On Ideal Lattices and Learning with Errors Over Rings[*]

Vadim Lyubashevsky[†] Chris Peikert[‡] Oded Regev[§]

June 25, 2013

Abstract

The "learning with errors" (LWE) problem is to distinguish random linear equations, which have been perturbed by a small amount of noise, from truly uniform ones. The problem has been shown to be as hard as worst-case lattice problems, and in recent years it has served as the foundation for a plethora of cryptographic applications. Unfortunately, these applications are rather inefficient due to an inherent quadratic overhead in the use of LWE. A main open question was whether LWE and its applications could be made truly efficient by exploiting extra algebraic structure, as was done for lattice-based hash functions (and related primitives).

We resolve this question in the affirmative by introducing an algebraic variant of LWE called *ring-LWE*, and proving that it too enjoys very strong hardness guarantees. Specifically, we show that the ring-LWE distribution is pseudorandom, assuming that worst-case problems on ideal lattices are hard for polynomial-time quantum algorithms. Applications include the first truly practical lattice-based public-key cryptosystem with an efficient security reduction; moreover, many of the other applications of LWE can be made much more efficient through the use of ring-LWE.

1 Introduction

Over the last decade, lattices have emerged as a very attractive foundation for cryptography. The appeal of lattice-based primitives stems from the fact that their security can often be based on *worst-case* hardness assumptions, and that they appear to remain secure even against *quantum* computers.

Many lattice-based cryptographic schemes are based directly upon two natural average-case problems that have been shown to enjoy worst-case hardness guarantees. The *short integer solution* (SIS) problem was first shown in Ajtai's groundbreaking work [Ajt96] to be at least as hard as approximating several worst-case lattice problems, such as the (decision version of the) shortest vector problem, to within a polynomial factor

[*] A preliminary version of this paper appeared in EUROCRYPT 2010.

[†] INRIA and École Normale Supérieure, Paris. Part of this work was performed while at Tel Aviv University and also while visiting Georgia Tech. Partially supported by a European Research Council (ERC) Starting Grant.

[‡] School of Computer Science, College of Computing, Georgia Institute of Technology. This material is based upon work supported by the National Science Foundation under Grant CNS-0716786 and CAREER Award CCF-1054495, and by the Alfred P. Sloan Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the Sloan Foundation.

[§] Courant Institute, New York University. Part of the work done while the author was with the School of Computer Science, Tel Aviv University, Israel and CNRS, DI, ENS, Paris. Supported by the Binational Science Foundation, by the Israel Science Foundation, by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848, by the Wolfson Family Charitable Trust, and by a European Research Council (ERC) Starting Grant.

in the lattice dimension. More recently, Regev [Reg05] defined the *learning with errors* (LWE) problem and proved that it enjoys similar worst-case hardness properties, under a quantum reduction. (That is, an efficient algorithm for LWE would imply efficient quantum algorithms for approximate lattice problems.) As shown in two recent results [Pei09, BLP$^+$13], establishing the hardness of LWE under classical (non-quantum) reductions is also possible, but currently this is based on more restricted lattice problems.

The SIS problem may be seen as a variant of subset-sum over a particular additive group. In more detail, let $n \geq 1$ be an integer dimension and $q \geq 2$ be an integer modulus; the problem is, given polynomially many random and independent $\mathbf{a}_i \in \mathbb{Z}_q^n$, to find a 'small' integer combination of them that sums to $\mathbf{0} \in \mathbb{Z}_q^n$. The LWE problem is closely related to SIS, and can be stated succinctly as the task of distinguishing 'noisy linear equations' from truly random ones. More specifically, the goal is to distinguish polynomially many pairs of the form $(\mathbf{a}_i, b_i \approx \langle \mathbf{a}_i, \mathbf{s} \rangle) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ from *uniformly random* and independent pairs. Here $\mathbf{s} \in \mathbb{Z}_q^n$ is a uniformly random secret (which is kept the same for all pairs), each $\mathbf{a}_i \in \mathbb{Z}_q^n$ is uniformly random and independent, and each inner product $\langle \mathbf{a}_i, \mathbf{s} \rangle \in \mathbb{Z}_q$ is perturbed by a fresh random error term that is typically distributed like a (rounded) normal variable.

