

How (Not) to Instantiate Ring-LWE

Chris Peikert*

July 18, 2016

Abstract

The *learning with errors over rings* (Ring-LWE) problem—or more accurately, family of problems—has emerged as a promising foundation for cryptography due to its practical efficiency, conjectured quantum resistance, and provable *worst-case hardness*: breaking certain instantiations of Ring-LWE is at least as hard as quantumly approximating the Shortest Vector Problem on *any* ideal lattice in the ring.

Despite this hardness guarantee, several recent works have shown that certain instantiations of Ring-LWE can be broken by relatively simple attacks. While the affected instantiations are not supported by worst-case hardness theorems (and were not ever proposed for cryptographic purposes), this state of affairs raises natural questions about what other instantiations might be vulnerable, and in particular whether certain classes of rings are inherently unsafe for Ring-LWE.

This work comprehensively reviews the known attacks on Ring-LWE and vulnerable instantiations. We give a new, unified exposition which reveals an elementary geometric reason why the attacks work, and provide rigorous analysis to explain certain phenomena that were previously only exhibited by experiments. In all cases, the insecurity of an instantiation is due to the fact that the error distribution is insufficiently “well spread” relative to the ring. In particular, the insecure instantiations use the so-called *non-dual* form of Ring-LWE, together with *spherical* error distributions that are much narrower and of a very different shape than the ones supported by hardness proofs.

On the positive side, we show that any Ring-LWE instantiation which satisfies (or only almost satisfies) the hypotheses of the “worst-case hardness of search” theorem is *provably immune* to broad generalizations of the above-described attacks: the running time divided by advantage is at least exponential in the degree of the ring. This holds for the ring of integers in *any* number field, so the rings themselves are not the source of insecurity in the vulnerable instantiations. Moreover, the hypotheses of the worst-case hardness theorem are *nearly minimal* ones which provide these immunity guarantees.

*Computer Science and Engineering, University of Michigan. Email: cpeikert@umich.edu. This material is based upon work supported by the National Science Foundation under CAREER Award CCF-1054495 and CNS-1606362, the Alfred P. Sloan Foundation, and by a Google Research Award. The views expressed are those of the authors and do not necessarily reflect the official policy or position of the National Science Foundation, the Sloan Foundation, or Google.

1 Introduction

Cryptography based on *lattices* is an exciting and fast-developing area of research, due in part to conjectured security against quantum attacks, asymptotic efficiency and parallelism, and a wide range of applications spanning from basic tasks like key exchange, to powerful objects like fully homomorphic encryption. A large fraction of lattice-based cryptographic constructions are built upon the average-case *learning with errors* (LWE) problem [Reg05] or its more efficient variant *learning with errors over rings* (Ring-LWE) [LPR10, LPR13]. These are actually *families* of problems, which are instantiated by choosing a particular dimension or ring, an integer modulus, and an error distribution.

A main attraction of the (Ring-)LWE problems is their *worst-case hardness theorems*, also known as worst-case/average-case reductions. These say that breaking certain instantiations of (Ring-)LWE (or the cryptosystems based upon them) is provably at least as hard as quantumly solving *any* instance of certain lattice problems, i.e., in the worst case. For Ring-LWE, the underlying lattice problem is the approximate Shortest Vector Problem (approx-SVP) on *ideal lattices*, which are algebraically structured lattices corresponding to ideals in the ring. To date, no known quantum (or classical) algorithm for approx-SVP has significantly better worst-case performance on ideal lattices (in any concrete ring of interest) than on general lattices of the same dimension. For the polynomial approximation factors often used in cryptography, the fastest known algorithms require exponential time and space in the dimension (see, e.g., [AKS01, MV10, ADRS15]).

Despite the above-described hardness guarantees, several recent works [EHL14, ELOS15, CLS15, CIV16, CLS16] have shown that certain *ad-hoc* instantiations of Ring-LWE are *insecure*, via relatively simple attacks. How should we interpret such results? We emphasize that the vulnerable instantiations were not previously proposed for any cryptographic application, and were specifically sought out for their insecurity.¹ In addition, the attacks do not appear to translate to any improved algorithms for ideal-lattice problems, because the vulnerable instantiations do not satisfy the hypotheses of the worst-case hardness theorems. Yet this explanation leaves several natural questions unanswered, such as:

- How “close” are these insecure instantiations to ones that do enjoy worst-case hardness?
- Do these vulnerable instantiations imply anything about what *rings* might be more or less secure for Ring-LWE?
- How can we evaluate other instantiations that may not be backed by worst-case hardness theorems?

The goals of this work are twofold: first, to shed further light on recent attacks and vulnerable Ring-LWE instantiations; and second, to articulate a general set of principles by which we can systematically evaluate the (in)security of an instantiation. Toward this end, we provide the following main contributions.

Review of attacks. We comprehensively review the attacks and insecure Ring-LWE instantiations from the above-cited works. We give a new, unified exposition of the attacks in terms of the geometry of *dual* ideals, and provide formal analysis to explain certain phenomena that were previously only exhibited by experiments. In all cases, the heart of the insecurity is the use of a non-standard, “*non-dual*” form of Ring-LWE with relatively narrow spherical error, rather than the “*dual*” form that was defined and proved to have worst-case hardness in [LPR10]. (See, e.g., Figures 3 and 7.) Using a simple “tweak” that allows for a direct comparison of the dual and non-dual forms, we find that the *error distributions* in the insecure instantiations are much narrower than those in the provably hard ones, which is why they are vulnerable to attacks.

¹Indeed, it is easy to design trivially insecure (Ring-)LWE instantiations for any choice of dimension or ring: just define the error distribution to always output zero. However, the vulnerable instantiations in question do involve some nontrivial error.

In a bit more technical detail: for the instantiations from [ELOS15], we give a simpler proof of the fact, first noticed and exploited in [CIV16], that the (discretized) errors lie in a very low-dimensional linear subspace of the ring. This means that every Ring-LWE sample reveals many errorless LWE samples, which leads to an elementary linear-algebraic attack (no ring algebra needed). We also show that the instantiations from [CLS15, CLS16], with slightly narrower error distributions, fall to the same kind of attack. Finally, we give formal analyses showing why the (unmodified) instantiations of [CLS15, CLS16] are broken by a different but closely related distinguishing attack.

Invulnerable instantiations. On the positive side, we consider Ring-LWE instantiations that satisfy, or only “almost” satisfy, the “worst-case hardness of search” theorem from [LPR10, Section 4]. We show that *any* such instantiation is *provably immune* to broad classes of attacks which include all those described above. By “immune” we mean that the attacks perform no better than known attacks (e.g., [BKW03, AG11]) against *plain* LWE when instantiated to have worst-case hardness; in particular, the running time divided by advantage is at least exponential in the LWE dimension (in the typical case of modulus $q = \text{poly}(n)$).

We stress that the worst-case hardness theorem from [LPR10, Section 4] works for the ring of integers (or more generally, any order) of *any* number field. Therefore, all the rings appearing in the insecure instantiations from the above-cited works do indeed admit provably hard instantiations—they just need different error distributions. In other words, the rings themselves are not the source of insecurity. For illustration, we describe and graphically depict some example hard instantiations in detail, including for prime cyclotomic rings and quadratic extensions thereof (see, e.g., Figures 6 and 7).

To be clear, in this work we do not propose *concrete* security estimates for particular (Ring-)LWE instantiations, e.g., “the m th cyclotomic with Gaussian error of width r offers at least λ bits of security” (see, e.g., [MR09, LP11, BCNS15, ADPS16] for representative works that do so). We are also not concerned with the *applicability* (or lack thereof) of instantiations for cryptographic purposes, nor with lower-level computational details or efficiency (see, e.g., [LPR13, CP15] for works along these lines). Our central focus is on understanding and evaluating the fundamental (in)security of Ring-LWE instantiations, which is a necessary prerequisite to these other important goals.

Discussion. The main conclusion from this work is that for the security of Ring-LWE, *proper choice of the error distribution is essential*, especially because there is so much more freedom of choice than in plain LWE. It should not be surprising that ad-hoc instantiations of Ring-LWE can be insecure—indeed, the same goes for LWE. For example, there is a roughly n^d -time attack (using roughly n^d samples) for d -bounded errors [AG11]. But this does not affect LWE’s conjectured $2^{\Omega(n)}$ hardness when instantiated according to its worst-case hardness theorems, which require Gaussian errors of standard deviation $\Omega(\sqrt{n})$. Indeed, it may even *increase* our confidence that this is the “right” error distribution for LWE, since the wide variety of known attack strategies all require $2^{\Omega(n)}$ time beyond this threshold.

On the positive side, the fact that worst-case-hard instantiations are immune to concrete attacks also should not be surprising, since any efficient attack would translate to a comparably efficient quantum algorithm for approx-SVP on any ideal lattice in the ring—which would be a major achievement in computational number theory. But it is instructive to understand what precisely gives the hard instantiations their immunity. In particular, some of the more peculiar aspects of Ring-LWE, like the width of the error distributions and especially the role of the “dual” ideal R^\vee , were adopted in order to obtain the strongest and tightest hardness theorems in general number fields. (See [LPR10, Section 3.3] for discussion.) Notably, these choices also turn out to be *nearly minimal* ones that provably withstand broad classes of attacks. We believe that this provides yet another example of the importance of worst-case hardness proofs in lattice cryptography.

Organization. The remainder of the paper is organized as follows.

Section 2 recalls the relevant mathematical background, the (Ring-)LWE problems and the formal relationship between them, and their known worst-case hardness theorems.

Section 3 gives a new exposition and unified framework for the Ring-LWE attacks developed in [EHL14, ELOS15, CLS15, CIV16, CLS16], focusing on the essential geometric reasons why they work.

Section 4 reviews the insecure Ring-LWE instantiations through the lens of the unified attack framework, and formally proves that the attacks work against them.

Section 5 gives a sufficient condition that makes a Ring-LWE instantiation *provably immune* to the attacks in the framework, and shows that the condition is satisfied for any instantiation supported by the worst-case hardness theorem of [LPR10, Section 4].

Acknowledgments. I thank Léo Ducas, Kristin Lauter, Vadim Lyubashevsky, Oded Regev, and Katherine Stange for many valuable discussions and comments on topics related to this work. I also thank the anonymous reviewers for helpful comments, and especially for pointing out a misinterpretation of the parameters in [ELOS15, Section 9].

2 Preliminaries

In this section we recall the necessary mathematical background on lattices, Gaussians, algebraic number theory, and (Ring-)LWE, including its “dual” and “non-dual” forms. We closely follow the presentation from [LPR10]; see that work for further details.

2.1 Lattices and Gaussians

In ring-based lattice cryptography, it is convenient to work in the space $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ for some nonnegative integers s_1, s_2 with $n = s_1 + 2s_2$, defined as

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

It is easy to check that H , with the inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_i x_i \overline{y_i}$ of the ambient space \mathbb{C}^n , is an n -dimensional real inner product space, i.e., it is isomorphic to \mathbb{R}^n via an appropriate rotation. Therefore, the reader may mentally replace H with \mathbb{R}^n in all that follows.

For the purposes of this work, a *lattice* \mathcal{L} is a discrete additive subgroup of H that is full dimensional, i.e., $\text{span}_{\mathbb{R}}(\mathcal{L}) = H$. Any lattice is generated as the set of all integer linear combinations of some (non-unique) linearly independent basis vectors $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$, as

$$\mathcal{L} = \mathcal{L}(\mathbf{B}) := \left\{ \mathbf{B}\mathbf{z} = \sum_i z_i \mathbf{b}_i : \mathbf{z} \in \mathbb{Z}^n \right\}.$$

The *volume* (or determinant) of a lattice \mathcal{L} is $\text{vol}(\mathcal{L}) := \text{vol}(H/\mathcal{L}) = |\det(\mathbf{B})|$, where \mathbf{B} denotes any basis of \mathcal{L} . The *minimum distance* $\lambda_1(\mathcal{L})$ is the length (in the Euclidean norm) of a shortest nonzero lattice vector: $\lambda_1(\mathcal{L}) = \min_{\mathbf{0} \neq \mathbf{x} \in \mathcal{L}} \|\mathbf{x}\|$. The *dual lattice* \mathcal{L}^\vee of \mathcal{L} is defined as the set of all points in H having integer inner products with every vector of the lattice: $\mathcal{L}^\vee = \{\mathbf{w} \in H : \langle \mathbf{w}, \mathcal{L} \rangle \subseteq \mathbb{Z}\}$. It is easy to verify that $(\mathcal{L}^\vee)^\vee = \mathcal{L}$.

Gaussians and smoothing. For $r > 0$, the Gaussian probability distribution D_r of parameter (or width) r over H is defined to have probability density function proportional to $\rho_r(\mathbf{x}) := \exp(-\pi\|\mathbf{x}\|^2/r^2)$. A standard fact is that $\langle \mathbf{w}, D_r \rangle = D_{r\|\mathbf{w}\|}$ (over \mathbb{R}) for any nonzero $\mathbf{w} \in H$. In addition, a one-dimensional Gaussian D_r over \mathbb{R} satisfies the tail bound $\Pr_{x \leftarrow D_r}[|x| \geq t] \leq 2 \exp(-\pi(t/r)^2)$ for any $t \geq 0$.

The *smoothing parameter* [MR04] is an important lattice quantity that is related to several other lattice parameters.

Definition 2.1. For a lattice \mathcal{L} and positive real $\varepsilon > 0$, the *smoothing parameter* $\eta_\varepsilon(\mathcal{L})$ is the smallest $r > 0$ such that $\rho_{1/r}(\mathcal{L}^\vee \setminus \{0\}) \leq \varepsilon$.

Lemma 2.2 ([MR04, Lemma 3.2]). For any n -dimensional lattice \mathcal{L} , we have $\eta_{2^{-2n}}(\mathcal{L}) \leq \sqrt{n}/\lambda_1(\mathcal{L}^\vee)$.²

The following lemma explains the name ‘‘smoothing parameter:’’ it says that a Gaussian whose width exceeds the smoothing parameter is essentially uniform modulo the lattice.

Lemma 2.3 ([MR04, Lemma 4.1]). For any lattice $\mathcal{L} \subset H$, $\varepsilon > 0$, and $r \geq \eta_\varepsilon(\mathcal{L})$, the statistical distance between $D_r \bmod \mathcal{L}$ and the uniform distribution over H/\mathcal{L} is at most $\varepsilon/2$.

2.2 Algebraic Number Theory

In this subsection we review standard concepts from algebraic number theory, including: number fields, their rings of integers, (fractional) ideals, the canonical embedding, ideal lattices, and dual ideals.

