although it seems appealing to use classical hard lattice problems such as the shortest vector problem or closest vector problem for cryptographic applications, the learning with errors problem (LWE) has proven much more versatile. This problem was introduced by Regev [12, 13] who showed that an efficient algorithm for LWE results in efficient quantum algorithms for approximate lattice problems. The *decision* version of LWE can be defined informally as the problem of distinguishing noisy linear equations from truly random ones. More precisely, let $n \geq 1$ be an integer dimension and $q \geq 2$ an integer modulus, then the problem is to distinguish polynomially many pairs of the form $(\mathbf{a}_i, b_i \approx \langle \mathbf{a}_i, \mathbf{s} \rangle)$ from uniformly random and independent pairs. The vectors $\mathbf{a}_i$ are chosen uniformly

random in $\mathbb{Z}_q^n$, the vector $\mathbf{s}$ is secret and the same for all pairs, and the element b_i is computed as $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i$ where e_i is a random error term drawn from an error distribution on $\mathbb{Z}_q$, such as a discretized Gaussian. The *search* version of LWE asks to recover the secret vector $\mathbf{s}$. The hardness of the LWE problem has been analyzed in $[12, 13, 11, 8, 3]$.

The main downside of LWE is that it is not very practical, basically due to the fact that each new $\mathbf{a}_i$ only gives rise to one element b_i (and not a vector of n elements as one could hope). The result is that the public key size and the computation time of LWE-based cryptosystems are typically quadratic in the security parameter. Lyubashevsky, Peikert and Regev [9] solved this issue by introducing the Ring-LWE (RLWE) problem and showing its hardness under wost-case assumptions on ideal lattices. Its flavour is distantly similar to that of NTRU [7]. Informally, the secret key space $\mathbb{Z}_q^n$ is replaced by $R_q = R/qR$ where R is the ring of integers in a number field $K = \mathbb{Q}[x]/(f)$ with f a monic irreducible integral polynomial of degree n and $q \geq 2$ an integer modulus. The inner product on $\mathbb{Z}_q^n$ is replaced by the ring product in R_q. In its *non-dual* form the *decision* version of RLWE is then roughly defined as follows: distinguish polynomially many samples of the form $(\mathbf{a}_i, \mathbf{b}_i \approx \mathbf{a}_i \cdot \mathbf{s})$ from uniformly random and independent pairs. Here the $\mathbf{a}_i \in R_q$ are uniformly random and independent, $\mathbf{s} \in R_q$ is a fixed random secret, and $\mathbf{b}_i$ is computed as $\mathbf{b}_i = \mathbf{a}_i \cdot \mathbf{s} + \mathbf{e}_i$ where $\mathbf{e}_i \in R_q$ is a short random error term that is drawn from a specific error distribution ψ on R_q. The *search* version of the problem is to recover the secret $\mathbf{s}$ from the list of samples. We stress that the actual problem described and analyzed in [9] is the *dual* RLWE problem, in which the secret and the error term are taken from the reduction modulo q of a certain fractional ideal of K, denoted by $R_q^\vee$; see Section 2 for more details.

As explained in [9], the search and decision problems are equivalent when K is Galois and q is a prime number that splits into prime ideals with small norm (polynomial in n). In general, no such reduction is known and it is easy to see that search RLWE is at least as hard as decision RLWE.

The definition of the error distribution ψ on R_q (or on $R_q^\vee$) plays a crucial role in RLWE and is obtained by pulling back a near-spherical Gaussian distribution under the canonical embedding of the number field. An alternative problem [5] is called Polynomial-LWE (PLWE) and uses an error distribution on R_q where each coordinate of the error term with respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$ is drawn independently from a fixed one-dimensional Gaussian distribution. Again we refer to Section 2 for more details.

In [5], Eisentraeger, Hallgren and Lauter presented families of defining polynomials $f \in \mathbb{Z}[x]$ and moduli q such that the *decision* version of PLWE is weak. The attack can be described in a nutshell as follows: assume that $f(1) \equiv 0 \bmod q$, then evaluation at 1 defines a ring homomorphism ϕ from R_q to $\mathbb{Z}_q$. Applying ϕ to the PLWE samples results in equations of the form $\mathbf{a}_i(1) \cdot \mathbf{s}(1) + \mathbf{e}_i(1) = \mathbf{b}_i(1)$. Therefore, if the images $\mathbf{e}_i(1)$ of the error terms can be distinguished from uniform, one can simply loop through all possibilities for $\mathbf{s}(1) \in \mathbb{Z}_q$ and determine if the corresponding $\mathbf{e}_i(1)$ are uniform on $\mathbb{Z}_q$ or not. So as long as q is small

enough (such that one can exhaustively run through $\mathbb{Z}_q$), $f(1) \equiv 0 \bmod q$, and the images $\mathbf{e}_i(1)$ do not wrap around too much modulo q, this attack breaks decision PLWE.

In [6], Elias, Lauter, Ozman and Stange extended this attack to the *decision* version of non-dual RLWE, rather than PLWE, by showing that for defining polynomials of the form

$$f_{n,a,b} = x^n + ax + b \in \mathbb{Z}[x]$$

where n, a, b are specifically chosen parameters such that i.a. $f_{n,a,b}(1) \equiv 0 \bmod q$, the distortion introduced by pulling back the Gaussian error terms through the canonical embedding is small enough such that the attack on PLWE still applies. This attack was executed for three parameter sets n, a, b, r where given 20 samples, non-dual decision RLWE could be solved with success rates ranging from 10% to 80% depending on the particular family considered [6, Section 9]. Here the parameter r determines the width of the Gaussian that is being pulled back, which Elias et al. chose to be spherical.

Our contributions in this paper are as follows. Firstly, we explain how to solve the *search* version of non-dual RLWE, which one might expect to be a harder problem than the decision version (due to the fact that the corresponding number fields are not Galois), for the same parameter sets, using only **7** samples with a success rate of 100%. The attack invokes simple linear algebra to recover the secret element $\mathbf{s}$ and does not use that $f_{n,a,b}(1) \equiv 0 \bmod q$: in fact, for the same defining polynomial and the same error parameter r our attack works for *every* modulus q'. Secondly, we show that if one tries to adjust r in order to obtain a hard instance of non-dual RLWE, the first few components of the noise wrap around modulo q and become indistinguishable from uniform, thereby obstructing certain cryptographic applications. Thirdly, we show that our observations also apply to the *dual* RLWE problem when set up for the same number fields: either the errors wrap around or linear algebra can be used to reveal the secret. The latter situation only occurs for error widths that are way too small for the hardness results of Lyubashekvsy, Peikert and Regev [9] to be applicable. Therefore neither the results from [6] nor our present attack seem to form a threat on RLWE, at least when set up along the guidelines in [9, 10].

Our observations are easiest to explain for $a = 0$, a case which covers two of the three parameter sets. From $f_{n,a,b}(1) = b + 1 \equiv 0 \bmod q$ and $f_{n,a,b}(1) \neq 0$ (by irreducibility) it follows that the roots of $f_{n,a,b}$ lie on a circle with radius

$$\rho \geq \sqrt[n]{q - 1} > 1.$$

With respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$, the canonical embedding matrix is essentially the Vandermonde matrix generated by these roots, whose column norms grow geometrically as

$$\sqrt{n}, \sqrt{n}\rho, \ldots, \sqrt{n}\rho^{n-1}.$$

This simple observation has major implications for the distortion introduced by the inverse of the canonical embedding: the distribution of the error terms will

be extremely stretched at the terms of low degree, whereas they will be squashed at the terms of high degree. For the parameter sets attacked by Elias et al., the latter are so small that after rounding they simply become zero, thereby resulting in exact linear equations in the coefficients of the secret element **s**. Given enough samples (in the cases considered, between 4 and 7 samples suffice), the secret **s** can be recovered using elementary linear algebra. Furthermore, since the ratio between the maximal and the minimal distortion is roughly $\rho^n \geq q - 1$, it is impossible to increase the width of the Gaussians used without causing the errors at the terms of low degree to wrap around modulo q.