In recent years, a multitude of cryptographic schemes have been proposed around the SIS and LWE problems. The SIS problem has been the foundation for one-way [Ajt96] and collision-resistant hash functions [GGH96], identification schemes [MV03, Lyu08, KTX08], and digital signatures [GPV08, CHKP10, Boy10, MP12, Lyu12]. The LWE problem has proved to be amazingly versatile, serving as the basis for secure public-key encryption under both chosen-plaintext [Reg05, PVW08, LP11] and chosen-ciphertext [PW08, Pei09, MP12] attacks, oblivious transfer [PVW08], identity-based encryption [GPV08, CHKP10, ABB10a, ABB10b], various forms of leakage-resilient cryptography (e.g., [AGV09, ACPS09, GKPV10]), fully homomorphic encryption [BV11, BGV12] (following the seminal work of Gentry [Gen09]), and much more.

A main drawback of schemes based on the SIS and LWE problems, however, is that they tend not to be efficient enough for practical applications. Even the simplest primitives, such as one-way and collision-resistant hash functions, have key sizes and require computation times that are at least *quadratic* in the main security parameter, which needs to be in the several hundreds for sufficient security against known attacks (see, e.g., [MR09, LP11]).

A promising approach for avoiding this intrinsic inefficiency is to use lattices that possess extra algebraic structure. Influenced by the heuristic design of the NTRU cryptosystem [HPS98], Micciancio [Mic02] proposed a 'compact,' efficient one-way function (though not a collision-resistant one; see [PR06, LM06]) using a ring-based variant of SIS that he proved is at least as hard as worst-case problems on *cyclic* lattices. Later, Peikert and Rosen [PR06] and Lyubashevsky and Micciancio [LM06] independently showed that a modified ring-SIS problem is as hard as worst-case problems on *ideal* lattices (a generalization of cyclic lattices), which led to constructions of collision-resistant hash functions with practical implementations [LMPR08]. These results paved the way for other efficient cryptographic constructions, including identification schemes [Lyu09] and signatures [LM08, Lyu09], though not any public-key encryption applications.

Despite its expected utility, a compact analogue of LWE with comparable security properties has not yet appeared in the literature (though see Section 1.4 for discussion of a recent related work). Indeed, the perspectives and techniques that have so far been employed for the ring-SIS problem appear insufficient for adapting the more involved hardness proofs for LWE to the ring setting. Our main contributions in this paper are to define a ring-based variant of LWE and to prove its hardness under worst-case assumptions on ideal lattices.

1.1 Results

Here we give an informal overview of the ring-LWE problem and our hardness results for it. See Section 1.2 below for a discussion of some of the technical points omitted from this overview.

Let $f(x) = x^n + 1 \in \mathbb{Z}[x]$, where the security parameter n is a power of 2, making $f(x)$ irreducible over the rationals. Let $R = \mathbb{Z}[x]/\langle f(x) \rangle$ be the ring of integer polynomials modulo $f(x)$. Elements of R (i.e., residues modulo $f(x)$) can be represented by integer polynomials of degree less than n. Let $q = 1 \bmod 2n$ be a sufficiently large public prime modulus (bounded by a polynomial in n), and let $R_q = R/\langle q \rangle = \mathbb{Z}_q[x]/\langle f(x) \rangle$ be the ring of integer polynomials modulo both $f(x)$ and q. The q^n elements of R_q may be represented by polynomials of degree less than n whose coefficients are from some set of canonical representatives of $\mathbb{Z}_q$, e.g., $\{0, \ldots, q-1\}$.

The ring-LWE problem in R, denoted R-LWE, may be informally defined as follows (the formal, more general definition is given in Section 3): fix a certain error distribution over R that is concentrated on 'small' elements—informally, those having small integer coefficients—and let $s = s(x) \in R_q$ be uniformly random. Analogously to LWE, the goal is to distinguish arbitrarily many independent 'random noisy ring equations' from truly uniform pairs. More specifically, the noisy equations are of the form $(a, b \approx a \cdot s) \in R_q \times R_q$, where each a is uniformly random, and each product $a \cdot s$ is perturbed by a term drawn independently from the error distribution over R.

Main Theorem 1 (Informal). *Suppose that it is hard for polynomial-time* quantum *algorithms to approximate the search version of the shortest vector problem (*SVP*) in the* worst case *on* ideal lattices *in R to within a fixed* $\mathrm{poly}(n)$ *factor. Then any* $\mathrm{poly}(n)$ *number of samples drawn from the R-*LWE *distribution are pseudorandom to any polynomial-time (possibly quantum) attacker.*

For the ring R defined above, the family of ideal lattices is essentially the family of all "anti-cyclic integer lattices," i.e., lattices in $\mathbb{Z}^n$ that are closed under the operation that cyclically rotates the coordinates and negates the cycled element (see below for the more general definition of ideal lattices). We stress that we rely here on the *search* version of SVP, which is important since contrary to the case of general lattices, the decision version is typically easy to approximate on ideal lattices (see Lemma 2.9). For the same reason, adopting the approach behind *classical* hardness reductions for LWE [Pei09, BLP$^+$13], all of which seem to inherently rely on the decision version of SVP, would not be meaningful in this context (but see the end of Section 4.1).