A *number field* K is a finite-degree field extension of the rationals \mathbb{Q} . More concretely, a number field can always be constructed as $K = \mathbb{Q}(\zeta) \cong \mathbb{Q}[X]/(f(X))$, where ζ denotes an element that satisfies the relation $f(\zeta) = 0$ for some monic irreducible polynomial $f(x) \in \mathbb{Q}[X]$, called the *minimal polynomial* of ζ . The degree n of K (over \mathbb{Q}) is the degree of f , and K can be seen as an n -dimensional vector space over \mathbb{Q} with *power basis* $\{1, \zeta, \dots, \zeta^{n-1}\}$; of course, this is just one possible basis among infinitely many.

Examples of number fields that we encounter later in this work include quadratic fields, which can be expressed as $K = \mathbb{Q}(\sqrt{d})$ for some square-free integer $d \neq 0, 1$, and cyclotomic fields, which can be expressed as $K = \mathbb{Q}(\zeta_m)$ for some positive integer m , where ζ_m denotes an element of multiplicative order m ; the degree of K over \mathbb{Q} is $n = \varphi(m)$, the totient of m .

2.2.1 Embeddings and Geometry

A number field $K = \mathbb{Q}(\zeta)$ of degree n has exactly n *ring embeddings* (i.e., injective ring homomorphisms) into the complex numbers, denoted $\sigma_i: K \rightarrow \mathbb{C}$. Concretely, each embedding is defined by mapping ζ to one of the (real or complex) roots of the minimal polynomial f of ζ . An embedding whose image lies in \mathbb{R} (corresponding to a real root of f) is called a *real embedding*, otherwise (for a complex root of f) it is a *complex embedding*. The complex embeddings come in conjugate pairs, just as the complex roots of f do. Numbering the s_1 real embeddings by $\sigma_1, \dots, \sigma_{s_1}$, and the $s_2 = (n - s_1)/2$ pairs of complex embeddings so that the $\sigma_{s_1+s_2+j} = \overline{\sigma_{s_1+j}}$ for $j = 1, \dots, s_2$, the *canonical embedding* $\sigma: K \rightarrow H$ (where $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ is as defined in Section 2.1) is defined as the concatenation of all the embeddings:

$$\sigma(x) := (\sigma_1(x), \dots, \sigma_n(x)).$$

²Note that we have $\varepsilon = 2^{-2n}$ instead of 2^{-n} as in [MR04], but the proof is exactly the same.

Notice that σ is a ring homomorphism from K to $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$, where addition and multiplication in the latter are both component-wise.

We endow K with a geometry by identifying elements with their canonical embeddings in H . For example, we define the Euclidean norm of $x \in K$ as $\|x\| := \|\sigma(x)\|$. Similarly, we can think of the Gaussian distribution D_r over H as a distribution over K as well, via σ^{-1} . To be formal, because $\sigma(K)$ does not equal H but is merely dense within it, the distribution $\sigma^{-1}(D_r)$ is over $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R}$, which is isomorphic to $H \cong \mathbb{R}^n$ as a real inner product space. For our purposes, there is no harm in viewing D_r as being over K with sufficiently large finite precision.

Notice that for the Gaussian D_r over H , the s_1 real coordinates are independent Gaussians of parameter r , and the real and imaginary parts of the $2s_2$ complex coordinates are independent Gaussians, up to conjugate symmetry, of parameter $r/\sqrt{2}$. (The $\sqrt{2}$ factor is due to the duplication in the conjugate pairs.) Because multiplication in $K_{\mathbb{R}}$ corresponds to coordinate-wise multiplication in H , for any $t \in K_{\mathbb{R}}$ the distribution $t \cdot D_r$ is simply D_r with its i th coordinate scaled by a $|\sigma_i(t)|$ factor.

2.2.2 Trace and Norm

For any $x \in K$, multiplication by x corresponds to a linear transform on K (viewed as a vector space over \mathbb{Q}); fixing a \mathbb{Q} -basis of K represents the transform by a concrete matrix in $\mathbb{Q}^{n \times n}$. The *trace* $\text{Tr}: K \rightarrow \mathbb{Q}$ and (algebraic) *norm* $N: K \rightarrow \mathbb{Q}$ of x are respectively the trace and determinant of this transform, or of any matrix representing it (recall that trace and determinant are invariant under change of basis). In particular, they correspond to the sum and product, respectively, of the embeddings: $\text{Tr}(x) = \sum_i \sigma_i(x)$ and $N(x) = \prod_i \sigma_i(x)$. Notice that for any $x, y \in K$, we have

$$\text{Tr}(x \cdot y) = \sum_i \sigma_i(x) \cdot \sigma_i(y) = \langle \sigma(x), \overline{\sigma(y)} \rangle,$$

so the “trace pairing” $\text{Tr}(x \cdot y)$ is a symmetric bilinear form on K , akin to the inner product on H .

2.2.3 Ring of Integers and (Fractional) Ideals

An *algebraic integer* is an element whose minimal polynomial over the rationals has integer coefficients. For a number field K , let $R = \mathcal{O}_K \subset K$ denote the set of all algebraic integers in K ; this set forms a subring of K , and is called its *ring of integers*. This ring is a free \mathbb{Z} -module of rank n , i.e., it is the set of all integer linear combinations of some basis elements $(b_1, \dots, b_n) \subset R$, which are \mathbb{Q} -linearly independent.

For example, in the quadratic number field $K = \mathbb{Q}(\sqrt{d})$ for square-free integer $d \neq 0, 1$, the ring of integers \mathcal{O}_K is $\mathbb{Z}[\sqrt{d}] = \{a + b\sqrt{d} : a, b \in \mathbb{Z}\}$ if $d \not\equiv 1 \pmod{4}$, otherwise it is $\mathbb{Z}[(1 + \sqrt{d})/2]$. For the m th cyclotomic number field $K = \mathbb{Q}(\zeta_m)$, the ring of integers happens to be $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$.

An (*integral*) *ideal* $\mathcal{I} \subseteq R$ is a nontrivial additive subgroup that is closed under multiplication by R , i.e., $r \cdot x \in \mathcal{I}$ for any $r \in R, x \in \mathcal{I}$. Like $R = \mathcal{O}_K$ itself, any ideal is a free \mathbb{Z} -module of rank n , i.e., it has a \mathbb{Z} -basis of size n . For two ideals \mathcal{I}, \mathcal{J} , their product ideal $\mathcal{I}\mathcal{J}$ is the set of all finite sums of terms xy for $x \in \mathcal{I}, y \in \mathcal{J}$. The *norm* of an ideal is its index as an additive subgroup of R , i.e., $N(\mathcal{I}) = |R/\mathcal{I}|$. This generalizes the algebraic norm defined above, in the sense that $N(xR) = |N(x)|$ for any $x \in R$, and $N(\mathcal{I}\mathcal{J}) = N(\mathcal{I})N(\mathcal{J})$.

More generally, a *fractional ideal* $\mathcal{I} \subset K$ is a set such that $d\mathcal{I} \subseteq R$ is an integral ideal for some $d \in R$. Its norm is defined as $N(\mathcal{I}) = N(d\mathcal{I})/|N(d)|$, and multiplication of fractional ideals is defined in the expected way, so the norm is multiplicative on fractional ideals. Finally, the fractional ideals form a group under multiplication, with R as the identity.

2.2.4 Ideal Lattices

Any fractional ideal $\mathcal{I} \subset K$ maps, via the canonical embedding σ , to a lattice $\mathcal{L} = \sigma(\mathcal{I}) \subset H$, which is called an *ideal lattice*. Naturally, if (b_1, \dots, b_n) is a \mathbb{Z} -basis of \mathcal{I} , then $\mathbf{B} = \sigma(B) := (\sigma(b_1), \dots, \sigma(b_n))$ is a basis of \mathcal{L} . We identify ideals with their corresponding ideal lattices, which lets us refer to the volume $\text{vol}(\mathcal{I}) := \text{vol}(\sigma(\mathcal{I}))$ of an ideal, the minimum distance $\lambda_1(\mathcal{I}) := \lambda_1(\sigma(\mathcal{I}))$, etc.

The (absolute) *discriminant* $\Delta_K = \text{vol}(R)^2$ of a number field K is the squared volume of its ring of integers $R = \mathcal{O}_K$, viewed as a lattice; equivalently, $\Delta_K = |\det(\text{Tr}(b_i \cdot b_j))| = |\det(\mathbf{B}^* \cdot \mathbf{B})|$ where $\mathbf{B} = \sigma(B)$ for an arbitrary \mathbb{Z} -basis $B = (b_1, \dots, b_n)$ of R . A useful dimension-normalized quantity is the *root discriminant*

$$\delta_K := \sqrt{\Delta_K}^{-1/n} = \text{vol}(R)^{1/n}$$

(sometimes also denoted δ_R), which is a measure of the ‘‘sparsity’’ of the algebraic integers in K . It follows directly from the definition that $\text{vol}(\mathcal{I}) = N(\mathcal{I}) \cdot \sqrt{\Delta_K}$ for any fractional ideal \mathcal{I} in K . The following standard fact is an immediate consequence of Minkowski’s first theorem (for the upper bound) and the arithmetic mean-geometric mean inequality (for the lower bound).

Lemma 2.4. *For any fractional ideal \mathcal{I} in a number field K of degree n ,*

$$\sqrt{n} \cdot N(\mathcal{I})^{1/n} \leq \lambda_1(\mathcal{I}) \leq \sqrt{n} \cdot N(\mathcal{I})^{1/n} \cdot \delta_K.$$

For example, a quadratic number field $K = \mathbb{Q}(\sqrt{d})$ for square-free integer $d \neq 0, 1$ has absolute discriminant $\Delta_K = 4|d|$ if $d \not\equiv 1 \pmod{4}$, and $\Delta_K = |d|$ otherwise. (The difference arises from the different form of the ring of integers in the two cases.) A cyclotomic number field $K = \mathbb{Q}(\zeta_p)$ for prime p has absolute discriminant $\Delta_K = p^{p-2}$, so $\delta_K = \sqrt{p}^{(p-2)/(p-1)} < \sqrt{p}$.

2.2.5 Duality

Here we recall the notion of the dual ideal under the trace pairing, and its connection to the dual lattice; see [Con09] for further details. For any fractional ideal $\mathcal{I} \subset K$, its *dual ideal* \mathcal{I}^\vee is defined as

$$\mathcal{I}^\vee = \{x \in K : \text{Tr}(x\mathcal{L}) \subseteq \mathbb{Z}\}.$$

Notice that the dual ideal \mathcal{I}^\vee embeds as the *conjugate* of the dual lattice of \mathcal{I} , i.e., $\sigma(\mathcal{I}^\vee) = \overline{\sigma(\mathcal{I})}^\vee$, because $\text{Tr}(x \cdot y) = \langle \sigma(x), \overline{\sigma(y)} \rangle$. It is easy to verify that \mathcal{I}^\vee is a fractional ideal, and that $(\mathcal{I}^\vee)^\vee = \mathcal{I}$. Also, if $B = (b_j)$ is a \mathbb{Z} -basis of \mathcal{I} , then its dual basis $B^\vee = (b_j^\vee)$, which is characterized by $\text{Tr}(b_j \cdot b_{j'}^\vee) = \delta_{j,j'}$, is a \mathbb{Z} -basis of \mathcal{I}^\vee . Finally, it turns out that $\mathcal{I}^\vee = \mathcal{I}^{-1} \cdot R^\vee$, where R^\vee is the dual ideal of the ring of integers $R = \mathcal{O}_K$.

For example, in a quadratic number field $K = \mathbb{Q}(\sqrt{d})$ for square-free integer $d \neq 0, 1$ with $R = \mathcal{O}_K$, one can verify that the dual ideal R^\vee is $(2\sqrt{d})^{-1}R$ if $d \not\equiv 1 \pmod{4}$, and is $(\sqrt{d})^{-1}R$ otherwise. As another example, in a cyclotomic number field $K = \mathbb{Q}(\zeta_p)$ for prime p with $R = \mathcal{O}_K$, it is not hard to verify that the dual ideal $R^\vee = p^{-1}(1 - \zeta_p)R$.

2.3 Learning With Errors (Over Rings)

In this section we review the learning with errors problem [Reg05] and its ring-based analogue [LPR10], describe the formal relationship between them, and recall their worst-case hardness theorems.

2.3.1 LWE

Informally, *learning with errors* (LWE) [Reg05] concerns “noisy” random inner products with a secret vector. More precisely, LWE is parameterized by a dimension n , a positive integer modulus q defining the quotient ring $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$, and an error distribution ψ over \mathbb{R} . (Throughout this work we mainly use continuous rather than discrete error distributions because they are easier to analyze, and because they more easily expose the essential ideas. This has no significant effect on the final results.)

Definition 2.5 (LWE, [Reg05]). The *search-LWE* $_{n,q,\psi}$ problem is to recover a uniformly random secret vector $\mathbf{s} \in \mathbb{Z}_q^n$, given many independent samples of the form

$$(\mathbf{a}_i, b_i = \langle \mathbf{s}, \mathbf{a}_i \rangle + e_i \bmod q) \in \mathbb{Z}_q^n \times \mathbb{R}/q\mathbb{Z},$$

where each $\mathbf{a}_i \leftarrow \mathbb{Z}_q^n$ is uniformly random and each $e_i \leftarrow \psi$ is drawn from the error distribution. The *decision-LWE* $_{n,q,\psi}$ problem is to distinguish, with some noticeable advantage, between samples generated as above, and uniformly random samples in $\mathbb{Z}_q^n \times (\mathbb{R}/q\mathbb{Z})$.

Sometimes the number m of available samples is also considered as an additional parameter of the LWE problems, but here we let it be arbitrarily large; this can only make the problems easier to solve, because samples can be ignored. It is often convenient to group the m samples (\mathbf{a}_i, b_i) into a matrix and vector

$$\mathbf{A} = [\mathbf{a}_1 \mid \mathbf{a}_2 \mid \cdots \mid \mathbf{a}_m] \in \mathbb{Z}_q^{n \times m} \quad \text{and} \quad \mathbf{b}^t = \mathbf{s}^t \mathbf{A} + \mathbf{e}^t \in (\mathbb{R}/q\mathbb{Z})^m,$$

where \mathbf{A} is uniformly random and $\mathbf{e} \in \mathbb{R}^m$ is distributed as ψ^m .