The remainder of the paper is organized as follows: in Section 2 we recall the definition of PLWE and of dual and non-dual RLWE, with particular focus on the error distributions involved. Section 3 reviews the attacks on decision PLWE by Eisenträger, Hallgren and Lauter and non-dual decision RLWE by Elias, Lauter, Ozman and Stange. Section 4 describes our attack on non-dual search RLWE by analyzing the singular value decomposition of the canonical embedding. We also report on an implementation of our attack in Magma [2], which shows that we can indeed easily break the families considered in [6] using less samples, with a higher success probability, and for every choice of modulus q' (instead of just the q that was used to define $f_{n,a,b}$). We also discuss how switching to dual RLWE affects these observations. In Section 5 we study the effect of increasing the error parameter as an attempt to counter our attack, and compare with the hardness results from [9]. Section 6 concludes the paper.

2 Preliminaries

In this section we briefly recall the necessary background on number fields, the canonical embedding and Gaussian distributions to give proper definitions of PLWE and dual and non-dual RLWE.

2.1 Number fields and the canonical embedding

Let $f \in \mathbb{Z}[x]$ be a monic irreducible polynomial of degree n and consider the number field $K = \mathbb{Q}[x]/(f)$ it defines. Let $R \subset K$ denote the ring of integers of K, i.e. the set of all algebraic integers that are contained in K. If f can be taken such that $R = \mathbb{Z}[x]/(f)$, then K is called a *monogenic* number field and f a monogenic polynomial.

The field K has exactly n embeddings into $\mathbb{C}$ denoted by $\sigma_i : K \to \mathbb{C}$ for $i = 1, \ldots, n$. These n embeddings correspond precisely to evaluation in each of the n distinct roots α_i of f, i.e. an element $a(x) \in K$ is mapped to $\sigma_i(a(x)) = a(\alpha_i) \in \mathbb{C}$. Assume that f has s_1 real roots and $n - s_1 = 2s_2$ complex conjugate roots and order the roots such that $\overline{\alpha_{s_1+k}} = \alpha_{s_1+s_2+k}$ for $k = 1, \ldots, s_2$. The *canonical embedding* (also known as the Minkowski embedding) $\sigma : K \to \mathbb{C}^n$ is then defined as:

$$\sigma(a) = (\sigma_1(a), \ldots, \sigma_{s_1}(a), \sigma_{s_1+1}(a), \ldots, \sigma_{s_1+s_2}(a), \overline{\sigma_{s_1+1}}(a), \ldots, \overline{\sigma_{s_1+s_2}}(a)) \,.$$

4

It is easy to see that the canonical embedding maps into the space $H \subset \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ given by

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

The space H is isomorphic to $\mathbb{R}^n$ as an inner product space by considering the orthonormal basis for H given by the columns of

$$B = \begin{pmatrix} I_{s_1 \times s_1} & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} I_{s_2 \times s_2} & \frac{\mathbf{i}}{\sqrt{2}} I_{s_2 \times s_2} \\ 0 & \frac{1}{\sqrt{2}} I_{s_2 \times s_2} & -\frac{\mathbf{i}}{\sqrt{2}} I_{s_2 \times s_2} \end{pmatrix}.$$

With respect to this basis, the coordinates of $\sigma(a)$ are given by a real vector

$$(\tilde{a}_1, \ldots, \tilde{a}_n) := (\sigma_1(a), \ldots, \sigma_{s_1}(a), \sqrt{2} \, \mathfrak{Re}(\sigma_{s_1+1}(a)), \ldots, \sqrt{2} \, \mathfrak{Re}(\sigma_{s_1+s_2}(a)),$$
$$\sqrt{2} \, \mathfrak{Im}(\sigma_{s_1+1}(a)), \ldots, \sqrt{2} \, \mathfrak{Im}(\sigma_{s_1+s_2}(a))).$$

Note that in [6] the authors did not include the factor $\sqrt{2}$, but we choose to keep it since it makes B unitary.

In summary, an element $a(x) \in K$ can be represented in the polynomial basis as $(a_0, \ldots, a_{n-1})$ where $a(x) = \sum_{i=0}^{n-1} a_i x^i$ but also by a real vector $(\tilde{a}_1, \ldots, \tilde{a}_n)$ where the canonical embedding of a is given by:

$$\sigma(a) = B \cdot (\tilde{a}_1, \ldots, \tilde{a}_n)^t.$$

Let M_f denote the Vandermonde matrix $(\alpha_i^{j-1})_{i,j}$ for $i, j = 1, \ldots, n$, then the polynomial basis representation is related to the (real) canonical embedding representation by the following transformation

$$(a_0, \ldots, a_{n-1})^t = M_f^{-1} \cdot B \cdot (\tilde{a}_1, \ldots, \tilde{a}_n)^t.$$

Since M_f^{-1} will play a crucial role in the following, we denote it with N_f. Later on, to ease notation we will just write M_f instead of $M_{f_{n,a,b}}$, and similarly for N_f.

2.2 Ideals of the ring of integers and their dual

An *integral ideal* $I \subseteq R$ is an additive subgroup of R closed under multiplication by elements of R, i.e. $rI \subset I$ for any $r \in R$. A *fractional ideal* $I \subset K$ is a set such that $dI \subseteq R$ is an integral ideal for some $d \in R$. A *principal* (fractional or integral) ideal I is one that is generated by some $u \in K$, i.e. $I = uR$; we denote it as $I = \langle u \rangle$. The sum $I + J$ of two (fractional or integral) ideals is the set of all $x + y$ with $x \in I, y \in J$ and the product $I \cdot J$ is the smallest (fractional or integral) ideal containing all products $x \cdot y$ with $x \in I, y \in J$. The set of non-zero fractional ideals forms a group under multiplication; this is not true for integral ideals. The inverse of a non-zero fractional ideal is denoted by I^{-1}. Every fractional ideal I is a free $\mathbb{Z}$-module of rank n, and therefore $I \otimes \mathbb{Q} = K$.

Its image $\sigma(I)$ under the canonical embedding is a lattice of rank n inside the space H.

The trace $\mathrm{Tr} = \mathrm{Tr}_{K/\mathbb{Q}} : K \rightarrow \mathbb{Q}$ maps an element x to the sum of its embeddings $\mathrm{Tr}\,(x) = \sum_{i=1}^{n} \sigma_i(x)$ and defines an additive homomorphism from R to $\mathbb{Z}$. The norm $\mathrm{No} = \mathrm{No}_{K/\mathbb{Q}} : K \rightarrow \mathbb{Q}$ takes the product of all embeddings $\mathrm{No}(x) = \prod_{i=1}^{n} \sigma_i(x)$ and is multiplicative.

For a fractional ideal I, its dual $I^{\vee}$ is defined as

$$I^{\vee} = \{x \in K : \mathrm{Tr}(xI) \subseteq \mathbb{Z}\}.$$

It is easy to see that $(I^{\vee})^{\vee} = I$ and that $I^{\vee}$ is also a fractional ideal. (Under the canonical embedding, this corresponds to the usual notion of dual lattice, modulo complex conjugation.) Furthermore, for any fractional ideal I, its dual is $I^{\vee} = I^{-1}R^{\vee}$. The factor $R^{\vee}$ is a fractional ideal called the *codifferent* and its inverse is called the *different ideal* which is integral. For a monogenic defining polynomial f, i.e. $R = \mathbb{Z}[x]/(f)$ we have that $R^{\vee} = \langle 1/f'(\alpha) \rangle$ where α is a root of f. Applying this fact to the cyclotomic number field of degree $n = 2^k$ with defining polynomial $f(x) = x^n + 1$, we get that $f'(\xi_{2n}) = n\xi_{2n}^{n-1}$ with ξ_{2n} a primitive $2n$-th root of unity. Thus $R^{\vee} = \langle n^{-1} \rangle$, since ξ_{2n}^{n-1} is a unit.