Our main theorem follows from two component results: the first one (proved in Section 4) is a quantum reduction from worst-case approximate SVP on ideal lattices to the *search* version of ring-LWE; the second one (proved in Section 5) shows that the R-LWE distribution is in fact *pseudorandom* assuming that the search problem is hard. More details on the proof are given in Section 1.3 below.

Efficiency. For cryptographic applications, the R-LWE problem has many attractive features. First note the cryptographic strength of R-LWE versus standard LWE (or, for that matter, any other common number-theoretic assumption): each noisy product $b \approx a \cdot s$ gives n simultaneously pseudorandom values over $\mathbb{Z}_q$, rather than just one scalar, yet the cost of generating it is quite small: polynomial multiplication can be performed in $O(n \log n)$ scalar operations, and in parallel depth $O(\log n)$, using the Fast Fourier Transform (FFT) or its variants, with highly optimized implementations in practice (see [LMPR08] and the companion paper [LPR13]). Finally, in most applications each sample $(a, b) \in R_q \times R_q$ from the R-LWE distribution can replace n samples $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ from the standard LWE distribution, thus reducing the size of the public key (and often the secret key as well) by a factor of n. This is especially beneficial because key size has probably been the main barrier to practical lattice-based cryptosystems enjoying rigorous security analysis.

Security. Given the utility, flexibility, and efficiency of the ring-LWE problem, a natural question is: how plausible is the underlying assumption? All of the algebraic and algorithmic tools (including quantum computation) that we employ in our hardness reductions can also be brought to bear against SVP and other problems on ideal lattices. Yet despite considerable effort, no significant progress in attacking these problems has been made. The best known algorithms for ideal lattices perform essentially no better than their generic counterparts, both in theory and in practice. In particular, the asymptotically fastest known algorithms for obtaining an approximation to SVP on ideal lattices to within polynomial factors require time $2^{\Omega(n)}$, just as in the case of general lattices [AKS01, MV10].

We also gain some confidence in the hardness of ideal lattices from the fact that they arise naturally in algebraic number theory, a deep and well-studied branch of mathematics that has been investigated reasonably thoroughly from a computational point of view (see, e.g., [Coh93]). Due to their recent application in the design of cryptographic schemes, however, it is probably still too early to say anything about their security with great confidence. Further study is certainly a very important research direction.

Applications. As mentioned above, a remarkable number and variety of cryptographic constructions have been based on standard LWE. Most of these applications can be made more efficient, and sometimes even practical for real-world usage, by adapting them to ring-LWE. This process is often straightforward, but in some cases it requires additional technical tools to obtain the tightest and most efficient results. In a companion paper [LPR13] we give a collection of such tools, which include a strong 'regularity' lemma for the ring setting, tight bounds on the growth of error terms under ring operations, and fast special-purpose algorithms for important operations like generating error terms according to the appropriate distributions. We also construct several ring-LWE-based cryptosystems using these tools.

As one example application, here we sketch a simple and efficient semantically secure public-key cryptosystem, but defer a precise analysis and generalization to arbitrary cyclotomics to [LPR13]. For concreteness, fix the ring $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$ for n a power of 2. The key-generation algorithm chooses a uniformly random element $a \in R_q$ as well as two random 'small' elements $s, e \in R$ from the error distribution. It outputs s as the secret key and the pair $(a, b = a \cdot s + e) \in R_q^2$ as the public key. To encrypt an n-bit message $z \in \{0,1\}^n$, we view it as an element of R by using its bits as the 0-1 coefficients of a polynomial. The encryption algorithm then chooses three random 'small' elements $r, e_1, e_2 \in R$ from the error distribution and outputs the pair $(u, v) \in R_q^2$ as the encryption of z, where

$$u = a \cdot r + e_1 \bmod q \quad \text{and} \quad v = b \cdot r + e_2 + \lfloor q/2 \rceil \cdot z \bmod q.$$

The decryption algorithm simply computes

$$v - u \cdot s = (r \cdot e - s \cdot e_1 + e_2) + \lfloor q/2 \rceil \cdot z \bmod q.$$

For an appropriate choice of parameters, the coefficients of $r \cdot e - s \cdot e_1 + e_2 \in R$ have magnitudes less than $q/4$, so the bits of z can be recovered by rounding each coefficient of $v - u \cdot s$ back to either 0 or $\lfloor q/2 \rceil$, whichever is closest modulo q. Notice the system's resemblance to the Diffie-Hellman key-agreement protocol [DH76] and ElGamal cryptosystem [ElG84], where $a \in R_q$ is analogous to the generator of a (multiplicative) cyclic group, and taking noisy products is analogous to exponentiation.