Insecure instantiations. Certain instantiations of LWE are trivially easy to solve. For instance, if the error distribution ψ always outputs 0—i.e., no error at all—then the problem is easily solved by standard linear algebra: as long as the rows of \mathbf{A} are linearly independent over \mathbb{Z}_q (which holds with high probability once m is a little more than n), we can easily recover \mathbf{s} given \mathbf{A} and $\mathbf{b}^t = \mathbf{s}^t \mathbf{A}$. More generally, the same holds if ψ (almost) always outputs a value in the interval $z + [-\frac{1}{2}, \frac{1}{2})$ for some fixed integer z , because we can “round away” the non-integral part and subtract z to remove the error from every sample.

Now, suppose we generalize LWE to allow potentially non-independent errors, i.e., each group of k samples has an error vector drawn from some distribution κ over \mathbb{R}^k . Then this form of LWE is easy if, e.g., some (discretized) error coordinate is always zero under κ (just ignore the samples corresponding to the other coordinates), or if the sum of the k error coordinates is always zero (just sum the samples in each group to get an errorless sample).

Other instantiations of LWE can be solved by less obvious means. For example, if all the (discretized) errors in our samples lie in a known set of size d , then we can solve search-LWE in roughly n^d time and space, using roughly n^d samples [AG11]. For any $d = O(1)$ this yields a polynomial-time attack, and for $d = n^{1-\epsilon}$ it yields a subexponential-time and -space attack, in both cases assuming we have enough samples.

Hard instantiations. Certain instantiations of LWE appear computationally hard, and have strong “worst-case hardness” theorems in support of this belief. Specifically, for a Gaussian error distribution $\psi = D_r$ with $r \geq 2\sqrt{n}$, solving search-LWE $_{n,q,\psi}$ is at least as hard as *quantumly* approximating certain well-studied “short vector” problems on *any* n -dimensional lattice to within $\tilde{O}(n \cdot q/r)$ factors, i.e., there is a quantum reduction from worst-case lattice problems to search-LWE [Reg05]. Moreover, for $q \geq 2^{n/2}$ there is a *classical* reduction from a subset of these problems, for essentially the same approximation factors [Pei09].

Finally, under mild conditions on the modulus q and the Gaussian parameter r , the search and decision problems are equivalent, i.e., there are reductions from search to decision. See, e.g., [Reg05, Pei09, MM11, MP12, BLP⁺13].

2.3.2 Ring-LWE

Analogously to LWE, *learning with errors over rings* (Ring-LWE) [LPR10, LPR13] concerns “noisy” random *ring* products with a secret ring element. Formally, it is parameterized by a ring R , which is the ring of integers (or more generally, an order) of a number field K , a positive integer modulus q , and an error distribution ψ over $K_{\mathbb{R}}$. Recall that $R^{\vee} = \{x \in K : \text{Tr}(xR) \subseteq \mathbb{Z}\}$ is the (fractional) dual ideal of R , and for any fractional ideal \mathcal{I} define the quotient $\mathcal{I}_q := \mathcal{I}/q\mathcal{I}$.

Definition 2.6 (Ring-LWE, [LPR10]). The *search- R -LWE* $_{q,\psi}$ problem is to find a uniformly random secret $s \in R_q^{\vee}$ given many independent samples of the form

$$(a_i, b_i = s \cdot a_i + e_i \bmod qR^{\vee}) \in R_q \times K_{\mathbb{R}}/qR^{\vee},$$

where each $a_i \leftarrow R_q$ is uniformly random and each $e_i \leftarrow \psi$ is drawn from the error distribution. (Observe that each $s \cdot a_i \in R_q^{\vee}$.) The *decision- R -LWE* $_{q,\psi}$ problem is to distinguish, with some noticeable advantage, between samples generated as above, and uniformly random samples in $R_q \times K_{\mathbb{R}}/qR^{\vee}$.

The above definition is sometimes called the “dual” form of Ring-LWE owing to the appearance of R^{\vee} , whose role might appear somewhat mysterious. However, its importance for obtaining the “right” definition of Ring-LWE is discussed at length in [LPR10, Section 3.3]. In short, the combination of R^{\vee} and (nearly) *spherical* Gaussian error ψ yields both the tightest connection with worst-case problems on ideal lattices, and the best error tolerance and computational efficiency in cryptographic applications. (See [LPR10, LPR13] for full details.) Nevertheless, for various reasons it may be more convenient to work with a “non-dual” form of Ring-LWE, where the secret is a uniformly random $s \in R_q$ (not R_q^{\vee}), and samples are of the form

$$(a_i, b_i = s \cdot a_i + e_i \bmod qR) \in R_q \times K_{\mathbb{R}}/qR,$$

where each $a_i \leftarrow R_q$ is uniform and each $e_i \leftarrow \psi$.

It turns out that the dual and non-dual forms of Ring-LWE are in fact *equivalent up to the choice of error distribution* ψ —so it does not really matter which syntactic form we use, as long as long as we also use an appropriate error distribution. This is because we can always convert one form to another using an appropriate “tweak” factor, as described at the end of this section. (Such a “tweaked” form of Ring-LWE, which replaces R^{\vee} by R , has been used in [AP13, Pei14, CP15].) However, it is important to note that the transformation may in general convert spherical Gaussian error to (highly) non-spherical error.

Because R is usually much sparser than R^{\vee} (viewed as lattices in the canonical embedding), for non-dual Ring-LWE some prior works have used a volume-normalized parameter $r_0 = r/\text{vol}(R)^{1/n} = r/\delta_R$ as a rough measure of how “wide” Gaussian error D_r is relative to R . We caution, however, that this measure does not account for any lack of “well-roundedness” in R (i.e., gaps among its successive minima), so even a large value of r_0 does not necessarily mean that the error is “well-spread” relative to R .

Hard instantiations. Much like LWE, certain instantiations of Ring-LWE are supported by worst-case hardness theorems; see [LPR10] for formal statements, which we summarize here. For $r \geq 2 \cdot \omega(\sqrt{\log n})$, [LPR10, Theorem 4.1] says that for *any* number field K and $R = \mathcal{O}_K$, solving search- R -LWE for all

continuous Gaussian error distributions $\psi = D_{\mathbf{r}}$, where each $r_i \leq r$, is at least as hard as quantumly approximating certain “short vector” problems on any ideal lattice in K , to within $\tilde{O}(\sqrt{n} \cdot q/r)$ factors. (The distribution $D_{\mathbf{r}}$ over H is essentially an elliptical Gaussian with parameter r_i in the i th coordinate.) Moreover, [LPR10, Section 5] shows that for any *cyclotomic* number field, and for appropriate moduli q , decision is classically at least as hard as search for any *spherical* Gaussian error distribution. (The proof immediately generalizes to any Galois number field [EHL14].) Alternatively, decision for spherical error of parameter roughly $r \cdot n^{1/4}$ is classically at least as hard as search for the class of elliptical distributions $D_{\mathbf{r}}$ described above. (The conditions on the modulus q have subsequently been weakened, and hardness of decision is now known for essentially any large enough modulus, via “modulus switching;” see, e.g., [BLP⁺13].)

Connection to LWE. Ring-LWE can be seen as a special case of LWE, in the following sense. For simplicity we describe a reduction for the “non-dual” form, but it easily generalizes to the dual form from Definition 2.6, either via the dual/non-dual equivalence described below, or directly using a \mathbb{Z} -basis of R^\vee .

Fix any \mathbb{Z} -basis B of the ring R , which is also a \mathbb{Z}_q -basis of R_q and an \mathbb{R} -basis of $K_{\mathbb{R}}$. Then for any $a \in R_q$, multiplication by a corresponds to a matrix $\mathbf{A}_a \in \mathbb{Z}_q^{n \times n}$ with respect to B , i.e., for any $s \in R_q$ having coefficient vector $\mathbf{s} \in \mathbb{Z}_q^n$ w.r.t. B , the coefficient vector of $s \cdot a$ w.r.t. B is $\mathbf{s}^t \mathbf{A}_a$. Moreover, if $a \in R_q$ is uniformly random then so is every column of \mathbf{A}_a (though the columns are maximally dependent).

Given a Ring-LWE sample $(a \in R_q, b = s \cdot a + e \in K_{\mathbb{R}}/qR)$, we can transform it to n LWE samples

$$(\mathbf{A}_a \in \mathbb{Z}_q^{n \times n}, \mathbf{b}^t = \mathbf{s}^t \mathbf{A}_a + \mathbf{e}^t \in (\mathbb{R}/q\mathbb{Z})^n),$$

where $\mathbf{s} \in \mathbb{Z}_q^n$, $\mathbf{e} \in \mathbb{R}^n$ are respectively the coefficient vectors of s, e w.r.t. B . The distribution of \mathbf{e} is $\sigma(B)^{-1} \cdot \sigma(\psi)$, which is “narrow” if ψ itself is narrow and B is chosen appropriately. Note that the columns of \mathbf{A}_a are not independent, nor are the entries of \mathbf{e} in general; if this is a concern then we can throw away all but one sample to get one LWE sample per Ring-LWE sample.

Equivalence of dual and non-dual forms. Here we show that the two syntactic forms of Ring-LWE (dual and non-dual) are equivalent up to the choice of error distribution. First, a consequence the Chinese Remainder Theorem (see [LPR10, Lemma 2.15]) is that for any fractional ideals \mathcal{I}, \mathcal{J} , there exists an efficiently computable $t \in K$ such that multiplication by t induces an efficiently invertible bijection $\theta_t: \mathcal{I}_q \rightarrow \mathcal{J}_q$, and also a function $\kappa_t: K_{\mathbb{R}}/q\mathcal{I} \rightarrow K_{\mathbb{R}}/q\mathcal{J}$ that maps the uniform distribution to the uniform distribution.³ For example, in many cases of interest $\mathcal{I} = R^\vee$ is *principal*, so for $\mathcal{J} = R$ we can let t be the inverse of any generator, so that $tR^\vee = R$.⁴

Using t and the associated functions θ_t, κ_t , we can transform one form of Ring-LWE to another. For concreteness we focus on converting the dual form to the non-dual form, i.e., $\mathcal{I} = R^\vee$ and $\mathcal{J} = R$, but the technique works for any \mathcal{I}, \mathcal{J} . Given dual-form Ring-LWE samples $(a_i, b_i = s \cdot a_i + e_i) \in R_q \times K_{\mathbb{R}}/qR^\vee$ for secret $s \in R_q^\vee$, where each $a_i \in R_q$ is uniform and each $e_i \leftarrow \psi$, we simply multiply each b_i by t (and reduce modulo qR if necessary) to obtain

$$b'_i := \kappa_t(b_i) = t \cdot b_i = s' \cdot a_i + e'_i \bmod qR,$$

where $s' = t \cdot s = \theta_t(s) \in R_q$ and the error term $e'_i = t \cdot e_i \in K_{\mathbb{R}}$ has “tweaked” distribution $t \cdot \psi$.

Clearly, the $(a_i, b'_i) \in R_q \times K_{\mathbb{R}}/qR$ are properly generated Ring-LWE samples for error distribution $t \cdot \psi$ with uniformly random secret $s' \in R_q$ (because $s \in R_q^\vee$ is uniform and θ_t is a bijection), and finding s'

³The “efficiency computable” part of the claim assumes that the factorization of q is known, which is typically the case.

⁴We stress that \mathcal{I} and \mathcal{J} need not be principal to obtain the desired functions, it just makes their existence easier to see.

immediately yields s (because θ_t is efficiently invertible). Therefore, search for the non-dual form with error $t \cdot \psi$ is at least as hard as for the dual form with error ψ . Moreover, because κ_t sends uniform input to uniform output, the same goes for the decision problems.

Lastly, we point out that the tweaked error distribution $t \cdot \psi$ is essentially a scaling of ψ by $\sigma_i(t)$ in the i th coordinate of the canonical embedding. Therefore, if ψ is a spherical Gaussian, the tweaked error is an elliptical Gaussian which may have different widths in each of the coordinates of the canonical embedding.

2.4 Tensor Products

Some of our analysis in Sections 4.3 and 4.4 makes use of *tensor products*. For vectors (ordered tuples) \mathbf{a}, \mathbf{b} over a common domain, their tensor product $\mathbf{a} \otimes \mathbf{b}$ is the vector of all $a_i \cdot b_j$, arranged in blocks $a_i \cdot \mathbf{b}$ in the same order as the entries a_i of \mathbf{a} . Similarly, for matrices \mathbf{A}, \mathbf{B} their tensor product $\mathbf{A} \otimes \mathbf{B}$ is the matrix made up of all blocks $a_{i,j} \mathbf{B}$, arranged analogously. Many linear-algebraic operations are multiplicative under the tensor product, e.g., the Euclidean norm $\|\mathbf{a} \otimes \mathbf{b}\| = \|\mathbf{a}\| \cdot \|\mathbf{b}\|$ and the dual $(\mathbf{A} \otimes \mathbf{B})^\vee = \mathbf{A}^\vee \otimes \mathbf{B}^\vee$.

For two vector spaces V, W over the rationals \mathbb{Q} , their tensor product $V \otimes W$ is the vector space over \mathbb{Q} consisting of all \mathbb{Q} -linear combinations of *pure tensors* $v \otimes w$ for $v \in V, w \in W$, which satisfy the relations

$$\begin{aligned} c(v \otimes w) &= (cv) \otimes w = v \otimes (cw) \\ (v_1 + v_2) \otimes w &= (v_1 \otimes w) + (v_2 \otimes w) \\ v \otimes (w_1 + w_2) &= (v \otimes w_1) + (v \otimes w_2). \end{aligned}$$

If V, W respectively have \mathbb{Q} -bases $(v_i), (w_j)$, then $(v_i \otimes w_j)$ is a \mathbb{Q} -basis of $V \otimes W$. If V, W are rings then $V \otimes W$ is a ring with multiplication defined via the *mixed-product property*

$$(v_1 \otimes w_1) \cdot (v_2 \otimes w_2) = (v_1 v_2) \otimes (w_1 w_2).$$

If V, W are fields, then in general $V \otimes W$ may not be a field, though it will be one in all cases we examine.

The tensor product of two \mathbb{Z} -modules (i.e., additive groups) is defined analogously, with a ring structure defined by the mixed-product property if the original modules are themselves rings.

3 Attack Framework

In this section we give a new exposition of the Ring-LWE attacks described in [EHL14, ELOS15, CLS15, CIV16, CLS16], focusing on the essential geometric reasons why they work. All the attacks fall into one of two classes: reduction to *errorless* LWE, for which search is trivially solvable; and reduction modulo an *ideal divisor* \mathfrak{q} of the modulus qR , for which decision can be solved under certain conditions on \mathfrak{q} and the error distribution. In this section we describe a simple, unified framework that encompasses both classes of attack. Then in Section 4 we show how certain concrete instantiations are vulnerable to the attacks, and in Section 5 we show that worst-case-hard instantiations are provably immune to them.