2.3 Gaussian distributions and discretization

Denote by Γ_r the normal Gaussian distribution on $\mathbb{R}$ with mean 0 and parameter r given by $\Gamma_r(x) = r^{-1}\exp(-\pi x^2/r^2)$. Note that we have $r = \sqrt{2\pi}\rho$ with ρ the standard deviation. We can define an elliptical Gaussian distribution $\Gamma_{\mathbf{r}}$ on H as follows: let $\mathbf{r} = (r_1, \ldots, r_n) \in (\mathbb{R}^+)^n$ be a vector of n positive real numbers, then a sample of $\Gamma_{\mathbf{r}}$ is given by $B \cdot (x_1, \ldots, x_n)^t$ where each x_i is sampled independently from Γ_{r_i} on $\mathbb{R}$. Note that via the inverse of the canonical embedding this also defines a distribution $\Psi_{\mathbf{r}}$ on $K \otimes \mathbb{R}$, in other words

$$N_f \cdot B \cdot (x_1, \ldots, x_n)^t$$

gives us the coordinates of $\Gamma_{\mathbf{r}} \leftarrow (x_1, \ldots, x_n)$ with respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$.

In practice we sample from the continuous distribution $\Gamma_{\mathbf{r}}$ modulo some finite but sufficiently high precision (e.g. using the Box-Muller method). In particular our samples live over $\mathbb{Q}$ rather than $\mathbb{R}$, so that an element sampled from $\Psi_{\mathbf{r}}$ can be truly seen as an element of the field K. For use in RLWE one even wants to draw elements from I for some fixed fractional ideal $I \subset K$, where $I = R$ (non-dual RLWE) and $I = R^{\vee}$ (dual RLWE) are the main examples. In this case one should discretize the Gaussian distribution $\Gamma_{\mathbf{r}}$ to the lattice $\sigma(I)$. There are several ways of doing this, e.g. by rounding coordinates with respect to some given $\mathbb{Z}$-module basis; see [10, 9] and the references therein. But for our conclusions this discretization is not relevant, and because it would needlessly complicate things we will just omit it.

2.4 The Polynomial-LWE and Ring-LWE problem

In this section we provide formal definitions of PLWE [5] and RLWE [9, 4], both in its dual and its non-dual version [6]. We stress that it is the dual version of RLWE that was introduced in [9] and for which certain hardness results are available, one of which is recalled in Theorem 1 below.

Let $f \in \mathbb{Z}[x]$ be a monic irreducible polynomial of degree n and let $q \geq 2$ be an integer modulus. Consider the quotient ring $P = \mathbb{Z}[x]/(f)$ and denote with P_q the residue ring P/qP. Denote with Γ_r^n the spherical Gaussian on $\mathbb{R}^n$ with parameter r and interpret this as a distribution on $P \otimes \mathbb{R}$ by mapping the standard basis of $\mathbb{R}^n$ to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$ of P. In particular, elements $\mathbf{e}(x) = \sum_{i=0}^{n-1} e_i x^i \leftarrow \Gamma_r^n$ have each coefficient e_i drawn independently from Γ_r. Let $\mathfrak{U}(P_q)$ denote the uniform distribution on P_q and let $\mathfrak{U}(P_{q,\mathbb{R}})$ be the uniform distribution on the torus $P_{q,\mathbb{R}} = (P \otimes \mathbb{R})/qP$.

With these ingredients we can define the decision and search PLWE problems.

Definition 1 (PLWE distribution). *For $\mathbf{s}(x) \in P_q$ and $r \in \mathbb{R}^+$, a sample from the PLWE distribution $A_{\mathbf{s}(x),r}$ over $P_q \times P_{q,\mathbb{R}}$ is generated by choosing $\mathbf{a}(x) \leftarrow \mathfrak{U}(P_q)$, choosing $\mathbf{e}(x) \leftarrow \Gamma_r^n$ and outputting $(\mathbf{a}(x), \mathbf{b}(x) = \mathbf{a}(x) \cdot \mathbf{s}(x) + \mathbf{e}(x) \bmod qP)$.*

Definition 2 (Decision PLWE). *The decision PLWE problem is to distinguish, for a random but fixed choice of $\mathbf{s}(x) \leftarrow \mathfrak{U}(P_q)$, with non-negligible advantage between arbitrarily many independent samples from $A_{\mathbf{s}(x),r}$ and the same number of independent samples from $\mathfrak{U}(P_q) \times \mathfrak{U}(P_{q,\mathbb{R}})$.*

Definition 3 (Search PLWE). *For a random but fixed choice of $\mathbf{s}(x) \leftarrow \mathfrak{U}(P_q)$, the search PLWE problem is to recover $\mathbf{s}(x)$ with non-negligible probability from arbitrarily many independent samples from $A_{\mathbf{s}(x),r}$.*

To define the dual and non-dual RLWE problems we require a degree n number field K with ring of integers R. We also fix a fractional ideal $I \subset K$, for which two choices are available: in the *dual* RLWE problems we let $I = R^\vee$, while in the *non-dual* RLWE problems we take $I = R$. Note that $I \otimes \mathbb{R} = K \otimes \mathbb{R}$, so we can view the distribution $\Psi_{\mathbf{r}}$ from the previous section as a distribution on $I \otimes \mathbb{R}$. We let I_q denote I/qI and write $I_{q,\mathbb{R}}$ for the torus $(I \otimes \mathbb{R})/qI$. As before we let $\mathfrak{U}(I_q)$ denote the uniform distribution on I_q and let $\mathfrak{U}(I_{q,\mathbb{R}})$ be the uniform distribution on $I_{q,\mathbb{R}}$.

Definition 4 (RLWE distribution). *For $\mathbf{s}(x) \in I_q$ and $\mathbf{r} \in (\mathbb{R}^+)^n$, a sample from the RLWE distribution $A_{\mathbf{s}(x),\mathbf{r}}$ over $R_q \times I_{q,\mathbb{R}}$ is generated by choosing $\mathbf{a}(x) \leftarrow \mathfrak{U}(R_q)$, choosing $\mathbf{e}(x) \leftarrow \Psi_{\mathbf{r}}$ and returning $(\mathbf{a}(x), \mathbf{b}(x) = \mathbf{a}(x) \cdot \mathbf{s}(x) + \mathbf{e}(x) \bmod qI)$.*

Definition 5 (Decision RLWE). *The decision RLWE problem is to distinguish, for a random but fixed choice of $\mathbf{s}(x) \leftarrow \mathfrak{U}(I_q)$, with non-negligible advantage between arbitrarily many independent samples from $A_{\mathbf{s}(x),\mathbf{r}}$ and the same number of independent samples from $\mathfrak{U}(R_q) \times \mathfrak{U}(I_{q,\mathbb{R}})$.*

Definition 6 (Search RLWE). *For a random but fixed choice of* $\mathbf{s}(x) \leftarrow \mathfrak{U}(I_q)$, *the search* RLWE *problem is to recover* $\mathbf{s}(x)$ *with non-negligible probability from arbitrarily many independent samples from* $A_{\mathbf{s}(x),\mathbf{r}}$.