Semantic security follows from two easy applications of the pseudorandomness of ring-LWE. First we note that ring-LWE samples are pseudorandom even when the *secret* is also chosen from the error distribution, by a transformation to the "(Hermite) normal form" analogous to the one for standard LWE [MR09, ACPS09]. Therefore, the public key $(a, b) \in R_q$ is pseudorandom, so as a thought experiment we may replace it with a truly uniform pair. Then we see that (ignoring the message component $\lfloor q/2 \rceil \cdot z$) the pairs $(a, u), (b, v) \in R_q^2$, which constitute the entire view of a passive adversary, are ring-LWE samples with secret r and hence are also pseudorandom, which implies semantic security.

1.2 More Details

Here we fill in some of the missing details in the high-level description above.

The underlying ring. Our main focus in this work is on the rings $\mathbb{Z}[x]/\langle \Phi_m(x) \rangle$ of integer polynomials modulo a *cyclotomic polynomial* $\Phi_m(x)$. To recall, the mth cyclotomic polynomial $\Phi_m(x) \in \mathbb{Z}[x]$ is the polynomial of degree $n = \varphi(m)$ whose roots are all the primitive mth roots of unity $\omega_m^i \in \mathbb{C}$, where $\omega_m = \exp(2\pi\sqrt{-1}/m)$ and $1 \leq i < m$ with i coprime to m. For instance, when $m \geq 2$ is a power of 2, we have $\Phi_m(x) = x^n + 1$ where $n = m/2$. From an algebraic point of view, it is more natural to view these rings as the rings of algebraic integers in cyclotomic number fields (as opposed to rings of polynomials), and this is indeed the perspective we adopt.

Rings of integers in (not necessarily cyclotomic) number fields have some nice algebraic properties that are essential to our results. For instance, they have unique factorization of ideals, and their fractional ideals form a multiplicative group; in general, neither property holds in $\mathbb{Z}[x]/\langle f(x) \rangle$ for monic irreducible $f(x)$, as demonstrated by the ring $\mathbb{Z}[x]/\langle x^2+3 \rangle = \mathbb{Z}[\sqrt{-3}]$. (For example, in this ring $4 = 2^2 = (1+\sqrt{-3})(1-\sqrt{-3})$, but 2, $1 + \sqrt{-3}$, and $1 - \sqrt{-3}$ are all irreducible.) In addition, *cyclotomic* number fields have Galois (i.e., automorphism) groups that 'act transitively' on the prime ideals dividing a given prime integer while still preserving the LWE error distribution, which is an essential fact used in the second component of our proof. The first component of our proof does not need this latter property, and therefore applies more generally to rings of integers in arbitrary (not necessarily cyclotomic) number fields. It is likely that our second component can also be somewhat extended beyond cyclotomic number fields, but we do not attempt to do so here.

Ideal lattices and the canonical embedding. Fix some underlying ring R, e.g., the ring of algebraic integers in a cyclotomic number field as above. Recall that an *ideal* is an additive subgroup that is closed under multiplication by elements of R. In the case of cyclotomic rings $\mathbb{Z}[x]/\langle \Phi_m(x) \rangle$, the latter constraint is equivalent to requiring closure under multiplication by x. Fix also an additive isomorphism σ mapping the ring R to some lattice $\sigma(R)$ in $\mathbb{R}^n$. For instance, the naïve "coefficient embedding" maps any element of R to the integer vector in $\mathbb{Z}^n$ whose coordinates are exactly the coefficients of that element when viewed as a polynomial residue. The family of ideal lattices (for the ring R and embedding σ) is the set of all lattices $\sigma(\mathcal{I})$ for ideals $\mathcal{I}$ in R. For example, when working with the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$ for n a power of 2 and the coefficient embedding, one obtains the family of all "anti-cyclic integer lattices" mentioned above.