Following the above-cited works, throughout this section we restrict our attention to the so-called “non-dual” form of Ring-LWE, which involves spherical Gaussian error relative to R (in the canonical embedding). We mainly work with *continuous* rather than discrete error, which more clearly exposes the essential ideas without the extra complication of discretization. We contend that a successful attack against an instantiation with continuous Gaussian error should be enough to reject the corresponding discrete version, since we should not rely on (nor expect) discretization itself to provide any significant security. Indeed, we shall see that the attacks described here frequently work in the discrete setting as well.

The rest of the section is organized as follows. As a warm-up, in Section 3.1 we start by describing some simple attacks on plain LWE, then in Section 3.2 we show how they naturally extend to Ring-LWE. In a nutshell, the attacks exploit the existence of one or more sufficiently short elements in the *dual ideal* of some small-norm ideal divisor of qR . Then in Section 3.3, for completeness we describe the effect of discretization on the attacks, which is typically minor (or nothing at all).

3.1 Warm-Up: Attacking Plain LWE

To start, we describe some folklore attacks against plain LWE and the conditions required for them to work. (Similar ideas have appeared in search-to-decision reductions for LWE in, e.g., [Pei09, ACPS09, MP12].) Suppose that $q' \geq 1$ is a divisor of the LWE modulus q , and let ψ be a (continuous) error distribution over \mathbb{R} . Then given LWE samples

$$(\mathbf{a}_i, b_i = \langle \mathbf{s}, \mathbf{a}_i \rangle + e_i \bmod q) \in \mathbb{Z}_q^n \times \mathbb{R}/q\mathbb{Z}$$

where $\mathbf{a}_i \leftarrow \mathbb{Z}_q^n$ and $e_i \leftarrow \psi$, we can reduce them modulo q' to obtain samples

$$(\mathbf{a}'_i = \mathbf{a}_i \bmod q', b'_i = b_i \bmod q') \in \mathbb{Z}_{q'}^n \times \mathbb{R}/q'\mathbb{Z}.$$

Notice that $b'_i = \langle \mathbf{s}', \mathbf{a}'_i \rangle + e_i \bmod q'$ where $\mathbf{s}' = \mathbf{s} \bmod q'$, so these are LWE samples with the same error distribution ψ , but now the secret now lies in a space of size $(q')^n \leq q^n$. For example, in the extreme case where $q' = 1$, all the $\langle \mathbf{s}', \mathbf{a}'_i \rangle = 0 \bmod \mathbb{Z}$, so $b'_i = e_i \bmod \mathbb{Z}$. Finally, observe that modular reduction transforms uniformly random samples modulo q to uniformly random ones modulo q' .

The above observations easily translate to potential attacks. For concreteness, let $q' = 1$.

- If ψ is *detectably non-uniform* modulo \mathbb{Z} , we immediately have a distinguishing attack: simply test whether the $b'_i \in \mathbb{R}/\mathbb{Z}$ are non-uniform.
- Alternatively, if ψ *usually does not “wrap around”* modulo \mathbb{Z} —i.e., if $\Pr_{e \leftarrow \psi}[e \notin [-\frac{1}{2}, \frac{1}{2}]]$ is small enough—then we immediately have an attack on search: simply recover the unreduced errors $e_i \in [-\frac{1}{2}, \frac{1}{2}]$ as the distinguished representatives of the $b'_i = e_i \bmod \mathbb{Z}$, then subtract these errors from the original values of $b_i \in \mathbb{R}/q\mathbb{Z}$ to obtain errorless LWE samples, which can be solved by linear algebra.

More generally, the same ideas apply for $q' > 1$, where now we care about the distribution of the $b'_i - \langle \mathbf{s}', \mathbf{a}'_i \rangle = e_i \bmod q'\mathbb{Z}$. However, we need to account for the exponentially many $(q')^n$ possible values of \mathbf{s}' , which we do not know how to handle efficiently in general.

On the positive side, suppose $\psi = D_r$ is a Gaussian distribution with parameter $r \geq \eta_\varepsilon(\mathbb{Z})$ exceeding the smoothing parameter of \mathbb{Z} for some very small ε , e.g., $r > \sqrt{n} \geq \eta_{2^{-n}}(\mathbb{Z})$ as in the worst-case hardness theorems for LWE. Such errors “wrap around” modulo \mathbb{Z} , so we cannot reliably recover them from their residues in the search attack. Moreover, the reduced samples $b'_i = e_i \bmod \mathbb{Z}$ are statistically close to uniform, and are therefore useless in the distinguishing attack. (More precisely, the distinguishing advantage is at most $m \cdot \varepsilon$, where m is the number of samples consumed.)

3.2 Attacking Ring-LWE

The authors of [EHL14, ELOS15, CLS15, CIV16, CLS16] describe analogous attacks on Ring-LWE that can yield rather small search spaces for the reduced secret, even for nontrivial target moduli. (The approaches are closely related to the search-to-decision reduction for Ring-LWE from [LPR10].) As we shall see, the attacks are all instances of the following framework.

Let $\mathfrak{q} \subseteq R$ be an ideal divisor of qR , having norm $N(\mathfrak{q}) := |R/\mathfrak{q}|$, and let ψ be a continuous error distribution over $K_{\mathbb{R}}$. Given Ring-LWE samples

$$(a_i, b_i = s \cdot a_i + e_i) \in R_{\mathfrak{q}} \times K_{\mathbb{R}}/qR$$

where $a_i \leftarrow R_{\mathfrak{q}}$ and $e_i \leftarrow \psi$, we can reduce them modulo \mathfrak{q} to obtain samples

$$(a'_i = a_i \bmod \mathfrak{q}, b'_i = b_i \bmod \mathfrak{q}) \in R/\mathfrak{q} \times K_{\mathbb{R}}/\mathfrak{q}.$$

As above, we have $b'_i = s' \cdot a'_i + e_i \bmod \mathfrak{q}$ where $s' = s \bmod \mathfrak{q}$, so these are Ring-LWE samples with error distribution ψ , but now the secret lies in a space of size $N(\mathfrak{q})$. Also observe that reduction modulo \mathfrak{q} maps uniform samples to uniform samples.

When $N(\mathfrak{q})$ is not too large, the preceding observations potentially yield attacks:

- If $\psi \bmod \mathfrak{q}$ is detectably non-uniform, then we immediately have an attack against the decision problem: try all candidates $\hat{s} \in R/\mathfrak{q}$ for s' , and for each one test whether the $b'_i - \hat{s} \cdot a'_i \in K_{\mathbb{R}}/\mathfrak{q}$ are statistically non-uniform; accept if such an \hat{s} exists, otherwise reject. In Section 3.2.1 below we describe a standard method of distinguishing reduced spherical Gaussians $D_r \bmod \mathfrak{q}$ from uniform.
- Similarly, if ψ has one or more coefficients (relative to some fixed \mathbb{Z} -basis of \mathfrak{q}) that usually do not “wrap around” modulo \mathbb{Z} , then we can attack search by reducing to errorless LWE. See Section 3.2.2 below for further details.
- On the positive side, if $\psi = D_r$ is a continuous Gaussian of parameter $r \geq \eta_{\varepsilon}(\mathfrak{q})$ for some very small ε , then neither of the attacks work, because every coefficient of the error “wraps around,” and moreover, the reduced error $\psi \bmod \mathfrak{q}$ is statistically close to uniform. We return to these points in Section 5.

3.2.1 Distinguisher

To run the distinguishing attack, we need a way of efficiently distinguishing $\psi \bmod \mathfrak{q}$ from uniform over $K_{\mathbb{R}}/\mathfrak{q}$, for spherical Gaussian error $\psi = D_r$. A variety of statistical tests have been proposed in [EHL14, ELOS15, CLS15, CLS16], but in this work it suffices to consider a standard method that uses a sufficiently short nonzero element w in the *dual ideal* \mathfrak{q}^{\vee} of \mathfrak{q} , or equivalently, a short nonzero vector $\mathbf{w} = \overline{\sigma(w)}$ in the dual lattice $\mathcal{L}^{\vee} = \sigma(\mathfrak{q})^{\vee}$ of $\mathcal{L} = \sigma(\mathfrak{q})$. More generally, one can use multiple linearly independent dual vectors—e.g., to reduce the number of consumed samples, or to improve the effectiveness of the statistical test—but we do not pursue such optimizations because we will not need them (see the discussion below).

Lemma 3.1. *Let \mathcal{L} be any lattice, $\mathbf{w} \in \mathcal{L}^{\vee} \setminus \{\mathbf{0}\}$ be any nonzero element of its dual lattice, and $r > 0$. Then for $\mathbf{x} \leftarrow D_r \bmod \mathcal{L}$, the distribution of $\langle \mathbf{w}, \mathbf{x} \rangle \bmod \mathbb{Z}$ is $D_{r\|\mathbf{w}\|} \bmod \mathbb{Z}$, and*

$$\mathbb{E}_{\mathbf{x}}[\cos(2\pi\langle \mathbf{w}, \mathbf{x} \rangle)] = \exp(-\pi(r\|\mathbf{w}\|)^2).$$

In particular, if $r\|\mathbf{w}\| = O(1)$, then the expectation is $\Omega(1)$.

Proof. Because $\mathbf{w} \in \mathcal{L}^{\vee}$ we have $\langle \mathbf{w}, \mathcal{L} \rangle \subseteq \mathbb{Z}$, so the distribution of $\langle \mathbf{w}, D_r \bmod \mathcal{L} \rangle \bmod \mathbb{Z}$ is $\langle \mathbf{w}, D_r \rangle \bmod \mathbb{Z} = D_s \bmod \mathbb{Z}$, where $s = r\|\mathbf{w}\|$. The expectation $\mathbb{E}_{\mathbf{x} \leftarrow D_s}[\cos(2\pi x)]$ is merely the Fourier coefficient at 1 of $D_s \bmod \mathbb{Z}$, which by a routine calculation is $\exp(-\pi s^2)$. \square

It is easy to see that for uniformly random \mathbf{x} modulo \mathcal{L} , the inner product $\langle \mathbf{w}, \mathbf{x} \rangle \bmod \mathbb{Z}$ is uniform, so $\mathbb{E}[\cos(2\pi\langle \mathbf{w}, \mathbf{x} \rangle)] = 0$. With Lemma 3.1, this immediately yields an efficient distinguisher between $D_r \bmod \mathcal{L}$ and uniform when $r\|\mathbf{w}\| = O(1)$: given many samples \mathbf{x}_i , compute the average of $\cos(2\pi\langle \mathbf{w}, \mathbf{x}_i \rangle)$ and accept if it exceeds an appropriate threshold $t = \Omega(1)$. (See, e.g., [Reg05, Lemma 3.6] for more details.)

Relation to prior distinguishing attacks. The distinguishing attacks from [EHL14, ELOS15, CLS15, CLS16] are not described using dual ideals and the trace pairing, but as we now show, they all can be expressed in these terms. Letting the modulus q be prime, the prior attacks use a prime ideal divisor \mathfrak{q} of qR having norm $N(\mathfrak{q}) = q^d$ for some small positive integer d , and a ring homomorphism (actually, an isomorphism) $h: R/\mathfrak{q} \rightarrow \mathbb{F}_{q^d}$. For a discrete error distribution ψ over R , the attacks use some statistical test of non-uniformity for $h(\psi \bmod \mathfrak{q}) \in \mathbb{F}_{q^d}$ to solve decision as described above.

Any such attack can be reformulated in terms of duality and trace pairings, as follows. Recall that \mathbb{F}_{q^d} is a d -dimensional vector space over $\mathbb{F}_q = \mathbb{Z}_q$, and fix some arbitrary \mathbb{Z}_q -basis. Observe that the d induced functions $h_j: R/\mathfrak{q} \rightarrow \mathbb{Z}_q$, where $h_j(r)$ is defined as the j th coefficient of $h(r) \in \mathbb{F}_{q^d}$ (with respect to the chosen \mathbb{Z}_q -basis), are additive group homomorphisms. A basic fact of duality is that any such homomorphism can be expressed as $h_j(x) = q \cdot \text{Tr}(w \cdot x)$ for some efficiently computable $w_j \in \mathfrak{q}^\vee$, where the q factor just scales the discrete torus $q^{-1}\mathbb{Z}/\mathbb{Z}$ to yield the range \mathbb{Z}_q . (Indeed, the dual lattice is sometimes *defined* as the group of such homomorphisms.) Moreover, because the domain R/\mathfrak{q} of h_j is discrete, the w_j need only be specified modulo R^\vee , i.e., we can take $w_j \in \mathfrak{q}^\vee/R^\vee$.

By the above, any ring homomorphism $h: R/\mathfrak{q} \rightarrow \mathbb{F}_{q^d}$ corresponds to taking the trace pairing with some fixed d -tuple of dual elements $w_j \in \mathfrak{q}^\vee$. However, not every tuple of dual elements (even \mathbb{Z}_q -linearly independent ones) corresponds to such a ring homomorphism, so the present framework is strictly more general. Moreover, as we will see in Section 4, all the previously considered attacks that used $d > 1$ can be made to work using just *one* appropriately chosen $w \in \mathfrak{q}^\vee$, so we do not need the extra generality of multiple dual elements.

3.2.2 Search Attack

We now describe the details of the attack on search, focusing on the case $\mathfrak{q} = R$ for simplicity.⁵ This generalizes the attack from [CIV16], and gives a simpler analysis with the same ultimate results (see below for a comparison, and Section 4 for concrete examples). Let $B = (b_j)_j$ be a fixed \mathbb{Z} -basis of R for which one or more coefficients of ψ do not wrap around, i.e.,

$$\Pr_{e \leftarrow \psi} [e_j \notin [-\frac{1}{2}, \frac{1}{2})] \approx 0 \quad (3.1)$$

for some j , where $e = \sum_j e_j \cdot b_j$ for $e_j \in \mathbb{R}$ is the unique representation of e with respect to B . Lemma 3.2 below shows that for spherical Gaussian error, this condition holds for the index j of any sufficiently short element of the *dual basis* of B .