A hardness statement on the search RLWE problem in its dual form (i.e. with $I = R^\vee$) was provided by Lyubashevsky, Peikert and Regev. For proof-technical reasons their result actually deals with a slight variant called the search RLWE$_{\leq r}$ problem, where $r \in \mathbb{R}^+$. In this variant each sample is taken from $A_{\mathbf{s}(x),\mathbf{r}}$ for a new choice of $\mathbf{r}$, chosen uniformly at random from $\{(r_1, \ldots, r_n) \in (\mathbb{R}^+)^n \,|\, r_i \leq r \text{ for all } i\}$. Think of this parameter r and the modulus $q \geq 2$ as quantities that vary with n, and let ω be a superlinear function. Then Lyubashevsky et al. proved:

Theorem 1 ([9, Theorem 4.1]). *If* $r \geq 2\omega(\sqrt{\log n})$ *then for some negligible* ε *(depending on n) there is a probabilistic polynomial-time quantum reduction from* KDGS$_\gamma$ *to* RLWE$_{\leq r}$, *where*

$$\gamma : I \mapsto \max\left\{ \eta_\varepsilon(I) \cdot (\sqrt{2}q/r) \cdot \omega(\sqrt{\log n}), \sqrt{2n}/\lambda_1(I^\vee) \right\}.$$

Here $\eta_\varepsilon(I)$ *is the smoothing parameter of* $\sigma(I)$ *with threshold* ε, *and* $\lambda_1(I^\vee)$ *is the length of a shortest vector of* $\sigma(I^\vee)$.

In the above statement KDGS$_\gamma$ refers to the *discrete Gaussian sampling problem*, which is about producing samples from a spherical Gaussian in H with parameter r', discretized to the lattice $\sigma(I)$, for any given non-zero ideal $I \subset R$ and any $r' \geq \gamma(I)$. As discussed in [9] there are easy reductions from certain standard lattice problems to the discrete Gaussian sampling problem.

As an intermediate step in their proof Lyubashevsky et al. obtain a classical (i.e. non-quantum) reduction from an instance of the bounded distance decoding problem in ideal lattices to RLWE$_{\leq r}$; see [9, Lemma 4.5].

In contrast, Elias, Lauter, Ozman and Stange [6] study RLWE in its non-dual version, and for the sake of comparison our main focus will also be on that setting, i.e. we will mostly take $I = R$. In Section 4.3 we will look at the effect of switching to the dual case where $I = R^\vee$, and in Section 5 we will include the above hardness result in the discussion. Moreover, again as in [6], the noise parameter $\mathbf{r} = (r_1, \ldots, r_n)$ will usually be taken fixed and spherical, i.e. $r_1 = \cdots = r_n = r$.

3   Provably weak instances of non-dual decision RLWE

In [5], Eisentraeger, Hallgren and Lauter presented families of defining polynomials $f \in \mathbb{Z}[x]$ such that the *decision* version of PLWE is weak. This attack was later extended to non-dual decision RLWE [6] by Elias, Lauter, Ozman and Stange. In this section we recall the attack, first for PLWE and then how it transfers to non-dual RLWE. We provide a detailed analysis of the singular value decomposition of the matrix N_f for these polynomial families, since this will play an instructive role in our exposition.

3.1 Attack on decision PLWE

The simplest form of the attack on decision PLWE requires that the defining polynomial f of P and the modulus q satisfy the relation $f(1) \equiv 0 \bmod q$. This implies that evaluation at 1 induces a ring homomorphism $\phi : P_q \to \mathbb{Z}_q :$ $a(x) \mapsto a(1) \bmod q$. By applying ϕ to the PLWE samples $(\mathbf{a}_i, \mathbf{b}_i = \mathbf{a}_i \cdot \mathbf{s} + \mathbf{e}_i)$ we obtain tuples in $\mathbb{Z}_q^2$ namely $(\phi(\mathbf{a}_i), \phi(\mathbf{a}_i) \cdot \phi(\mathbf{s}) + \phi(\lfloor \mathbf{e}_i \rceil))$. Here $\lfloor \mathbf{e}_i \rceil$ denotes the polynomial obtained by rounding each coefficient of $\mathbf{e}_i$ to the nearest integer (with ties broken upward, say).

Assuming that the images of the error terms $\mathbf{e}_i$ under the homomorphism ϕ can be distinguished from uniform with sufficiently high probability, one obtains the following straightforward attack: for each guess $s \in \mathbb{Z}_q$ for the value of $\phi(\mathbf{s}) = \mathbf{s}(1) \bmod q$, compute the corresponding image of the (rounded) error term $\phi(\lfloor \mathbf{e}_i \rceil)$ as $\phi(\lfloor \mathbf{b}_i \rceil) - \phi(\mathbf{a}_i)s$, assuming that the guess is correct. If there exists an s such that the corresponding images $\phi(\lfloor \mathbf{e}_i \rceil)$ are more or less distributed like a discretized Gaussian, rather than uniform, the samples were indeed likely to be actual PLWE samples and the secret $\mathbf{s}$ satisfies $\mathbf{s}(1) = s$. If no such guess is found, the samples were likely to be uniform samples. The attack succeeds if the following three conditions are met:

1. $f(1) \equiv 0 \bmod q$,
2. q is small enough that $\mathbb{Z}_q$ can be enumerated,
3. $\phi(\lfloor \Gamma_r^n \rceil)$ is distinguishable from uniform $\mathfrak{U}(\mathbb{Z}_q)$.

Note that if $\mathbf{e}_i$ is sampled from Γ_r^n, then the $\mathbf{e}_i(1)$ are also Gaussian distributed but with parameter $\sqrt{n} \cdot r$. Therefore, as long as $\sqrt{n} \cdot r$ is sufficiently smaller than q, it should be possible to distinguish $\phi(\lfloor \Gamma_r^n \rceil)$ from uniform.

3.2 Attack on non-dual decision RLWE

The attack of Elias et al. on non-dual decision RLWE basically works by interpreting the RLWE samples as PLWE samples and then executing the above attack. For this approach to work, two requirements need to be fulfilled. Firstly, the ring of integers R of the number field K should be a quotient ring of the form $R = \mathbb{Z}[x]/(f)$, i.e. the number field should be monogenic.

The second condition deals with the difference between the error distributions of PLWE and non-dual RLWE. For PLWE one simply uses a spherical Gaussian Γ_r^n on $R \otimes \mathbb{R}$ with respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$, whereas the RLWE distribution $\Psi_{\mathbf{r}}$ is obtained by pulling back a near-spherical Gaussian distribution on H through the canonical embedding σ. With respect to the polynomial basis one can view $\Psi_{\mathbf{r}}$ as a near-spherical Gaussion that got distorted by $N_f \cdot B$. Since B is a unitary transformation, the only actual distortion comes from N_f.

The maximum distortion of N_f is captured by its spectral norm $s_1(N_f)$, i.e. its largest singular value. The other singular values are denoted by $s_i(N_f)$ ordered by size such that $s_n(N_f)$ denotes its smallest singular value. A spherical Gaussian distribution on H of parameter $\mathbf{r} = (r, r, \ldots, r)$ will therefore be

transformed into an elliptical Gaussian distribution on $R \otimes \mathbb{R} = K \otimes \mathbb{R}$ where the maximum parameter will be given by $s_1(N_f) \cdot r$. The attack on non-dual decision RLWE then proceeds by considering the samples with errors coming from $\Psi_{\mathbf{r}}$ as PLWE samples where the error is bounded by a spherical Gaussian with deviation $s_1(N_f) \cdot r$, with $r = \max(\mathbf{r})$.

For the attack to succeed we therefore need the following four conditions:

1. K is monogenic,
2. $f(1) \equiv 0 \bmod q$,
3. q is small enough that $\mathbb{Z}_q$ can be enumerated,
4. $r' = s_1(N_f) \cdot r$ is small enough such that $\phi(\lfloor \Gamma_{r'}^n \rceil)$ can be distinguished from uniform.

Again note that if $\mathbf{e}_i$ is bounded by Γ_r^n then the $\mathbf{e}_i(1)$ are bounded by a Gaussian with parameter $\sqrt{n} \cdot r' = \sqrt{n} \cdot s_1(N_f) \cdot r$. So the requirement is that the latter quantity is sufficiently smaller than q. In fact this is a very rough estimate, and indeed Elias et al. empirically observe in [6, §9] that their attack works more often than this bound predicts. We will explain this observation in Section 4.1.