Unlike almost all previous works in the area (e.g., [Mic02, PR06, LM06, LM08, Gen09, Lyu09, SSTX09]), we choose not to use the naïve coefficient embedding, but instead exclusively use the so-called *canonical embedding* from algebraic number theory (also adopted in the prior work [PR07]), whose definition will appear later.

By definition, any two embeddings are related to each other simply by a fixed linear transformation on $\mathbb{R}^n$. Moreover, in many cases the distortion introduced by this transformation is small; for example, in the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$ where n is a power of 2, the transformation is even an isometry (i.e., a scaled rotation). In particular, worst-case lattice problems like approximate-SVP are equivalent under any two embeddings, up to a factor corresponding to the distortion between them. Hence, one might wonder why bother with the canonical embedding and not just work with the naïve one. Yet due to its central role in the study of number fields and many useful geometric properties, we contend that the canonical embedding is the 'right' notion to use in the study of ideal lattices. We elaborate on this point in the last part of this section.

Error distribution. One important issue we have ignored so far is the precise error distribution in the definition of ring-LWE for which our hardness results hold. As in the standard LWE problem, the error distribution we use is a (centered) Gaussian. However, unlike the standard LWE problem where the error is a one-dimensional Gaussian (and hence the distribution can be specified by just one parameter, the standard deviation), here the error is an n-dimensional Gaussian. While in general specifying an n-dimensional centered Gaussian distribution requires an entire n-by-n covariance matrix, our error distributions are always diagonal in the canonical embedding. In other words, when viewed in the canonical embedding, our error distributions are product distributions in which each component is a (one-dimensional, centered) normal distribution with a certain standard deviation, and hence an entire error distribution is defined by just n parameters. When all these parameters are equal, we say that the distribution is spherical.

Notice that all of the above is under the assumption that we are using the canonical embedding. When using another embedding (say, the naïve coefficient embedding), the error distribution is still a multivariate Gaussian (since a linear transformation of a Gaussian is Gaussian), but its coordinates need no longer be independent. (One exception is spherical error with an underlying ring $\mathbb{Z}[x]/\langle x^n + 1\rangle$ for n a power of 2; here the two embeddings are isometric, and hence the error also has i.i.d. coordinates in the coefficient embedding, albeit with a different standard deviation due to the scaling involved.)

For our ring-LWE hardness results, the search problem requires a solution for *any* Gaussian error distribution whose n parameters are all at most some parameter α. For the average-case decision problem, the n parameters are themselves chosen at random and kept secret (see Definition 3.5). This situation is in contrast with standard LWE, where the error distribution, being one-dimensional, is simply a fixed normal distribution.

The above non-spherical error distributions might be an artifact of our proof technique, and although they typically do not cause any serious problems, they might make certain applications and their proofs more cumbersome. Fortunately, if we restrict the ring-LWE problem (in either its search or decision form) to any *bounded* number ℓ of samples, then we can prove hardness for a *fixed, spherical* error distribution that is only about an $\ell^{1/4}$ factor wider than the non-spherical one with random parameters (see Theorem 5.2). Because the security reductions for most ring-LWE-based cryptographic schemes use only a small number of samples (often, $\ell = O(1)$ or $\ell = O(\log n)$), it is appropriate and simpler to use spherical error in those applications. Finally, we mention that if one assumes the hardness of the search problem with a fixed spherical Gaussian error distribution and unbounded samples (which seems plausible, but is not implied by our worst-case hardness proof), then the average-case decision problem for the very same error distribution (and unbounded samples) is also hard (see Theorem 5.3).

In praise of the canonical embedding. While the number-theoretic perspective on ideal lattices (and in particular the use of the canonical embedding) requires some investment in the mathematical background, we find that it delivers many nice geometric and algebraic properties that pay dividends in the ease of working with the objects, and in the strength and generality of results that can be obtained. We now describe a few examples of this.

First, unlike the coefficient embedding, under the canonical embedding both addition *and* multiplication of ring elements are simply coordinate-wise. As a result, both operations have simple geometric interpretations that lead to tight bounds, and product distributions (such as Gaussians) behave very nicely under both addition and multiplication. By contrast, analyzing multiplication under the coefficient embedding required previous works to use rather crude quantities like the "expansion factor" of the ring. The expansion factor bounds the *worst-case* ratio of $\|\sigma(a \cdot b)\|$ to $\|\sigma(a)\| \cdot \|\sigma(b)\|$ over all $a, b \in R$, but on average (over the random choice of a, b from natural distributions), it is often quite loose. Moreover, it does not provide any more detailed information about how $a \cdot b$ relates to a and b geometrically, e.g., for analyzing probability distributions.