To perform the attack, as described in Section 2.3.2 we transform each Ring-LWE sample ($a \in R_q, b = s \cdot a + e \in K_{\mathbb{R}}/qR$) to n LWE samples

$$(\mathbf{A}_a \in \mathbb{Z}_q^{n \times n}, \mathbf{b}^t = \mathbf{s}^t \mathbf{A}_a + \mathbf{e}^t \in (\mathbb{R}/q\mathbb{Z})^n),$$

where \mathbf{A}_a denotes the matrix of multiplication by a (whose columns are uniformly random but maximally dependent), and \mathbf{s}, \mathbf{e} are respectively the coefficient vectors of s, e (all with respect to basis B). Now, for each index j for which Equation (3.1) holds, we can with high probability obtain $e_j \in [-\frac{1}{2}, \frac{1}{2})$ as the distinguished representative of the j th entry of $\mathbf{b} \bmod \mathbb{Z}$. This yields an errorless LWE sample for the j th column of \mathbf{A}_a . Given enough errorless samples, we can solve for \mathbf{s} by standard linear algebra.

⁵The attack easily generalizes to arbitrary ideal divisors $\mathfrak{q}|qR$ of not-too-large norm; we omit the details, because the present form will be enough for our purposes.

Lemma 3.2. Let $D = (d_j)$ be the dual basis of $B = (b_j)$, i.e., $\mathbf{D} = \mathbf{B}^\vee := (\mathbf{B}^{-1})^*$ where $\mathbf{B} = \sigma(B)$, $\mathbf{D} = \sigma(D)$. Then for any $r, \varepsilon > 0$, if

$$\|d_j\| \leq \left(2r \sqrt{\log(2/\varepsilon)/\pi}\right)^{-1}$$

then $\Pr_{e \leftarrow D_r}[e_j \in [-\frac{1}{2}, \frac{1}{2}]] \geq 1 - \varepsilon$, where $e = \sum_j e_j \cdot b_j$ for $e_j \in \mathbb{R}$.

Proof. By definition, the distribution of \mathbf{e} , the coefficient vector of e with respect to B , is $\mathbf{B}^{-1} \cdot D_r$, where D_r is the spherical Gaussian of parameter r over H . Therefore, e_j is distributed as a Gaussian of parameter $s = r \cdot \|d_j\| \leq (2\sqrt{\log(2/\varepsilon)/\pi})^{-1}$. The claim then follows directly by the standard Gaussian tail bound $\Pr_{x \leftarrow D_s}[|x| \geq t] \leq 2 \exp(-\pi(t/s)^2)$ for any $t \geq 0$. \square

Comparison with [CIV16]. The authors of [CIV16] also attack search using a reduction to errorless LWE, but use a different approach for showing that error coefficients are zero. In brief, they consider the matrix $\mathbf{B} = \sigma(B)$ of the linear transformation that maps from a basis B (e.g., the power basis $B = (1, X, \dots, X^{n-1})$) to the canonical embedding. Using its singular value decomposition, they analyze the “skewness” of the transformation via its singular values, and the “alignment” of the basis elements with the singular vectors, to show that certain error coefficients are usually small. By contrast, the approach described above only needs to analyze the lengths of the dual vectors, i.e., the rows of \mathbf{B}^{-1} .

3.3 Discretization

We contend that a successful attack against a Ring-LWE instantiation with *continuous* Gaussian error should already be enough to reject the corresponding *discrete* instantiation, because we should not expect discretization itself to provide any significant security. Indeed, the above-described attacks can frequently be made to work in the discrete setting as well, as we now describe. Here the attacker is given samples $(a, b = s \cdot a + \bar{e}) \in R_q \times R_q$, where $\bar{e} \in R$ is the discretization of some $e \in K_{\mathbb{R}}$ drawn from ψ . A standard discretization method, which in particular is used in [ELOS15, CLS15, CLS16], writes $e = \sum_j e_j \cdot b_j$ with respect to some fixed \mathbb{Z} -basis $B = (b_j)$ of R , rounds each e_j to the nearest integer $\bar{e}_j = \lfloor e_j \rfloor \in \mathbb{Z}$, and outputs $\bar{e} = \sum_j \bar{e}_j \cdot b_j \in R$. Such discretization has the following effect on the above attacks:

- For the search attack via reduction to errorless LWE, if the discretization basis B is also the basis used for the reduction (which is indeed the case for prior instantiations), then discretization has no effect at all: the small real error coefficients $e_j \in [-\frac{1}{2}, \frac{1}{2})$ are simply rounded off to zero before being given to the attacker, instead of by the attacker.
- For the distinguishing attack that uses a short $w \in \mathfrak{q}^\vee$, discretizing $x \in K_{\mathbb{R}}/\mathfrak{q}$ to $\bar{x} \in R/\mathfrak{q}$ has the effect of making $\text{Tr}(w \cdot \bar{x}) \approx \text{Tr}(w \cdot x) \in \mathbb{R}/\mathbb{Z}$ lie in $\text{Tr}(wR)/\mathbb{Z} \subseteq (N(\mathfrak{q})^{-1}\mathbb{Z})/\mathbb{Z}$, because

$$w \in \mathfrak{q}^\vee = \mathfrak{q}^{-1}R^\vee \subseteq N(\mathfrak{q})^{-1}R^\vee.$$

So if $w \in R^\vee$ (which in particular holds if $\mathfrak{q} = R$) then the attack is trivially nullified, because $\text{Tr}(w \cdot \bar{x}) = 0 \in \mathbb{Z}/\mathbb{Z}$ always. However, for $N(\mathfrak{q}) \gg 1$ and $w \in \mathfrak{q}^\vee \setminus R^\vee$, the distinguishing attack often still works, because the elements b_j of the discretization basis B typically have small trace pairings $\text{Tr}(w \cdot b_j)$, like 0 or $\pm N(\mathfrak{q})^{-1}$. Therefore, discretization does little to “smooth out” the non-uniform distribution of $\text{Tr}(w \cdot x) \in \mathbb{R}/\mathbb{Z}$.

4 Insecure Instantiations

In this section we show how the attack framework from Section 3 applies to the concrete insecure Ring-LWE instantiations defined in [EHL14, ELOS15, CLS15, CLS16] (among others). In all cases, the core reason for the insecurity is that the error distributions are insufficiently “well spread” relative to the rings, viewed as lattices. (See, e.g., Figure 3.) To prove this formally, it suffices by Lemmas 3.1 and 3.2 to demonstrate sufficiently short nonzero elements in the dual ideal \mathfrak{q}^\vee of some ideal divisor \mathfrak{q} of qR (possibly $\mathfrak{q} = R$ itself) whose algebraic norm $N(\mathfrak{q})$ is not too large.

We stress that all these insecure instantiations—excepting [EHL14], for which the following conclusions still apply—are for the “non-dual” version of Ring-LWE with spherical Gaussian errors relative to R (in the canonical embedding). By contrast, the definition of Ring-LWE from [LPR10], and the instantiations having worst-case hardness, involve spherical errors relative to the *dual ideal* R^\vee (see Section 2.3.2). When the insecure and hard instantiations are transformed to be directly comparable, the resulting error distributions turn out to have very different widths and shapes. We return to this point in Section 5, where we show that the hard instantiations are immune to the attacks from Section 3.

4.1 Rings $\mathbb{Z}[X]/(X^n + aX + b)$

The instantiations defined in [ELOS15] involve rings of the form $R = \mathbb{Z}[X]/(X^n + aX + b)$ for some nonnegative integers a, b , and spherical Gaussian error in the canonical embedding. The original attacks on these instantiations solved the *decision* problem for certain moduli q via (essentially) the distinguishing attack from Section 3.2.1, using 20 samples. Later work by [CIV16] successfully attacked *search* for any modulus q by reduction to errorless LWE, obtaining a 100 % success probability using only 4–7 samples (depending on the instantiation). The analysis of [CIV16] relied on the singular values of the embedded power basis elements, and the alignment of the singular vectors with those embedded powers. Here we use our simpler analysis in terms of the lengths of dual vectors, following the approach described in Section 3.2.2, to obtain the same conclusions; see Figure 1.⁶

$f(X)$	r_0	r	threshold	num. short \mathbf{d}_j	samples
$X^{192} + 4092$	8.87	5,440	6.32×10^{-5}	29 (15.1 %)	7
$X^{256} + 8190$	8.35	8,390	4.01×10^{-5}	47 (18.3 %)	6
$X^{128} + 524288X + 524285$	8.00	45,540	7.79×10^{-6}	32 (25.0 %)	4

Figure 1: Analysis of the instantiations from [ELOS15], for the power basis $B = (1, X, \dots, X^{n-1})$ and its dual $D = (\mathbf{d}_j)$. “Threshold” denotes the value $(2r\sqrt{\log(4n)/\pi})^{-1}$ from Lemma 3.2 (for $\varepsilon = 1/(2n)$) for the lengths of the dual vectors \mathbf{d}_j , below which the j th error coefficient is zero with probability at least $1 - \varepsilon$. “Num. short \mathbf{d}_j ” denotes the number (and percentage, out of n) of the dual vectors whose norms are below the threshold. “Samples” denotes the number of Ring-LWE samples that suffice to recover the secret via linear algebra on the errorless coefficients.

⁶A preliminary version of this work incorrectly concluded that for each instantiation, more than 90 % of the coordinates are errorless; this was due to a misinterpretation of the parameter w from [ELOS15, Section 9]. We thank an anonymous reviewer for pointing this out.

4.2 Prime Cyclotomics

Let the modulus q be a prime integer and let $R = \mathbb{Z}[\zeta_q]$ be the q th cyclotomic ring, where ζ_q denotes a primitive q th root of unity. It is well known (and easy to verify) that $qR = \mathfrak{q}^{q-1}$ and $qR^\vee = \mathfrak{q}$, where the ideal $\mathfrak{q} = (1 - \zeta_q)R + qR$ is prime and has norm $N(\mathfrak{q}) = q$.

In [CLS15, Section 6], the authors use (essentially) the approach from Section 3.2.1 to obtain distinguishing attacks that work in practice for the cases $q = 251, 503, 809$, using \mathfrak{q} as the ideal divisor of qR . Their experiments work for parameters

$$r \leq 1.53 \cdot \delta_R = 1.53 \sqrt{q^{(q-2)/(q-1)}} < 1.53 \sqrt{q}.$$

(Note that this corresponds to a volume-normalized parameter of $r_0 \leq 1.53$, which is considered quite small for LWE errors.) We remark that these distinguishing attacks are not known to translate to search, because no search-decision equivalence is known for this choice of parameters.

Our analysis. The following lemma formally proves why the experiments work, and additionally implies that search can be solved via errorless LWE for slightly smaller parameters. See Figure 2 for a graphical depiction in the third cyclotomic ring.

Lemma 4.1. *Let q , R , and \mathfrak{q} be as above. Then $q^{-1} \cdot (1, \zeta_q, \dots, \zeta_q^{q-2})$ is a \mathbb{Z} -basis of \mathfrak{q}^\vee , all of whose elements have length $\sqrt{q-1}/q$.*

Two immediate corollaries are that by Lemma 3.2, we can solve search by reducing to errorless LWE for, say, $r \leq \sqrt{q \cdot \pi / (4 \log(4q))} = \Theta(\sqrt{q / \log q})$; and by Lemma 3.1 and the associated distinguisher, we can efficiently solve decision for any $r = O(\sqrt{q})$.

Proof. Because $qR^\vee = \mathfrak{q}$, the dual ideal of \mathfrak{q} is $\mathfrak{q}^\vee = \mathfrak{q}^{-1}R^\vee = q^{-1}R$, for which $q^{-1} \cdot (1, \zeta_q, \dots, \zeta_q^{q-2})$ is a \mathbb{Z} -basis. Because every complex embedding of ζ_q^j is a root of unity, we have $\|q^{-1} \cdot \zeta_q^j\| = \sqrt{q-1}/q$. \square

4.3 Quadratic Extensions of Cyclotomics

In [CLS16], the authors consider non-dual Ring-LWE instantiations for certain quadratic extensions of cyclotomics, namely, $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$ where ζ_p denotes a primitive p th root of unity for an odd prime p , and $d > 1$ is a square-free integer that is coprime to p , and is 3 modulo 4. They prove that for appropriate moduli, and for spherical Gaussian error of parameter $r \approx \sqrt{d}$, which corresponds to a volume-normalized parameter of $r_0 = r/\delta_R \approx d^{1/4}/\sqrt{p}$, one can efficiently solve search by combining a distinguishing attack with known search-decision equivalences for Galois rings. In addition, their distinguishing attacks work in practice up to larger parameters $r \approx \sqrt{p \cdot d}$ (corresponding to $r_0 \approx d^{1/4}$), though no formal analysis was provided to explain why.

Our analysis. Here we prove that for the same class of rings, and for $r \approx \sqrt{p \cdot d / \log p}$ (i.e., $r_0 \approx d^{1/4}/\sqrt{\log p}$), we can solve search directly by reducing to errorless LWE, using the approach from Section 3.2.2. (As above, this works for any choice of modulus q .) Moreover, for any $r = O(\sqrt{p \cdot d})$ we can efficiently solve decision, and hence search, using the distinguishing attack from Section 3.2.1.

The basic reason why the attacks work on these instantiations is quite simple: $\mathbb{Z}[\sqrt{d}]$ has root discriminant $\approx d^{1/4}$, but its dual lattice has a very short vector of length $\approx 1/\sqrt{d}$. This means that error of parameter

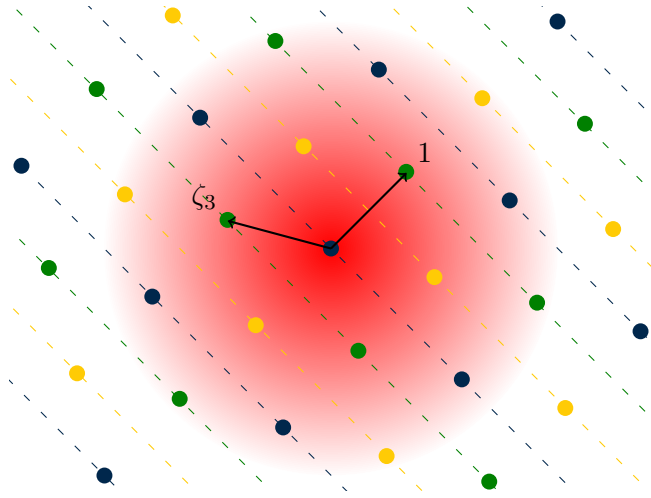


Figure 2: The canonical embedding of the third cyclotomic ring $R = \mathbb{Z}[\zeta_q]$ for $q = 3$, along with a continuous Gaussian of parameter $r = q/\sqrt{q-1}$. (We have depicted the two-dimensional real inner product space $H \subset \mathbb{C}^2$ as \mathbb{R}^2 .) The dashed colored lines show the hyperplanes that partition the lattice points according to their inner products with (the canonical embedding of) $q^{-1} \in \mathfrak{q}^\vee$; this also partitions according to the three cosets of $\mathfrak{q} = qR^\vee$. The zero coset of \mathfrak{q} (in blue) has noticeably more probability mass under the Gaussian than the ± 1 cosets (in maize and green).