In [6] the authors remark that given a parameter set (n, q, r) for PLWE, one cannot simply use the same parameter set for non-dual RLWE since the canonical embedding of the ring R into H might be very sparse, i.e. the covolume (volume of a fundamental domain) of $\sigma(R)$ in H might be very large. They therefore propose to scale up the parameter r by a factor of $|\det(M_f B)|^{1/n} = |\det(M_f)|^{1/n}$, which is the n-th root of the covolume. Thus given a PLWE parameter set (n, q, r), their corresponding RLWE parameter set reads $(n, q, \tilde{r})$ with $\tilde{r} = r \cdot |\det(M_f)|^{1/n}$.

3.3 Provably weak number fields for non-dual decision RLWE

The first type of polynomials to which the attack of [6] was applied are polynomials of the form $f_{n,a,b}$ with $a = 0$. More precisely they considered

$$f_{n,q} := f_{n,0,q-1} = x^n + q - 1 \,,$$

where $n \geq 1$ and q is a prime. Note that the roots of these polynomials are simply the primitive $2n$-th roots of unity scaled up by $(q-1)^{1/n}$. These polynomials satisfy $f_{n,q}(1) \equiv 0 \bmod q$ and are irreducible by Eisenstein's criterion whenever $q - 1$ has a prime factor with exponent one. As shown in [6, Proposition 3], the polynomials $f_{n,q}$ are monogenic whenever $q - 1$ is squarefree, n is a power of a prime ℓ, and $\ell^2 \nmid ((1-q)^n - (1-q))$. In particular it is easy to construct examples for $n = 2^k$.

The final missing ingredient is a bound on the spectral norm $s_1(N_f)$. In [6], a slightly different matrix M_f is used (it is a real matrix containing the real and imaginary parts of the roots of f). For use further down, we adapt the proof of [6, Proposition 4] to derive *all* singular values $s_i(N_f)$. Due to its practical importance we will only deal with the case where n is even, since we are particularly interested in the case where $n = 2^k$.

Proposition 1 (Adapted from [6, Proposition 4]). *Assume that $f_{n,q}$ is irreducible and that n is even, then the singular values $s_i(N_f)$ are given by*

$$s_i(N_f) = \frac{1}{\sqrt{n}(q-1)^{(i-1)/n}} \, .$$

PROOF: The roots of $f_{n,q}$ are given by $a \cdot \xi_{2n}^j$ for $0 < j < 2n$ and j odd, with $a = (q-1)^{1/n} \in \mathbb{R}^+$ and ξ_{2n} a primitive $2n$-th root of unity. To derive the singular values of $N_f = M_f^{-1}$ it suffices to derive the singular values of M_f. Recall that the u-th column of M_f (counting from 0) is given by

$$a^u \cdot (\xi_{2n}^u, \xi_{2n}^{3u}, \ldots, \xi_{2n}^{(2n-1)u})^t \, .$$

The (Hermitian) inner product of the u-th and v-th column is therefore given by

$$S = a^{u+v} \cdot \sum_{k=0}^{n-1} \xi_{2n}^{(2k+1)(u-v)} \, .$$

Since $\xi_{2n}^{2n+1} = \xi_{2n}$, we obtain that $\xi_{2n}^{2(u-v)} S = S$. For $u \neq v$ we have that $\xi_{2n}^{2(u-v)} \neq 1$, which implies that $S = 0$. For $u = v$ we obtain $S = na^{2u}$. This shows that the matrix M_f has columns that are orthogonal. The singular values of M_f can be read off from the diagonal of $\overline{M_f}^t \cdot M_f$, in particular $s_i(M_f) = \sqrt{n}a^{n-i}$ for $i = 1, \ldots, n$. This also shows that $s_i(N_f) = 1/(\sqrt{n}a^{i-1})$ for $i = 1, \ldots, n$. One finishes the proof by using that $a^n = q - 1$. $\qquad\square$

The above proposition gives $s_1(N_f) = 1/\sqrt{n}$ which is small enough for the attack described in Section 3.2 to apply. In [6, Section 9], two examples of this family were attacked, giving the following results:

$f_{n,q}$	q	r	$\tilde{r}$	samples per run	successful runs	time per run
$x^{192} + 4092$	4093	8.87	5440.28	20	1 of 10	25 sec
$x^{256} + 8190$	8191	8.35	8399.70	20	2 of 10	44 sec

Recall that $\tilde{r}$ is simply r scaled up by a factor $|\det(M_f)|^{1/n}$. We remark, as do Elias et al. [6, §9], that these two examples unfortunately do *not* satisfy that $q - 1$ is squarefree. As a consequence the RLWE problem is not set up in the full ring of integers of the number field $K = \mathbb{Q}[x]/(f)$. We will nevertheless keep using these examples for the sake of comparison; it should be clear from the exposition below that this is not a crucial issue.

As a second instance, the authors of [6] considered polynomials of the form $f_{n,a,b} = x^n + ax + b$ with $a \approx b$, again chosen such that $f_{n,a,b}(1) \equiv 0$ modulo q, which is assumed to be an odd prime. More precisely, they let $a = (q-1)/2 + \Delta$ and $b = (q-1)/2 - \Delta - 1$, or $a = q + \Delta$ and $b = q - \Delta - 1$, for a small value of Δ. Heuristically these polynomials also result in weak instances of non-dual

decision RLWE, even though the analysis cannot be made as precise as in the foregoing case. In particular, no explicit formula is known for the spectral norm $s_1(N_f)$, but in [6] a heuristic perturbation argument is given that implies that it is bounded by $\sqrt{\max(a, b)} \cdot \det(N_f)^{1/n}$ infinitely often. They ran their attack for the particular case where $q = 524287$, $\Delta = 1$, $a = q + \Delta$ and $b = q - \Delta - 1$:

$f_{n,a,b}$	q	r	$\tilde{r}$	samples per run	successful runs	time per run
$x^{128} + 524288x + 524285$	524287	8.00	45540	20	8 of 10	24 sec

4   A simple attack on search RLWE

We derive a very simple attack on *search* RLWE for the families and parameter sets considered by Elias, Lauter, Ozman and Stange in [6]. The attack is based on two observations.

Firstly, a unit ball in the H-space gets severely deformed when being pulled back to $K \otimes \mathbb{R}$ along the canonical embedding. With respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$ we end up with an ellipsoid whose axes have lengths $s_1(N_f), \ldots, s_n(N_f)$. For the first family of polynomials (i.e. where $a = 0$) this is a geometrically decreasing sequence, while for the second family this statement remains almost true. In particular a spherical Gaussian distribution $\Gamma_{\mathbf{r}}$ with $\mathbf{r} = (r, \ldots, r)$ on the H-space will result in a very skew elliptical Gaussian distribution on $K \otimes \mathbb{R}$ with parameters $s_1(N_f) \cdot r, \ldots, s_n(N_f) \cdot r$. For the choices of r (or in fact $\tilde{r}$) made by Elias et al., the errors along the shortest axes of the ellipsoid are so small that after rounding they become zero.

The second observation is that the axes of the error distribution ellipsoid coincide almost perfectly with the polynomial basis. Again for the first family this is exactly the case, while for the second family the distribution is consistent enough, in the sense that the axes do not line up perfectly, but the coordinates of the error samples with respect to $1, x, x^2, \ldots, x^{n-1}$ still tend to go down geometrically. The result is that the directions that get squashed simply correspond to the coefficients of the higher powers of x in the error terms $\mathbf{e}(x)$.

To make these statements precise we will compute the singular value decomposition of the whole transformation matrix $N_f \cdot B$. Recall that the singular value decomposition of an $n \times n$ matrix M is given by

$$M = U \Sigma \overline{V}^t,$$

where U, V are $n \times n$ unitary matrices and Σ is an $n \times n$ matrix with non-negative real numbers on the diagonal, namely the singular values. The image of a unit sphere under M will therefore result in an ellipsoid where the axes are given by the columns of U, with lengths equal to the corresponding singular values.