Second, although for many rings the canonical and coefficient embeddings are (nearly) isometric, in many other rings of interest the distortion between them can be very large—even *super-polynomial* in the dimension for some cyclotomic polynomial rings [Erd46]. This may explain why previous work was mostly restricted to $\mathbb{Z}[x]/\langle x^n + 1\rangle$ for n a power of 2, and a few other concrete rings, whereas we can prove tight geometric bounds and hardness results for *all* cyclotomic rings (regardless of their expansion factor).

A third point in favor of the canonical embedding is that it behaves very nicely under the automorphisms that are crucial to the second component of our proof: they simply permute the axes of the embedding.

1.3 Proof Outline and Techniques

As mentioned before, our main theorem consists of two component results, which we now describe in more detail. We note that the two parts are essentially independent, and can be read separately.

First component: worst-case hardness of the search problem. In the first component we give a quantum reduction from approximate SVP (in the worst case) on ideal lattices in R to the *search* version of ring-LWE, where the goal is to *recover* the secret $s \in R_q$ (with high probability, for any s) from arbitrarily many noisy products. This result is stated formally as Theorem 4.1, and is proved throughout Section 4. As already mentioned before, this reduction actually works in *general* (not necessarily cyclotomic) number fields.

Our reduction follows the general outline of Regev's iterative quantum reduction for general lattices [Reg05]. In fact, we use the quantum part of the reduction in [Reg05] essentially as a black box; the main effort is in the classical (non-quantum) part, and requires perspectives and tools from algebraic number theory such as the canonical embedding and the Chinese Remainder Theorem (CRT).

In particular, one of the main technical contributions is the use of the CRT for 'clearing the ideal' $\mathcal{I}$ from an arbitrary ideal lattice instance (see Lemmas 2.14 and 2.15). This involves mapping the quotient ring $\mathcal{I}/q\mathcal{I}$ to the fixed quotient ring R/qR in an 'algebraically consistent' way (formally, as an isomorphism of R-modules). We believe that this technique should be useful elsewhere; in particular, it implies simpler and slightly tighter hardness proofs for ring-SIS through the use of the 'discrete Gaussian' style of worst-case to average-case reduction from [GPV08]. Lacking this technique, prior reductions for ideal lattices following [Mic02] used samples from a *principal subideal* of $\mathcal{I}$ with known generator; however, this restriction does not seem compatible with the approaches of [Reg05, GPV08], where the reduction must deal with Gaussian samples from the full ideal $\mathcal{I}$.

Second component: search / decision equivalence. In the second component we give a reduction from the search problem (shown hard in the first component) to the decision variant, thereby showing that the R-LWE distribution is *pseudorandom*. As alluded to before, we actually provide two variants of the reduction: one to the decision problem with a nonspherical error distribution in the canonical embedding (Theorem 5.1), and one to the decision problem with a spherical error distribution but with a bounded number of samples (Theorem 5.2). We stress that these reductions are entirely classical (not quantum) and hence if one is willing to assume the classical hardness of the search problem, one gets classical hardness of the decision problem.

Moreover, if we assume hardness of the search problem under a *fixed* spherical Gaussian error distribution (which is not implied by our worst-case hardness proof), then an easy simplification of our search-to-decision reduction (Theorem 5.3) gives hardness of the decision variant under the same error distribution. The same can be proved for many other natural error distributions, which demonstrates that our second component is of value even without the first one.

Our approach is also inspired by analogous reductions for the standard LWE problem [BFKL93, Reg05], but again the ring context presents significant new obstacles, primarily related to proving that the entire n-dimensional quantity $b \approx a \cdot s$ is pseudorandom. Here again, the solution seems to rely inherently on tools from algebraic number theory: we develop new techniques that exploit special properties of cyclotomic number fields of degree n — namely, that they are *Galois* (i.e., have n automorphisms) — and our particular choice of modulus $q = 1 \bmod m$ — namely, that $\langle q \rangle$ 'splits completely' into n prime ideals $\mathfrak{q}_i$ each of norm $q = \mathrm{poly}(n)$, which are permuted transitively by the automorphisms. (Interestingly, this complete splitting is also useful for performing the ring operations very efficiently in practice; see [LPR13]).