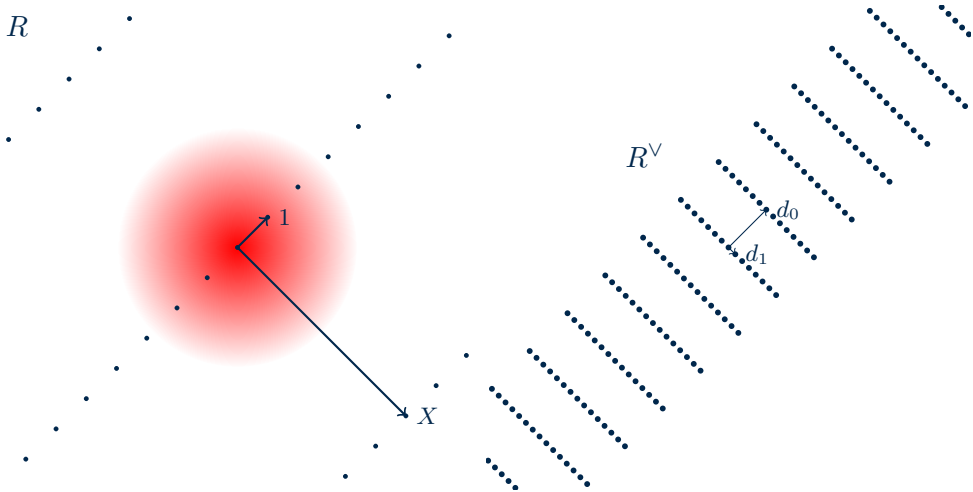


Figure 3: On the left: the canonical embedding $\mathcal{L} = \sigma(R)$ of $R = \mathbb{Z}[\sqrt{d}]$ for $d = 31$, along with a continuous spherical Gaussian distribution of parameter $r = \sqrt{d}/2$, which corresponds to a volume normalized parameter of $r_0 := r/\det(\mathcal{L})^{1/2} = d^{1/4}/(2\sqrt{2})$. Observe that discretizing an error term to R using the power basis $P = (1, X)$ usually results in a coefficient of zero for X . On the right: the dual lattice \mathcal{L}^\vee (corresponding to R^\vee), along with the dual basis $D = (d_0, d_1)$ of the power basis. Observe that d_1 is very short, which corresponds to the wide gap between integers multiples of X .

$r \approx \sqrt{d}$ (i.e., $r_0 \approx d^{1/4}$) is still so narrow relative to $\mathbb{Z}[\sqrt{d}]$ that discretizing yields a zero coefficient; see Figure 3. The same goes for the compositum ring $\mathbb{Z}[\zeta_p, \sqrt{d}] \cong \mathbb{Z}[\zeta_p] \otimes \mathbb{Z}[\sqrt{d}]$, because $\mathbb{Z}[\zeta_p]$ has many dual elements whose lengths are essentially the inverse of the root discriminant.

Lemma 4.2. *For p and d as described above, let $B = (1, \zeta_p, \dots, \zeta_p^{p-2}) \otimes (1, \sqrt{d})$, which is a \mathbb{Z} -basis of $R = \mathbb{Z}[\zeta_p, \sqrt{d}] \cong \mathbb{Z}[\zeta_p] \otimes \mathbb{Z}[\sqrt{d}]$. Then the dual basis $D = B^\vee$ has $p - 1$ elements of length $1/\sqrt{pd}$.*

An immediate corollary is that by Lemma 3.2, we can solve search via errorless LWE for, say, $r = \sqrt{p \cdot d \cdot \pi / (4 \log(8p))}$. Because R has root discriminant

$$\delta_R = \delta_{\mathbb{Z}[\zeta_p]} \cdot \delta_{\mathbb{Z}[\sqrt{d}]} = \sqrt{p^{(p-2)/(p-1)} \cdot (4d)^{1/2}} \leq \sqrt{p} \cdot (4d)^{1/4},$$

this corresponds to a volume-normalized parameter $r_0 \geq (d \cdot \pi^2 / 64)^{1/4} / \sqrt{\log(8p)} = \Theta(d^{1/4} / \sqrt{\log p})$. Another corollary is that by Lemma 3.1, we can solve decision for any $r = O(\sqrt{p \cdot d})$, which corresponds to $r_0 = O(d^{1/4})$.

Proof of Lemma 4.2. Let σ_p and σ_d respectively denote the canonical embeddings for $\mathbb{Z}[\zeta_p]$ and $\mathbb{Z}[\sqrt{d}]$, and let $B_p = (1, \zeta_p, \dots, \zeta_p^{p-2})$, $B_d = (1, \sqrt{d})$ and $D_p = B_p^\vee$, $D_d = B_d^\vee$. Then $D = D_p \otimes D_d$, and

$$\mathbf{D} := \sigma(D) = \sigma_p(D_p) \otimes \sigma_d(D_d) = \sigma_p(B_p)^\vee \otimes \sigma_d(B_d)^\vee.$$

We analyze the two components separately. First, $\sigma_d(B_d) = \begin{pmatrix} 1 & \sqrt{d} \\ 1 & -\sqrt{d} \end{pmatrix}$, so $\sigma_d(D_d) = \frac{1}{2} \begin{pmatrix} 1 & 1/\sqrt{d} \\ 1 & -1/\sqrt{d} \end{pmatrix}$; note that the second column has norm $1/\sqrt{2d}$.

Next, we show that all the columns of $\mathbf{D}_p = \sigma_p(D_p)$ have norm $\sqrt{2/p}$. This is equivalent to showing that all the diagonal entries of $\mathbf{D}_p^* \cdot \mathbf{D}_p$ are $2/p$. Indeed, letting $\mathbf{B}_p = \sigma_p(B_p)$ and noting that it is just the lower-left $(p - 1)$ -dimensional submatrix of the (non-normalized) p -dimensional Fourier matrix, whose columns are orthogonal with norm \sqrt{p} and whose top row is the all-ones vector, we have

$$\mathbf{D}_p^* \cdot \mathbf{D}_p = (\mathbf{B}_p^* \cdot \mathbf{B}_p)^{-1} = (p\mathbf{I}_{p-1} - \mathbf{1})^{-1} = p^{-1}\mathbf{X},$$

where $\mathbf{1}$ is the all-1s matrix and \mathbf{X} has 2s along the diagonal and 1s in every other entry.

The claim then follows from the fact that the columns of \mathbf{D} include the columns of $\mathbf{B}_p^\vee \otimes \frac{1}{2} \begin{pmatrix} 1/\sqrt{d} \\ -1/\sqrt{d} \end{pmatrix}$, which all have norm $1/\sqrt{pd}$. \square

4.4 Subfields of Cyclotomics

In [CLS15, Section 5], the authors consider non-dual Ring-LWE instantiations involving subfields K of cyclotomic fields $L = \mathbb{Q}(\zeta_m)$, namely, those that are fixed pointwise by the automorphisms in some subgroup of the Galois group of L/\mathbb{Q} . Letting $R = \mathcal{O}_K$ be the ring of integers in K , the instantiations involve spherical Gaussian error with volume-normalized parameter $r/\delta_R = r_0 = \sigma_0 \sqrt{2\pi} < 3.14$ (which is considered somewhat small for LWE errors). The authors' distinguishing attacks work in practice, and they provide some heuristics as potential explanations, but no formal analysis.⁷

⁷We remark that the ring dimensions in these instantiations are all at most 144, which is small enough that search is reasonably easy to solve using standard basis-reduction techniques. Here we restrict our attention to the class of attacks from Section 3.

p	d	r	r_0	$r' = 1.48\sqrt{pd}$	r'_0
31	4,967	148.5	2.38	580.7	9.30
43	4,871	168.1	2.27	677.3	9.14
61	4,643	189.8	2.15	787.6	8.94
83	4,903	222.0	2.12	944.1	8.99
103	4,951	244.4	2.08	1,056.9	8.98
109	4,919	249.6	2.06	1,083.7	8.95
151	100,447	1,296.2	4.26	5,763.9	18.94
181	100,267	1,400.0	4.20	6,304.9	18.89

Figure 4: Instantiations of non-dual Ring-LWE for rings $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$ where: (1) for spherical error of parameter $r = r_0 \cdot \delta_R$, search can be solved by reducing to errorless LWE, and (2) for spherical error of parameter $r' = r'_0 \cdot \delta_R$, decision (and hence search) can be solved efficiently using the distinguishing attack. The constant factor 1.48 is chosen (somewhat arbitrarily) to ensure a Fourier coefficient of at least 10^{-3} in the distinguishing attack.

Our analysis. For the sub-cyclotomic rings R considered in [CLS15, Section 5], it turns out that the dual ideal R^\vee contains many rather short nonzero elements, relative to the root discriminant δ_R . (See Figure 5.) By Lemma 3.1, this implies an efficient distinguishing attack on non-dual Ring-LWE for narrow enough spherical Gaussians, which in particular includes the parameters studied in [CLS15]. We note that the attack works for any choice of the modulus q , at least for continuous error. Fortunately, the analysis easily transfers to the discrete setting, as described below.

The dual ideals R^\vee contain short nonzero elements for nearly the same reason as for the quadratic extensions of cyclotomics studied in the previous subsection. More specifically, we show that representative sub-cyclotomic number fields K from [CLS15] turn out to be quadratic extensions $K = J(\sqrt{d}) \cong J \otimes \mathbb{Q}(\sqrt{d})$ of some subfield $J \subset K$, for some squarefree integer d of relatively large magnitude. The ring of integers is therefore $R = \mathcal{O}_K \cong S \otimes D$, where $S = \mathcal{O}_J$ and $D = \mathbb{Z}[\delta]$ for $\delta = \sqrt{d}$ if $d \not\equiv 1 \pmod{4}$, otherwise $\delta = (1 + \sqrt{d})/2$. In any case, the magnitude of d is large enough that D^\vee contains a very short nonzero element relative to the root discriminant δ_D . (See Figure 3.) This similarly carries over to $R^\vee = S^\vee \otimes D^\vee$, because we have many nonzero elements of S^\vee whose Euclidean norms are roughly the inverse of the root discriminant δ_S .

For completeness, we mention that all our analysis carries over to the setting of discrete error over R , as long as we have an ideal divisor $\mathfrak{q} \subset S$ of qS having sufficiently small norm. Then $(\mathfrak{q}R)^\vee = \mathfrak{q}^\vee \otimes D^\vee$, so for its dual element the distinguisher can use the pure tensor of some $w \in \mathfrak{q}^\vee \setminus S^\vee$ with a shortest nonzero element in D^\vee . The element w does not even need to be short: we can simply try representatives of all the $N(\mathfrak{q}) - 1$ nonzero cosets of \mathfrak{q}^\vee/S^\vee . Because one of these is congruent modulo S^\vee to a shortest nonzero element $v \in \mathfrak{q}^\vee$, and the error is discrete over R , the distinguisher behaves exactly as if it had used v .

4.4.1 Detailed Analysis of Instantiations

In what follows we analyze two representative instantiations from [CLS15, Section 5] in detail. The other instantiations can be handled similarly, but their number fields are much more complicated to write down

ID	m	generators of $H \subset \mathbb{Z}_m^*$	$\deg(K/\mathbb{Q})$	K
1	$3 \cdot 5 \cdot 11 \cdot 17$	1684, 1618	40	$\mathbb{Q}(\zeta_3, \zeta_{11}, \sqrt{5 \cdot 17})$
2	$3 \cdot 5 \cdot 7 \cdot 11 \cdot 13$	12286, 2003, 11936	60	$\mathbb{Q}(\zeta_7, \zeta_{11} + \zeta_{11}^{-1}, \sqrt{-3 \cdot 5 \cdot 13})$

ID	$\delta_K \approx$	$\lambda_1(R^\vee) \leq$	$\lambda_1(R^\vee) \cdot r_0 \cdot \delta_K \approx$
1	11.76	$\sqrt{8/2805} \approx 0.0534$	1.58
2	21.94	$\sqrt{12/15015} \approx 0.0283$	1.55

Figure 5: Example sub-cyclotomic number fields from [CLS15, Section 5]. For the m th cyclotomic field $L = \mathbb{Q}(\zeta_m)$, the subfield K is the fixed field of the Galois subgroup $H \subset \mathbb{Z}_m^*$ defined by the given generators. The second table displays the root discriminant δ_K of K , an upper bound on the minimum distance of the dual ideal R^\vee (where $R = \mathcal{O}_K$), and the resulting parameter of the real Gaussian that needs to be distinguished from uniform modulo \mathbb{Z} using the attack from Lemma 3.1. In all cases, the parameter is small enough to permit efficient distinguishing.

explicitly. (Instead, one can use computer search to find an appropriate dual element, which is sufficient to mount the distinguishing attack.) Figure 5 shows the results of our analysis.

We first recall some standard background about cyclotomic number fields; see, e.g., [Lan94] for proofs. First, for any positive integers $m = m_1 \cdot m_2$ for coprime m_1, m_2 , the m th cyclotomic field $\mathbb{Q}(\zeta_m)$ is isomorphic to the compositum field $\mathbb{Q}(\zeta_{m_1}, \zeta_{m_2}) \cong \mathbb{Q}(\zeta_{m_1}) \otimes \mathbb{Q}(\zeta_{m_2})$. The Galois group of $\mathbb{Q}(\zeta_m)/\mathbb{Q}$ is isomorphic to $\mathbb{Z}_m^* \cong \mathbb{Z}_{m_1}^* \times \mathbb{Z}_{m_2}^*$, where the i th automorphism is defined by $\tau_i(\zeta_m) = \zeta_m^i$ for $i \in \mathbb{Z}_m^*$; equivalently, $\tau_i(\zeta_{m_1}) = \zeta_{m_1}^{i_1}$ and $\tau_i(\zeta_{m_2}) = \zeta_{m_2}^{i_2}$ for $i = (i_1, i_2) \in \mathbb{Z}_{m_1}^* \times \mathbb{Z}_{m_2}^*$. Finally, for any odd prime p , the (unique) quadratic subfield of the p th cyclotomic $\mathbb{Q}(\zeta_p)$ is $\mathbb{Q}(\sqrt{\pm p})$, where the sign is such that $\pm p \equiv 1 \pmod{4}$.