4.1   Singular value decomposition and error distribution

For the first family of polynomials $f_{n,q}$ everything can be made totally explicit:

Proposition 2. *The singular value decomposition of $N_f \cdot B$ is*

$$I_{n \times n} \cdot \Sigma \cdot \overline{V}^t, \qquad where \quad V = \overline{B}^t \cdot M_f \cdot \Sigma$$

and Σ is the diagonal matrix containing the singular values of N_f.

PROOF: Recall from the proof of Proposition 1 that the Vandermonde matrix M_f has mutually orthogonal columns, where the ith column has norm $\sqrt{n}a^{i-1}$. Thus the normalized matrix

$$M_f \cdot \Sigma \qquad where \ \Sigma = \mathrm{diag}\left(1/(\sqrt{n}a^{i-1})\right)_i = \mathrm{diag}\left(s_i(N_f)\right)_i$$

is unitary. But then so is $V = \overline{B}^t \cdot M_f \cdot \Sigma$, and since $\Sigma = \Sigma^2 \cdot \Sigma^{-1} = N_f \overline{N_f}^t \cdot \Sigma^{-1}$, we see that

$$N_f \cdot B = I_{n \times n} \cdot \Sigma \cdot \overline{V}^t$$

is the singular value decomposition of our transformation matrix $N_f \cdot B$. $\qquad \square$

The factor $I_{n \times n}$ implies that the axes of our ellipsoid match perfectly with the polynomial basis $1, x, x^2, \ldots, x^{n-1}$. In other words, if we start from a spherical error distribution $\Gamma_{\mathbf{r}}$ on H, $\mathbf{r} = (r, r, \ldots, r)$, then the induced error distribution $\Psi_{\mathbf{r}}$ on $K \otimes \mathbb{R}$ in the ith coordinate (coefficient of x^{i-1}) is a Gaussian with parameter

$$s_i(N_f) \cdot r = \frac{r}{\sqrt{n} \cdot (q-1)^{(i-1)/n}}$$

by Proposition 1. This indeed decreases geometrically with i.

As a side remark, note that this implies that for $\mathbf{e}(x) \leftarrow \Psi_{\mathbf{r}}$ the evaluation $\mathbf{e}(1)$ is sampled from a Gaussian with parameter

$$\left(\sum_{i=1}^{n} s_i(N_f)^2\right)^{1/2} \cdot r = s_1(N_f)\sqrt{\frac{(q-1)^2 - 1}{(q-1)^2 - (q-1)^{2(n-1)/n}}} \cdot r.$$

This is considerably smaller than $\sqrt{n} \cdot s_1(N_f) \cdot r$ and explains why the attack from [6] works better than what their theory predicts [6, §9].

To illustrate the geometric behavior of the coordinates of the errors $\mathbf{e}(x)$ with respect to the polynomial basis, we have plotted the average and standard deviation of their high order coefficients for the second example $x^{256} + 8190$ from [6] in Figure 1 (the results for the first example are totally similar), using the error parameter that they used to attack non-dual decision RLWE. The plot shows that for the given parameter set, the highest $\lceil n/7 \rceil$ error coefficients in the polynomial basis of $K \otimes \mathbb{R}$ are all extremely likely to be smaller than $1/2$ (indicated by the dashed line) in absolute value and therefore become zero after rounding.

For the second family of polynomials $f_{n,a,b}$ with $a \neq 0$, we were not able to derive the singular value decomposition in such an explicit form. To get a handle on them, we have computed it explicitly for $f = x^{128} + 524288x + 524285$. For
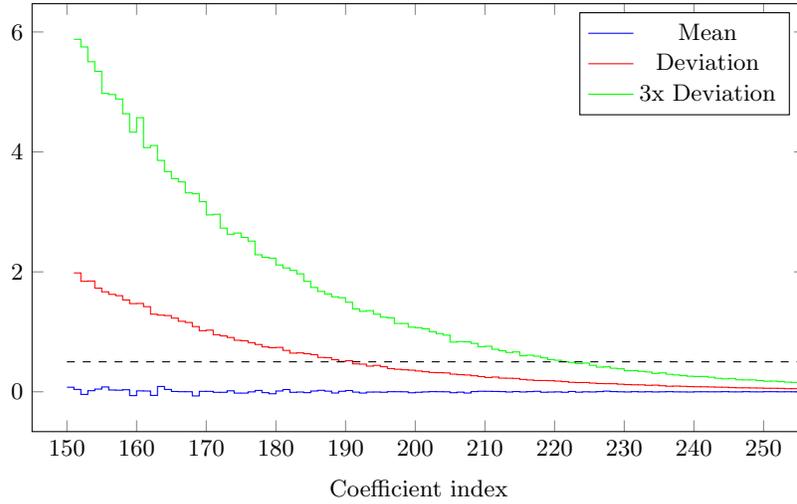
Fig. 1. Distribution of the error terms in the polynomial basis for $f = x^{256} + 8190$
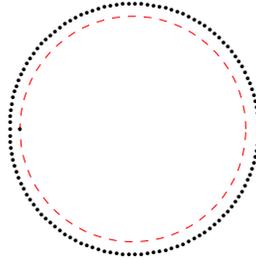

Fig. 2. Zeroes in $\mathbb{C}$ of $x^{128} + 524288x + 524285$, along with the unit circle (dashed)

this particular example, the roots of $f_{n,a,b}$ again lie roughly on a circle (except for the real root close to -1): see Figure 2. So through the Vandermonde matrix we again expect geometric growth of the singular values, as is confirmed by the explicit numerics in Figure 3, which shows a plot of their logarithms. There is only one outlier, caused by the real root of f close to -1.

The heat map in Figure 4 plots the norms of the entries in the U-matrix of the singular value decomposition of $N_f \cdot B$ and shows that U is close to being diagonal, implying that the axes of the ellipsoid are indeed lining up almost perfectly with the polynomial basis. Finally Figure 5 contains a similar plot as Figure 1, namely, the distribution of the errors terms (highest powers only) for the polynomial $f = x^{128} + 524288x + 524285$. Again we conclude that with very high probability, the last $\lceil n/6 \rceil$ coefficients of the error terms in the polynomial basis will be smaller than $1/2$, and therefore they become zero after rounding.
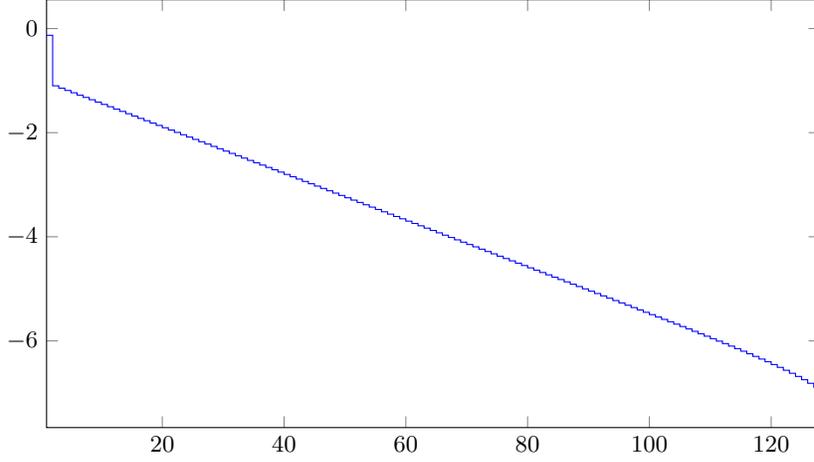
14

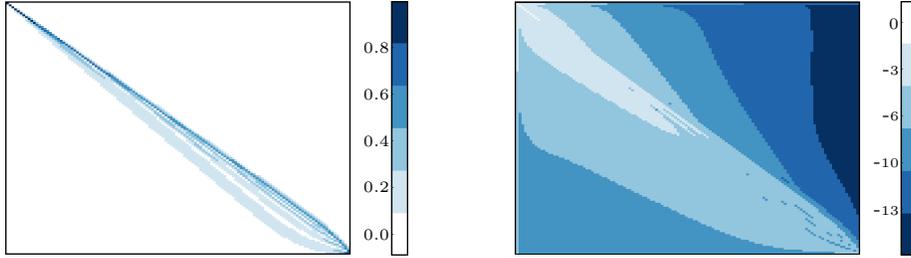**Fig. 3.** $\log_{10}$ of the singular values of N_f for $f = x^{128} + 524288x + 524285$