The basic outline of the reduction is as follows. First, by a hybrid argument we show that any distinguisher between the uniform distribution and the ring-LWE distribution with secret $s \in R_q$ must have some noticeable advantage relative to *some* prime ideal factor $\mathfrak{q}_i$ of $\langle q \rangle$ (of the distinguisher's choice); this advantage can be amplified using standard self-reduction techniques. Next, we give an efficient search-to-decision reduction that finds the value of s modulo $\mathfrak{q}_i$, using the fact that the ring modulo $\mathfrak{q}_i$ is a field of order $q = \mathrm{poly}(n)$. Then, because the automorphisms of the number field permute the $\mathfrak{q}_i$, we can find s modulo *every* $\mathfrak{q}_j$ by applying an appropriate automorphism to the ring-LWE distribution. (Crucially, the error distribution also remains legal under the automorphisms.) This lets us recover all of $s \bmod q$ using the Chinese Remainder Theorem.

1.4 Related Work

In a concurrent and independent work, Stehlé, Steinfeld, Tanaka, and Xagawa [SSTX09] formulated a variant of LWE quite similar to ours. We believe that our results subsume those of [SSTX09], although their techniques, being quite modular, are of independent interest and might have further applications.

In more detail, their main result is analogous to our first component, showing hardness of the search problem based on worst-case lattice problems and using a quantum reduction. However, whereas we show a quantum reduction directly from the worst-case lattice problems, Stehlé et al. show a quantum reduction from the ring-SIS problem, which they then combine with prior (classical) reductions from worst-case lattice problems to ring-SIS [PR06, LM06]. Their reduction highlights a nice duality between (ring-)LWE and (ring-)SIS (first observed in [GPV08]), and builds on the quantum machinery from [Reg05], together with some new observations.

Although both reductions show hardness of the search problem, there are a couple of notable differences. Whereas our reduction shows hardness of the search problem with an *unbounded* number of samples, the reduction of Stehlé et al. shows hardness of the search problem with any *a priori bounded* number of samples. It is probably possible to generate an unbounded number of samples from this bounded number of given samples by taking random combinations, although this would incur an additional loss in the parameters. Another difference is that their proof is presented only for the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$ for n a power of 2, whereas ours works for the ring of integers in any number field.

Probably the most significant difference between our work and that of [SSTX09] is that the latter has no analogue of our second component, namely the search-to-decision reduction. As a result, for cryptographic applications Stehlé et al. use hard-core bits obtained via the efficient Goldreich-Levin construction based on Toeplitz matrices [Gol04, Section 2.5]. This approach, however, induces a security reduction that runs in *exponential* time in the number of hard bits. In particular, encrypting in amortized $\tilde{O}(1)$ time per message bit induces the assumption that worst-case lattice problems are hard for $2^{o(n)}$-time quantum algorithms. This much stronger assumption is of course quite undesirable, and moreover, since it requires a higher dimension n for the same level of security, the efficiency of the resulting cryptosystem (which is the main reason for using ring-LWE in the first place) is harmed. In contrast, we obtain a linear number of hard bits by showing directly that the ring-LWE distribution is pseudorandom; in particular, this yields a cryptosystem with the same (or even slightly better) running times under a fully polynomial security reduction.

Subsequent work. Since the publication of a preliminary version of this paper, several works have appeared which use our results for cryptographic purposes. These include the work of Stehlé and Steinfeld [SS11] who show how a slight modification of the NTRU cryptosystem can be based on ring-LWE, constructions of fully homomorphic encryption schemes by Brakerski, Gentry, and Vaikuntanathan [BV11, BGV12] and multikey fully homomorphic schemes by López-Alt, Tromer, and Vaikuntanathan [LTV12], and more. Also, Langlois and Stehlé [LS13] extended the pseudorandomness of ring-LWE to essentially all choices of the modulus q.

Outline. We start in Section 2 with some background material on lattices and Gaussian measures, followed by an overview of concepts from algebraic number theory required for our proofs. Although the latter material is mostly standard, we are not aware of any single accessible reference that covers all the necessary background. Section 3 gives the formal definition of the ring-LWE problem, both in its search and average-case decision versions, and states our main theorem. In Section 4 we prove the hardness of the search ring-LWE problem. We continue in Section 5 with several reductions to the average-case decision problem. The latter two sections are the main contributions of the paper, and are essentially independent of each other.

Acknowledgments. We thank Damien Stehlé for useful discussions, and for sharing with us, together with Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa, an early draft of their result.