Instantiation #1. This instantiation (ID 1 in Figure 5) uses a subfield K of the cyclotomic field $L = \mathbb{Q}(\zeta_{2805}) \cong \mathbb{Q}(\zeta_3, \zeta_5, \zeta_{11}, \zeta_{17})$, namely, the fixed field of the Galois subgroup H generated by the automorphisms τ, ν respectively indexed by

$$(1, -1, 1, 1), (1, 3, 1, 3) \in \mathbb{Z}_3^* \times \mathbb{Z}_5^* \times \mathbb{Z}_{11}^* \times \mathbb{Z}_{17}^*.$$

The Galois group of L/\mathbb{Q} has order $1280 = 2 \cdot 4 \cdot 10 \cdot 16$, and the subgroup H has order 32, so the degree of K/\mathbb{Q} is 40.

Lemma 4.3. *The field K defined above is isomorphic to $\mathbb{Q}(\zeta_3, \zeta_{11}, \sqrt{5 \cdot 17}) \cong \mathbb{Q}(\zeta_3) \otimes \mathbb{Q}(\zeta_{11}) \otimes \mathbb{Q}(\sqrt{5 \cdot 17})$, and the dual ideal R^\vee of its ring of integers $R = \mathcal{O}_K$ has 20 known \mathbb{Q} -linearly independent elements of Euclidean norm $\sqrt{8/(3 \cdot 5 \cdot 11 \cdot 17)} \approx 0.0534$.*

It follows that the root discriminant of K is $\delta_K = 3^{1/4} \cdot 11^{9/20} \cdot 85^{1/4} \approx 11.76$. For $r = \sqrt{2\pi} \cdot \delta_K \approx 29.47$ as in [CLS15], using the above elements of R^\vee in the distinguishing attack of Lemma 3.1 results in real Gaussian error of parameter roughly $0.0534 \cdot r \approx 1.58$, which is easily distinguished from uniform modulo \mathbb{Z} .

Proof of Lemma 4.3. The latter isomorphism follows from the fact that the discriminants of the component fields are pairwise coprime. Because the degree of K (over \mathbb{Q}) matches that of the field from the claim (they

are both $2 \cdot 10 \cdot 2 = 40$), to prove the first claim it suffices to show that the two automorphisms τ, ν fix ζ_3, ζ_{11} , and $\sqrt{5 \cdot 17} \in L$.

Indeed, it is immediate that τ and ν fix both ζ_3 and ζ_{11} , because their \mathbb{Z}_3^* and \mathbb{Z}_{11}^* components are unity. Moreover, τ fixes $\sqrt{5 \cdot 17}$: it trivially fixes $\sqrt{17}$, and it also fixes $\sqrt{5}$ because $\mathbb{Q}(\sqrt{5})$ is the quadratic subfield of $\mathbb{Q}(\zeta_5)$, which is the fixed field of the order-two subgroup $\{\pm 1\} \subset \mathbb{Z}_5^*$. Finally, we claim that ν maps $\sqrt{5}$ to $-\sqrt{5}$ and maps $\sqrt{17}$ to $-\sqrt{17}$. Because 3 is a generator of \mathbb{Z}_5^* , the corresponding fixed subfield of $\mathbb{Q}(\zeta_5)$ is just \mathbb{Q} , therefore ν does not fix $\sqrt{5}$. Since any automorphism of L/\mathbb{Q} induces an automorphism on any subfield of L (because L/\mathbb{Q} is Galois, and hence normal), it follows that ν must map $\sqrt{5}$ to its only remaining conjugate, which is $-\sqrt{5}$. Identical reasoning shows that ν maps $\sqrt{17}$ to $-\sqrt{17}$, so ν fixes $\sqrt{5 \cdot 17}$, which proves the first claim.

For the second claim, $R = \mathcal{O}_K$ is isomorphic to $\mathbb{Z}[\zeta_3, \zeta_{11}] \otimes \mathbb{Z}[\delta]$, where $\delta = (1 + \sqrt{85})/2$. It follows that the dual ideal R^\vee is $S^\vee \otimes \mathbb{Z}[\delta]^\vee$, where $S = \mathbb{Z}[\zeta_3, \zeta_{11}]$. We know that $1/\sqrt{85} \in \mathbb{Z}[\delta]^\vee$, which has Euclidean norm $\sqrt{2/85}$ under the canonical embedding of $\mathbb{Q}(\sqrt{85})$. As in the proof of Lemma 4.2, the dual ideal S^\vee has a known \mathbb{Z} -basis whose 20 elements all have Euclidean norm $\sqrt{4/33}$ under the canonical embedding of $\mathbb{Q}(\zeta_3, \zeta_{11})$. The claim follows by taking the product of $1/\sqrt{85}$ with these basis elements. \square

Instantiation #2. This instantiation (ID 2 in Figure 5) uses a subfield K of the cyclotomic field $L = \mathbb{Q}(\zeta_{15015}) \cong \mathbb{Q}(\zeta_3, \zeta_5, \zeta_7, \zeta_{11}, \zeta_{13})$, namely, the fixed field of the Galois subgroup H generated by the three automorphisms τ, ν, ξ respectively indexed by

$$(1, 1, 1, -1, 1), (2, 3, 1, 1, 1), (2, 1, 1, 1, 2) \in \mathbb{Z}_3^* \times \mathbb{Z}_5^* \times \mathbb{Z}_7^* \times \mathbb{Z}_{11}^* \times \mathbb{Z}_{13}^*.$$

The Galois group of L/\mathbb{Q} has order $6720 = 2 \cdot 4 \cdot 7 \cdot 10 \cdot 12$, and the subgroup H has order 112, so the degree of K/\mathbb{Q} is 60.

Lemma 4.4. *The field K defined above is isomorphic to*

$$\mathbb{Q}(\zeta_7, \zeta_{11} + \zeta_{11}^{-1}, \sqrt{-3 \cdot 5 \cdot 13}) \cong \mathbb{Q}(\zeta_7) \otimes \mathbb{Q}(\zeta_{11} + \zeta_{11}^{-1}) \otimes \mathbb{Q}(\sqrt{-3 \cdot 5 \cdot 13}),$$

and the dual ideal R^\vee of its ring of integers $R = \mathcal{O}_K$ has 30 known \mathbb{Q} -linearly independent elements of Euclidean norm $\sqrt{12/(3 \cdot 5 \cdot 7 \cdot 11 \cdot 13)} \approx 0.02827$.

It follows that the root discriminant of K is $\delta_K = 7^{5/12} \cdot 11^{4/10} \cdot 195^{1/4} \approx 21.94$. For $r = \sqrt{2\pi} \cdot \delta_K \approx 55.00$ as in [CLS15], using the above elements of R^\vee in the distinguishing attack of Lemma 3.1 results in real Gaussian error of parameter roughly $0.02827 \cdot r \approx 1.55$, which is easily distinguished from uniform modulo \mathbb{Z} .

Proof of Lemma 4.4. The latter isomorphism follows from the fact that the discriminants of the components fields are coprime. Because the degree of K (over \mathbb{Q}) matches that of the field from the claim (they are both $6 \cdot 5 \cdot 2 = 60$), to prove the first claim it suffices to show that the three given automorphisms τ, ν, ξ fix $\zeta_7, \zeta_{11} + \zeta_{11}^{-1}$, and $\sqrt{-3 \cdot 5 \cdot 13} \in L$.

It is immediate that all three automorphisms fix ζ_7 , because their \mathbb{Z}_7^* components are all unity, and similarly for $\zeta_{11} + \zeta_{11}^{-1}$, because their \mathbb{Z}_{11}^* components are all ± 1 . It remains to show that they all fix $\sqrt{-3 \cdot 5 \cdot 13}$. For τ this is immediate, because it fixes $\sqrt{-3} \in \mathbb{Q}(\zeta_3)$, $\sqrt{5} \in \mathbb{Q}(\zeta_5)$, and $\sqrt{13} \in \mathbb{Z}(\zeta_{13})$. For ν and ξ , by similar reasoning as in the previous proof we have $\nu(\sqrt{-3}) = -\sqrt{-3}$, $\nu(\sqrt{5}) = -\sqrt{5}$, and $\nu(\sqrt{13}) = \sqrt{13}$; and $\xi(\sqrt{-3}) = -\sqrt{-3}$, $\xi(\sqrt{5}) = \sqrt{5}$, and $\xi(\sqrt{13}) = -\sqrt{13}$. This proves the first claim.

For the second claim, $R = \mathcal{O}_K$ is isomorphic to $\mathbb{Z}[\zeta_7] \otimes \mathbb{Z}[\zeta_{11} + \zeta_{11}^{-1}] \otimes \mathbb{Z}[\delta]$, where $\delta = (1 + \sqrt{-195})/2$. We know that $1/\sqrt{-195} \in \mathbb{Z}[\delta]^\vee$, which has Euclidean norm $\sqrt{2/195}$ under the canonical embedding

of $\mathbb{Q}(\sqrt{-195})$. As in the proof of Lemma 4.2, the dual ideal $\mathbb{Z}[\zeta_7]^\vee$ has a known \mathbb{Z} -basis whose 6 elements all have Euclidean norm $\sqrt{2/7}$. Similarly, $\mathbb{Z}[\zeta_{11} + \zeta_{11}^{-1}]^\vee$ has a known \mathbb{Z} -basis (namely, the dual of the basis made up of the conjugates of $\zeta_{11} + \zeta_{11}^{-1}$) whose 5 elements all have Euclidean norm $\sqrt{3/11}$. The claim follows by taking the product of $1/\sqrt{195}$ with the tensor product of the two just-described bases. \square

5 Invulnerable Instantiations

In this section we give sufficient conditions that make a Ring-LWE instantiation *provably immune* to all the attacks described in Section 3. By “immune” we mean that the attacks perform no better than known attacks (e.g., [BKW03, AG11]) against *plain* LWE when instantiated to have worst-case hardness, i.e., with Gaussian error of parameter $r \geq 2\sqrt{n}$. In particular, each attack’s running time divided by its advantage is at least $2^{\Omega(n)}$, in the typical case of polynomially bounded modulus $q = \text{poly}(n)$.

We focus on instantiations that satisfy, or only “almost” satisfy, the hypotheses of the “worst-case hardness of search” theorem from [LPR10, Section 4]. We show that any such instantiation, in *any* number field, satisfies the sufficient conditions, and is therefore immune to the attacks.

5.1 Class of Instantiations

Throughout the section, we consider instantiations of the “dual” Ring-LWE problem (Definition 2.6 and [LPR10, Section 3]) for the ring of integers R in a number field K of degree n (over \mathbb{Q}), with a continuous, spherical Gaussian error distribution $\psi = D_r$ over $K_{\mathbb{R}}$ for some $r > 0$. Recall from Section 2.3.2 that in this form of Ring-LWE,

$$s \in R_q^\vee := R^\vee/qR^\vee \quad \text{and} \quad a \in R_q := R/qR,$$

so $s \cdot a \in R_q^\vee$, and we have “noisy” products $b = s \cdot a + e \in K_{\mathbb{R}}/qR^\vee$ where $e \leftarrow \psi$.

For showing invulnerability to attacks, using continuous rather than discrete error yields stronger results that immediately transfer to the discrete setting. This is because the attacker can always discretize continuous samples, and thereby the underlying error, on its own if it so desires.⁸ We also note that all the results in this section apply (tautologically) to any equivalent form of Ring-LWE, e.g., the “tweaked” form that replaces R^\vee with R . For illustration, we depict some of these forms later in the section.

Invulnerability condition. We will show that a sufficient condition for invulnerability to the attacks from Section 3 is

$$r \geq 2. \tag{5.1}$$

While at first glance this bound may appear very small, remember that it should be compared against the high “density” of R^\vee , and in this respect the error is actually quite well spread relative to R^\vee . This will become apparent in the analysis and figures below.

We remark that Condition (5.1) is actually a bit weaker than what is required by [LPR10, Theorem 4.1] (worst-case hardness of search). Specifically, the theorem requires $r \geq 2 \cdot \omega(\sqrt{\log n})$, and moreover, it requires the search algorithm to work for *any* elliptical Gaussian error distribution whose parameter in each coordinate (of the canonical embedding) is bounded by r . These conditions may be artifacts of the proof technique, but in any case, they certainly require the attacker to succeed for spherical Gaussian error of some parameter $r \geq 2$, which is the case we study here.

⁸More precisely, this argument applies to any discretization $[\cdot]: K_{\mathbb{R}} \rightarrow R^\vee$ for which $[z + e] = z + [e]$ for any $z \in R^\vee$ and $e \in K_{\mathbb{R}}$, which is the case for any standard method. See [LPR13, Section 2.6] for further details.

5.2 Invulnerability to Attacks

Here we consider the two classes of attack described in Section 3.2: reducing to plain LWE, and reducing modulo an ideal divisor of qR . We prove that Condition (5.1) renders our class of instantiations invulnerable to both kinds of attack.

5.2.1 Reduction to LWE

As described in Section 2.3.2, this attack simply converts each Ring-LWE sample to n plain-LWE samples, and attempts to solve the resulting LWE instance. We emphasize that the attacker may use *arbitrary* \mathbb{Z} -bases of R and R^\vee to perform the transformation. More specifically, given each Ring-LWE sample

$$(a, b = s \cdot a + e) \in R_q \times K_{\mathbb{R}}/qR^\vee$$

where $e \leftarrow D_r$, we transform it to n LWE samples

$$(\mathbf{A}_a, \mathbf{b} = \mathbf{s}^t \mathbf{A}_a + \mathbf{e}^t),$$

where $\mathbf{b} \in (\mathbb{R}/q\mathbb{Z})^n$ and $\mathbf{e} \in \mathbb{R}^n$ are respectively the coefficient vectors of $b \in K_{\mathbb{R}}/qR^\vee$ and $e \in K_{\mathbb{R}}$ (with respect to the chosen basis of R^\vee), and $\mathbf{A}_a \in \mathbb{Z}_q^{n \times n}$ is the matrix of multiplication by $a \in R_q$ with any element of R_q^\vee (with respect to the chosen bases of R, R^\vee).

The following shows that the entries of the resulting error vector \mathbf{e} are Gaussians of parameter at least $2\sqrt{n}$, which is the exactly the lower bound from the worst-case hardness theorems for plain LWE [Reg05, Pei09].