Fig. 4. Heat maps of the norms of the entries of U (left) and $\log_{10}$ of the norms of the entries of $U\Sigma$ (right), where $U\Sigma\overline{V}^t$ is the singular value decomposition of $N_f B$

4.2 Linear algebra attack on non-dual search RLWE

Turning the above observations into an attack on non-dual search RLWE for these families is straightforward. Recall that the samples are of the form $(\mathbf{a}, \mathbf{b} = \mathbf{a} \cdot \mathbf{s} + \mathbf{e} \bmod q)$ where the errors were sampled from the distribution $\Psi_{\mathbf{r}}$ on $K \otimes \mathbb{R}$. Since $\mathbf{a}$ is known, we can express multiplication by $\mathbf{a}$ as a linear operation, i.e. we can compute the $n \times n$ matrix $M_{\mathbf{a}}$ that corresponds to multiplication by $\mathbf{a}$ with respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$. Each RLWE sample can therefore be written as a linear algebra problem as follows:

$$M_{\mathbf{a}} \cdot (s_0, s_1, \ldots, s_{n-1})^t = (b_0, b_1, \ldots, b_{n-1})^t - (e_0, e_1, \ldots, e_{n-1})^t \qquad (1)$$

where the s_i (resp. b_i, e_i) are the coefficients of $\mathbf{s}$ (resp. $\mathbf{b}$ and $\mathbf{e}$) with respect to the polynomial basis. By rounding the coefficients of the right-hand side, we effectively remove the error terms of high index, which implies that the last equations in the linear system become *exact* equations in the unknown coefficients of
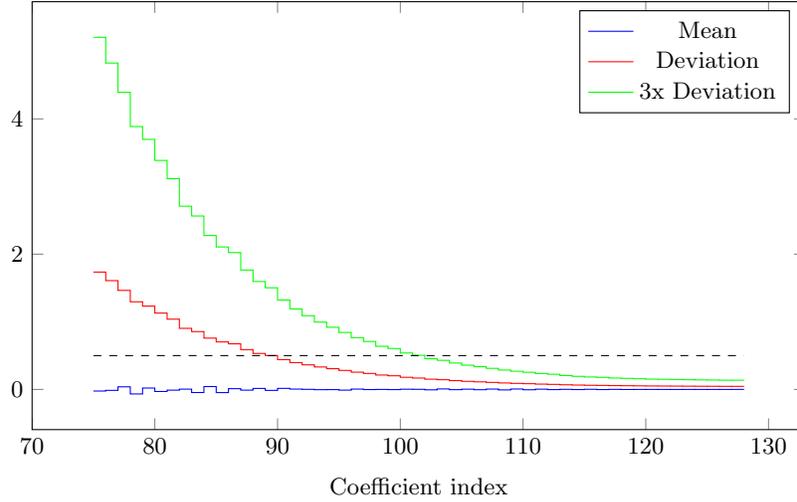
Fig. 5. Distribution of errors of high index for $f = x^{128} + 524288x + 524285$

s. Assuming that the highest $\lceil n/k \rceil$ error terms round to zero, we only require k samples to recover the secret **s** using simple linear algebra with a 100% success rate.

We have implemented this attack in Magma [2] with the following results.

$f_{n,a,b}$	q	r	$\tilde{r}$	samples per run	successful runs	time per run
$x^{192} + 4092$	4093	8.87	5440	7	10 of 10	8.37 sec
$x^{256} + 8190$	8191	8.35	8390	6	10 of 10	17.2 sec
$x^{128} + 524288x + 524285$	524287	8.00	45540	4	10 of 10	1.96 sec

We note that using less samples per run is also possible, but results in a lower than 100% success rate. A more elaborate strategy would construct several linear systems of equations by discarding some of the equations of lower index (which are most likely to be off by 1) and running exhaustively through the kernel of the resulting underdetermined system of equations. However, we did not implement this strategy since it needlessly complicates the attack.

In fact for errors of the above size one can also use the linearization technique developed by Arora and Ge [1, Theorem 3.1] to retrieve $\mathbf{s}(x)$, but this requires a lot more samples.

We stress that our attack does not use that $f(1) \equiv 0 \bmod q$. For the above defining polynomials our attack works modulo *every* modulus q', as long as the same error parameters are used (or smaller ones).

4.3 Modifications for dual search RLWE

In this section we discuss how switching from non-dual RLWE (i.e. from $I = R$) to dual RLWE (where one takes $I = R^\vee$) affects our observations. Recall that in the case of a monogenic defining polynomial f, the codifferent $R^\vee$ is generated as a fractional ideal by $1/f'(\alpha)$ with $\alpha \in \mathbb{C}$ a root of f. We will again work with respect to the polynomial basis $1, x, x^2, \ldots, x^{n-1}$ of $K = \mathbb{Q}[x]/(f)$ over $\mathbb{Q}$, which is also a basis of R over $\mathbb{Z}$, and take $\alpha = x$. For technical reasons we will only do the analysis for the first family of polynomials, namely those of the form

$$f_{n,q} = f_{n,0,q-1} = x^n + q - 1,$$

where one has $f'_{n,q} = nx^{n-1}$. Since

$$1 = \frac{1}{q-1}f_{n,q} - \frac{x}{n(q-1)}f'_{n,q}$$

we find that

$$R^\vee = R\frac{x}{n(q-1)}.$$

Proposition 3. *The elements*

$$\frac{1}{n}, \frac{x}{n(q-1)}, \frac{x^2}{n(q-1)}, \frac{x^3}{n(q-1)}, \ldots, \frac{x^{n-1}}{n(q-1)} \tag{2}$$

form a $\mathbb{Z}$-basis of $R^\vee$.

Proof. It is immediate that

$$\frac{x}{n(q-1)}, \frac{x^2}{n(q-1)}, \frac{x^3}{n(q-1)}, \ldots, \frac{x^{n-1}}{n(q-1)}, \frac{x^n}{n(q-1)}$$

form a $\mathbb{Z}$-basis. But modulo $f_{n,q}$ the last element is just $-1/n$.

Thus we can think of our secret $\mathbf{s}(x) \in R_q^\vee$ as a $\mathbb{Z}$-linear combination of the elements in (2), where the coefficients are considered modulo q. A corresponding RLWE-sample is then of the form $(\mathbf{a}(x), \mathbf{a}(x) \cdot \mathbf{s}(x) + \mathbf{e}(x) \bmod qR^\vee)$ with $\mathbf{e}(x) \in R^\vee \otimes \mathbb{R} = K \otimes \mathbb{R}$ sampled from $\Psi_\mathbf{r}$ for an appropriate choice of $\mathbf{r} \in (\mathbb{R}^+)^n$. To make a comparison with our attack in the non-dual case, involving the parameters from [6], we have to make an honest choice of $\mathbf{r}$, which we again take spherical. Note that the lattice $\sigma(R^\vee)$ is much denser than $\sigma(R)$: the covolume gets scaled down by a factor

$$\left|\mathrm{No}(f'_{n,q}(\alpha))\right| = n^n(q-1)^{n-1}.$$

Therefore, in view of the discussion concluding Section 3.2, we scale down our scaled-up error parameter $\tilde{r}$ by a factor

$$\sqrt[n]{n^n(q-1)^{n-1}} \approx n(q-1).$$

Let us denote the result by $\tilde{r}^\vee$.

It follows that the dual setting is essentially just a scaled version of its non-dual counterpart: both the errors and the basis elements become divided by a factor of roughly $n(q-1)$. In particular, for the same choices of r we again find that with near certainty the highest $\lceil n/7 \rceil$ error coefficients are all smaller than

$$\frac{1}{2} \cdot \frac{1}{n(q-1)}$$

in absolute value, and therefore become zero after rounding to the nearest multiple of $1/(n(q-1))$. This then again results in exact equations in the coefficients of the secret $\mathbf{s}(x) \in R_q^\vee$ with respect to the basis (2), that can be solved using linear algebra.

Here too, the attack does not use that $f(1) \equiv 0 \bmod q$ so it works for whatever choice of modulus q' instead of q, as long as the same error parameters are used (or smaller ones).

5 Range of applicability

One obvious way of countering our attack is by modifying the error parameter. In principle the skewness of $N_f \cdot B$ could be addressed by using an equally distorted elliptical Gaussian rather than a near-spherical one, but that conflicts with the philosophy of RLWE (as opposed to PLWE), namely that the more natural way of viewing a number field is through its canonical embedding. So we will not discuss this option and stick to spherical distributions. Then the only remaining way out is to enlarge the width of the distribution. Again for technical reasons we will restrict our discussion to the first family of polynomials, namely those of the form $f_{n,q} = x^n + q - 1$; the conclusions for the second family should be similar.

In the non-dual case we see that a version of the attack works as long as a sample drawn from a univariate Gaussian with parameter $s_n(N_f) \cdot \tilde{r}$ has absolute value less than $1/2$ with non-negligible probability: then by rounding one obtains at least one exact equation in the unknown secret $\mathbf{s}(x)$. For this one needs that

$$s_n(N_f) \cdot \tilde{r} \leq \frac{C}{2}$$

for some absolute constant $C > 0$ that quantifies what it means to be 'non-negligible'.

Remark 1. In order to recover the *entire* secret, one even wants a non-negligible probability for n consecutive samples to be less than $1/2$, for which one should replace $s_n(N_f) \cdot \tilde{r}$ by $s_n(N_f) \cdot \tilde{r} \cdot \sqrt{\log n}$ (roughly). In fact a slightly better approach is to find the optimal $1 \leq k \leq n$ for which $s_{n-k+1}(N_f) \cdot \tilde{r}$ is likely to be less than $1/2$, thereby yielding at least k exact equations at once, for $\lceil n/k \rceil$ consecutive times.

Let us take $C = 1$ in what follows: for this choice meeting the upper bound corresponds to a chance of about 98.78% of recovering at least one exact equation. Using Proposition 1 this can be rewritten as

$$\tilde{r} \leq \frac{1}{2} \cdot \sqrt{n} \cdot (q-1)^{1-1/n}. \tag{3}$$

For our two specific polynomials $x^{192} + 4092$ and $x^{256} + 8190$ the right-hand side reads 27148.39 and 63253.95 whereas Elias et al. took $\tilde{r}$ to be 5440.28 and 8399.70, respectively.

Note that the bound in (3) does not depend on the modulus q' that is being used: the q that appears there is just part of the data defining our number field. In other words, whenever $\tilde{r}$ satisfies (3) then for every choice of modulus q' we are very likely to recover at least one exact equation in the coefficients of the secret $\mathbf{s}(x)$.

Unfortunately the bound (3) does not allow for an immediate comparison with the hardness result of Lyubashevsky, Peikert and Regev (see Theorem 1), which was formulated for dual RLWE only. But for dual RLWE one can make a similar analysis. From Section 4.3 it follows that we want error coefficients that are smaller than $1/(2n(q-1))$ with a non-negligible probability. The same discussion then leads to the bound

$$\tilde{r}^{\vee} \leq \frac{1}{2} \cdot \frac{1}{\sqrt{n} \cdot (q-1)^{\frac{1}{n}}} \tag{4}$$

which is highly incompatible with the condition $\tilde{r}^{\vee} \geq 2\omega(\sqrt{\log n})$ from Theorem 1. Thus we conclude that it is impossible to enlarge the error parameter up to a range where our attack would form an actual threat to RLWE, as defined in [9, §3].

Another issue with modifying the error parameter is decodability. In the non-dual case, from (3) we see that $s_n(N_f) \cdot \tilde{r} \gg 1$ is needed to avoid being vulnerable to our skewness attack. But it automatically follows that $s_1(N_f) \cdot \tilde{r} \gg q$. Indeed, this is implied by the fact that the condition number $k(N_f) := s_1(N_f)/s_n(N_f)$ equals

$$(q-1)^{1-1/n} \approx q$$

by Proposition 1. This causes the errors at the terms of low degree to wrap around modulo q. In the dual case the same observation applies, where now the error terms of low degree tend to wrap around modulo multiples of $q \cdot 1/(n(q-1))$. In both cases the effect is that several of these terms become indistinguishable from uniform, requiring more samples for the RLWE problem to become information theoretically solvable. This obstructs, or at least complicates, certain cryptographic applications.

So overall, the conclusion is that the defining polynomials $f_{n,a,b}$ are just not well-suited for use in RLWE: either the error parameter is too small for the RLWE problem to be hard, or the error parameter is too large for the problem to be convenient for use in cryptography. But we stress once more that neither the attack from [6] nor our attack form a genuine threat to RLWE, as it was defined in [9, §3].

6   Conclusions

In this paper we have shown that non-dual *search* RLWE can be solved efficiently for the families of polynomials and parameter sets from [6] which were shown to be weak for the *decision* version of the problem. The central reason for this weakness lies in the (exponential) skewness of the canonical embedding transformation. We analyzed the singular value decomposition of this transformation and showed that the singular values form an (approximate) geometric series. Furthermore, we also showed that the axes of the error ellipsoid are consistent with the polynomial basis, allowing us to readily identify very small noise coefficients. The attack applies to wider ranges of moduli, and also applies to the dual version, but does not contradict any statement in the work of Lyubashevsky, Peikert and Regev [9].

It is worth remarking that while we used the language of singular value decomposition, for our skewness attack it merely suffices that $N_f \cdot B$ has a very short row, so that the corresponding error coefficient e_i vanishes after rounding and (1) provides an exact equation in the coefficients of the secret. For general number fields this is a strictly weaker condition than having a very small singular value whose corresponding axis lines up perfectly with one of the polynomial basis vectors. But for the particular families of [6] the singular value decomposition turned out to be a convenient tool in proving this, and in visualizing how the RLWE errors are transformed under pull-back along the canonical embedding.

Acknowledgments

References

1. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In *Automata, languages and programming. Part I*, volume 6755 of *Lecture Notes in Comput. Sci.*, pages 403–415. Springer, Heidelberg, 2011.
2. Wieb Bosma, John Cannon, and Catherine Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).
3. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Symposium on Theory of Computing Conference, STOC'13*, pages 575–584. ACM, 2013.
4. Léo Ducas and Alain Durmus. Ring-LWE in polynomial rings. In *Public Key Cryptography - PKC 2012*, volume 7293 of *Lecture Notes in Computer Science*, pages 34–51. Springer, 2012.
5. Kirsten Eisenträger, Sean Hallgren, and Kristin E. Lauter. Weak instances of PLWE. In *Selected Areas in Cryptography - SAC 2014*, volume 8781 of *Lecture Notes in Computer Science*, pages 183–194. Springer, 2014.

6. Yara Elias, Kristin E. Lauter, Ekin Ozman, and Katherine E. Stange. Provably weak instances of Ring-LWE. In *Advances in Cryptology - CRYPTO 2015 - Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 63–92. Springer, 2015.

7. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *Algorithmic Number Theory, ANTS-III*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

8. 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 Computer Science*, pages 319–339. Springer, 2011.

9. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *Advances in Cryptology - EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.

10. 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 Computer Science*, pages 35–54. Springer, 2013.

11. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Symposium on Theory of Computing, STOC 2009*, pages 333–342. ACM, 2009.

12. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Symposium on Theory of Computing*, pages 84–93. ACM, 2005.

13. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.