Theorem 5.1. *For any \mathbb{Z} -basis $B^\vee = (b_j^\vee)$ of R^\vee used in the above reduction, each entry of \mathbf{e} is a continuous Gaussian of parameter at least $r\sqrt{n} \geq 2\sqrt{n}$.*

Proof. Let $B = (b_j)_j = (B^\vee)^\vee$ be the ordered \mathbb{Z} -basis of R that is dual to B^\vee , i.e., $\sigma(B)^* = \sigma(B^\vee)^{-1}$. Because $\mathbf{e} \in \mathbb{R}^n$ is the coefficient vector of $e \in K_{\mathbb{R}}$ with respect to basis B^\vee , by definition we have

$$\mathbf{e} = \sigma(B^\vee)^{-1} \cdot \sigma(e) = \sigma(B)^* \cdot \sigma(e).$$

Now because $B \subseteq R$ is a \mathbb{Z} -basis of R , all its elements are nonzero, so $\|\sigma(b_j)\| \geq \sqrt{n}$ by Lemma 2.4. Because the j th row of $\sigma(B)^*$ is $\sigma(b_j)^*$, the j th entry of \mathbf{e} is a continuous Gaussian of parameter $r\|\sigma(b_j)\| \geq r\sqrt{n} \geq 2\sqrt{n}$, as claimed. \square

We point out that while the Gaussian entries of \mathbf{e} have large width, they are not necessarily *independent*. It follows from the above proof that \mathbf{e} is distributed as a Gaussian with covariance matrix $r^2 \cdot \sigma(B)^* \cdot \sigma(B)/(2\pi)$. For example, when $B = (1, \zeta_p, \dots, \zeta_p^{p-2})$ is the power basis of the p th cyclotomic for prime p , the covariance matrix of \mathbf{e} is $r^2 \cdot (p\mathbf{I}_{p-1} - \mathbf{1})/(2\pi)$. Whether there are better attacks for this or other regimes that arise from reducing Ring-LWE to LWE is an interesting open question.

5.2.2 Reducing Modulo an Ideal

This attack uses an ideal divisor \mathfrak{q} of qR to attempt to solve decision-Ring-LWE, analogously to the attack described in Section 3.2. More specifically, we are given independent samples $(a_i \in R_q, b_i \in K_{\mathbb{R}}/qR^\vee)$, which are distributed either uniformly or according to the Ring-LWE distribution with some secret $s \in R_q^\vee$. We first reduce the samples to

$$(a'_i = a_i \bmod \mathfrak{q}, b'_i = b_i \bmod \mathfrak{q}R^\vee) \in R/\mathfrak{q} \times K_{\mathbb{R}}/\mathfrak{q}R^\vee,$$

and for each of the $N(\mathfrak{q})$ candidate reduced secrets $s' \in R^\vee/\mathfrak{q}R^\vee$, we test whether the $b'_i - a'_i \cdot s' \in K_{\mathbb{R}}/\mathfrak{q}R^\vee$ are non-uniform. The exact implementation of this test is not important for our purposes, because we will show that no test can meaningfully succeed.

For the attack to work, the reduced error distribution $D_r \bmod \mathfrak{q}R^\vee$ needs to have noticeable statistical distance from uniform; otherwise, the $b'_i - a'_i \cdot s'$ are close to uniform regardless of the form of the original samples. However, the following theorem shows that for *any* ideal \mathfrak{q} whose norm is not too large, and for error satisfying Condition (5.1), the statistical distance from uniform is exponentially small.

Theorem 5.2. *Let $\mathfrak{q} \subseteq R$ be any ideal of norm $N(\mathfrak{q}) \leq 2^n$, and let the error parameter $r \geq 2$ satisfy Condition (5.1). Then the reduced error distribution $D_r \bmod \mathfrak{q}R^\vee$ is within statistical distance 2^{-2n} of uniform over $K_{\mathbb{R}}/\mathfrak{q}R^\vee$.*

Proof. The dual ideal of $\mathfrak{q}R^\vee$ is $(\mathfrak{q}R^\vee)^\vee = \mathfrak{q}^{-1}$, which has norm $N(\mathfrak{q}^{-1}) = N(\mathfrak{q})^{-1} \geq 2^{-n}$. By Lemma 2.4, its minimum distance is

$$\lambda_1(\mathfrak{q}^{-1}) \geq \sqrt{n} \cdot N(\mathfrak{q}^{-1})^{1/n} \geq \sqrt{n}/2.$$

Then by Lemma 2.2, the smoothing parameter of $\mathfrak{q}R^\vee$ for $\varepsilon = 2^{-2n}$ is $\eta_\varepsilon(\mathfrak{q}R^\vee) \leq \sqrt{n}/\lambda_1(\mathfrak{q}^{-1}) \leq 2 \leq r$. The theorem then follows by Lemma 2.3. \square

5.3 Examples

For illustration, in this section we study some example invulnerable instantiations in detail, and contrast them with related insecure instantiations that were studied in Section 4.

5.3.1 Prime Cyclotomics

Let p be a prime integer and let $K = \mathbb{Q}(\zeta_p)$ and $R = \mathbb{Z}[\zeta_p]$ respectively denote the p th cyclotomic field and ring, which have degree $n = p - 1$ (over \mathbb{Q} and \mathbb{Z} , respectively). It is well known, and easy to check, that $R^\vee = p^{-1}(1 - \zeta_p)R$, and that $(R^\vee)^\vee = R$ has minimum distance $\lambda_1(R) = \sqrt{n}$ (witnessed by any power of ζ_p), which happens to be as small as possible relative to its norm. Therefore, $\eta_{2^{-2n}}(R^\vee) \leq \sqrt{n}/\lambda_1(R) = 1$, so spherical Gaussian error D_r of parameter $r \geq 2$ is extremely “smooth” modulo R^\vee (when p is large enough), and similarly for any ideal $\mathfrak{q}R^\vee$ where $N(\mathfrak{q}) \leq 2^n$, by Theorem 5.2. See Figure 6 for a depiction in the third cyclotomic.

Comparison with insecure instantiations. It is instructive to compare these invulnerable instantiations to the insecure non-dual prime-cyclotomic instantiations from [CLS15], as studied in Section 4.2. To make a direct comparison, we convert the dual form to an equivalent “tweaked” form, replacing R^\vee by R using a “tweak” factor of $t = p(1 - \zeta_p)^{-1}$, which satisfies $tR^\vee = R$. More specifically, we simply multiply each $b = s \cdot a + e \in K_{\mathbb{R}}/\mathfrak{q}R^\vee$ (where $s \in R_q^\vee$ and $a \in R_q$) by t to get

$$b' = t \cdot b = s' \cdot a + e' \in K_{\mathbb{R}}/\mathfrak{q}R,$$

where $s' = t \cdot s \in R_q$ and $e' = t \cdot e \in K_{\mathbb{R}}$.

Notice that for $p > 3$, the tweaked error distribution $t \cdot D_r$ of e' is (highly) non-spherical: the i th complex coordinate of D_r (for $i \in \mathbb{Z}_p^*$) is scaled by a factor of $\sigma_i(t) = p/(1 - \omega_p^i)$ for $\omega_p = \exp(2\pi\sqrt{-1}/p)$, which has magnitude ranging from about $p/2$ (for $i \approx p/2$) to $\Omega(p^2)$ (for $i \approx 0$). In comparison with the insecure error distributions from [CLS15], which involve spherical error of parameter $r = O(\sqrt{p})$, the tweaked error




Figure 6: The canonical embedding of: (in blue) the third cyclotomic ring $R = \mathbb{Z}[\zeta_3]$ for $q = 3$, (in maize) the dual ideal $R^\vee = q^{-1}(1 - \zeta_q)R$, and (in red) a spherical Gaussian of parameter $r = 2$. (We have depicted the two-dimensional real inner product space $H \subset \mathbb{C}^2$ as \mathbb{R}^2 .) Notice that the Gaussian is “well spread” relative to R^\vee , as implied by Theorem 5.2.

distribution has coordinates whose widths are larger by $\Omega(\sqrt{p})$ to $\Omega(p^{3/2})$ factors, as well as a very different non-spherical shape. (We remark that the non-spherical error is not a problem for applications, because the tweaked form is computationally equivalent to the dual form, which admits a fast error-sampling algorithm; see [LPR13, CP15] for full details.)

5.3.2 Quadratic Extensions (of Cyclotomics)

Let $d > 1$ be a square-free integer that is 3 modulo 4, define the totally real number field $K = \mathbb{Q}(\sqrt{d})$, and let $R = \mathbb{Z}[\sqrt{d}]$ denote its ring of integers; these have degree $n = 2$ over \mathbb{Q} and \mathbb{Z} , respectively. (If $d = 1 \pmod{4}$, then the ring of integers is $\mathbb{Z}[(1 + \sqrt{d})/2]$; what follows is easily adapted for this case.) It is easy to check that $R^\vee = (2\sqrt{d})^{-1}R$, and that $(R^\vee)^\vee = R$ has minimum distance $\lambda_1(R) = \sqrt{n}$, witnessed by $1 \in R$. Therefore, $\eta_{2-2n}(R^\vee) \leq \sqrt{n}/\lambda_1(R) = 1$, so spherical Gaussian error D_r of parameter $r \geq 2$ is “smooth” modulo R^\vee . See Figure 7 for a depiction for $d = 11$.

Of course, in the above example the degree $n = 2$ is a constant, so the smoothness error $\varepsilon = 2^{-2n}$ is actually not very small (and Ring-LWE is easily solved by brute force anyway). To increase the degree, as in Section 4.3 we can consider the quadratic extension $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$ of the p th cyclotomic, where p is prime and coprime with d ; here the degree is $n = 2(p - 1)$ over \mathbb{Z} . Then $R^\vee = (2p\sqrt{d})^{-1} \cdot (1 - \zeta_p)R$ and $\lambda_1(R) = \sqrt{n}$, which is witnessed by any power of ζ_p . Therefore, $\eta_{2-2n}(R^\vee) \leq 1$, so spherical Gaussian error D_r of parameter $r \geq 2$ is extremely “smooth” modulo R^\vee (when p is large enough).

Comparison with insecure instantiations. We now compare the above instantiations to the insecure non-dual instantiations from [CLS16] for the same class of rings, as studied in Section 4.3. To make a direct comparison, we convert the dual form to an equivalent “tweaked” form, replacing R^\vee by R using a “tweak” factor of $t = 2p\sqrt{d} \cdot (1 - \zeta_p)^{-1}$, which satisfies $tR^\vee = R$. Similarly to the cyclotomic case from Section 5.3.1 above, the coordinates of D_r are scaled by various factors having magnitudes that range from at least $p\sqrt{d}$ to




Figure 7: The canonical embedding of: (in maize) the dual ideal $R^\vee = (2\zeta)^{-1}R$ of the ring $R = \mathbb{Z}[\zeta]$ for $\zeta = \sqrt{11}$, (in blue) the ideal $\mathfrak{q}R^\vee$ where $\mathfrak{q} = (1 + \zeta)R + 5R$ is a prime ideal of norm $N(\mathfrak{q}) = 5$, and (in red) a spherical Gaussian of parameter $r = 2$. Notice that the Gaussian is “well spread” relative to R^\vee and even $\mathfrak{q}R^\vee$, as implied by Theorem 5.2. The “tweak” factor 2ζ that maps R^\vee to R has canonical embedding $(2\sqrt{11}, -2\sqrt{11})$, so it simply scales everything by a $2\sqrt{11}$ factor, and reflects over the horizontal axis.

$\Omega(p^2\sqrt{d})$. In comparison with the insecure error distributions from [CLS16], which involve spherical error of parameter $r = O(\sqrt{p \cdot d})$, the tweaked error distribution again has coordinates whose widths are larger by $\Omega(\sqrt{p})$ to $\Omega(p^{3/2})$ factors, and has a very different non-spherical shape.

References

- [ACPS09] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618. 2009.
- [ADPS16] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange - a new hope. In *USENIX Security Symposium*. 2016. To appear.
- [ADRS15] D. Aggarwal, D. Dadush, O. Regev, and N. Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling. In *STOC*, pages 733–742. 2015.
- [AG11] S. Arora and R. Ge. New algorithms for learning in presence of errors. In *ICALP (1)*, pages 403–415. 2011.
- [AKS01] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610. 2001.
- [AP13] J. Alperin-Sheriff and C. Peikert. Practical bootstrapping in quasilinear time. In *CRYPTO*, pages 1–20. 2013.

- [BCNS15] J. W. Bos, C. Costello, M. Naehrig, and D. Stebila. Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In *IEEE Symposium on Security and Privacy*, pages 553–570. 2015.
- [BKW03] A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
- [BLP⁺13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. 2013.
- [CIV16] W. Castryck, I. Iliashenko, and F. Vercauteren. Provably weak instances of ring-LWE revisited. In *EUROCRYPT*, pages 147–167. 2016.
- [CLS15] H. Chen, K. Lauter, and K. E. Stange. Attacks on search RLWE. Cryptology ePrint Archive, Report 2015/971, 2015. <http://eprint.iacr.org/>.
- [CLS16] H. Chen, K. Lauter, and K. E. Stange. Vulnerable Galois RLWE families and improved attacks. Cryptology ePrint Archive, Report 2016/193, 2016. <http://eprint.iacr.org/>.
- [Con09] K. Conrad. The different ideal, 2009. Available at <http://www.math.uconn.edu/~kconrad/blurbs/>, last accessed 12 Oct 2009.
- [CP15] E. Crockett and C. Peikert. $\Lambda \circ \lambda$: A functional library for lattice cryptography. Cryptology ePrint Archive, Report 2015/1134, 2015. <http://eprint.iacr.org/>.
- [EHL14] K. Eisenträger, S. Hallgren, and K. E. Lauter. Weak instances of PLWE. In *SAC*, pages 183–194. 2014.
- [ELOS15] Y. Elias, K. E. Lauter, E. Ozman, and K. E. Stange. Provably weak instances of ring-LWE. In *CRYPTO*, pages 63–92. 2015.
- [Lan94] S. Lang. *Algebraic number theory*, volume 110 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1994.
- [LP11] R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In *CT-RSA*, pages 319–339. 2011.
- [LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in Eurocrypt 2010.
- [LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In *EUROCRYPT*, pages 35–54. 2013.
- [MM11] D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *CRYPTO*, pages 465–484. 2011.
- [MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718. 2012.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.

- [MR09] D. Micciancio and O. Regev. Lattice-based cryptography. In *Post Quantum Cryptography*, pages 147–191. Springer, February 2009.
- [MV10] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358. 2010.
- [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009.
- [Pei14] C. Peikert. Lattice cryptography for the Internet. In *PQCrypto*, pages 197–219. 2014.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.