2 Preliminaries

For a vector $\mathbf{x}$ in $\mathbb{R}^n$ or $\mathbb{C}^n$ and $p \in [1, \infty]$, we define the ℓ_p norm as $\|\mathbf{x}\|_p = (\sum_{i \in [n]} |x_i|^p)^{1/p}$ when $p < \infty$, and $\|\mathbf{x}\|_\infty = \max_{i \in [n]} |x_i|$ when $p = \infty$.

2.1 The Space H

When working with number fields and ideal lattices, it is convenient to work with the space $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ for some numbers $s_1 + 2s_2 = n$, defined as

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

It is not difficult to verify that H (with the inner product induced on it by $\mathbb{C}^n$) is isomorphic to $\mathbb{R}^n$ as an inner product space. This can seen via the orthonormal basis $\{\mathbf{h}_i\}_{i \in [n]}$, defined as follows: for $j \in [n]$, let $\mathbf{e}_j \in \mathbb{C}^n$ be the vector with 1 in its jth (complex) coordinate, and 0 elsewhere; then for $j \in [s_1]$, we take $\mathbf{h}_j = \mathbf{e}_j \in \mathbb{C}^n$ and for $s_1 < j \leq s_1 + s_2$ we take $\mathbf{h}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_j + \mathbf{e}_{j+s_2})$ and $\mathbf{h}_{j+s_2} = \frac{\sqrt{-1}}{\sqrt{2}}(\mathbf{e}_j - \mathbf{e}_{j+s_2})$. Note that the complex conjugation operation (which maps H to itself) acts in the $\{\mathbf{h}_i\}_{i \in [n]}$ basis by flipping the sign of all coordinates in $\{s_1 + s_2 + 1, \ldots, n\}$.

We will also equip H with the ℓ_p norm induced on it from $\mathbb{C}^n$. Namely, for any $a_1, \ldots, a_n \in \mathbb{R}$, the ℓ_p norm of the element $\sum a_i \mathbf{h}_i \in H$ is given by

$$\Big\| \sum_{i=1}^{n} a_i \mathbf{h}_i \Big\|_p = \Big(\sum_{i=1}^{s_1} |a_i|^p + 2 \sum_{i=s_1+1}^{s_1+s_2} \Big(\tfrac{a_i^2 + a_{i+s_2}^2}{2} \Big)^{p/2} \Big)^{1/p}.$$

We note that for any $p \in [1, \infty]$, this norm is equal within a factor of $\sqrt{2}$ to $(\sum_{i=1}^{n} |a_i|^p)^{1/p}$, which is the ℓ_p norm induced on H from the isomorphism with $\mathbb{R}^n$ described above; for the ℓ_2 norm, we in fact have an equality. This near equivalence between H and $\mathbb{R}^n$ will allow us to use known definitions and results on lattices in our setting, the only minor caveat being the $\sqrt{2}$ factor when dealing with ℓ_p norms for $p \neq 2$.

2.2 Lattice Background

We define a *lattice* as a discrete additive subgroup of H. We deal exclusively with full-rank lattices, which are generated as the set of all integer linear combinations of some set of n linearly independent *basis* vectors $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\} \subset H$:

$$\Lambda = \mathcal{L}(\mathbf{B}) = \Big\{ \sum_{i \in [n]} z_i \mathbf{b}_i \ : \ \mathbf{z} \in \mathbb{Z}^n \Big\}.$$

The *minimum distance* $\lambda_1(\Lambda)$ of a lattice Λ in a given norm $\|\cdot\|$ is the length of a shortest nonzero lattice vector: $\lambda_1(\Lambda) = \min_{\mathbf{0} \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\|$. We use the Euclidean norm unless stated otherwise.

The *dual lattice* of $\Lambda \subset H$ is defined as $\Lambda^* = \{\mathbf{x} \in H : \langle \Lambda, \mathbf{x} \rangle \subseteq \mathbb{Z}\}$. It is easy to see that $(\Lambda^*)^* = \Lambda$.

2.2.1 Gaussian Measures

For $r > 0$, define the Gaussian function $\rho_r \colon H \to (0, 1]$ as $\rho_r(\mathbf{x}) = \exp(-\pi \langle \mathbf{x}, \mathbf{x} \rangle / r^2) = \exp(-\pi \|\mathbf{x}\|_2^2 / r^2)$. By normalizing this function we obtain the *continuous* Gaussian probability distribution D_r of width r, whose density is given by $r^{-n} \cdot \rho_r(\mathbf{x})$. We extend this to elliptical (non-spherical) Gaussian distributions in the basis $\{\mathbf{h}_i\}_{i \in [n]}$ as follows. Let $\mathbf{r} = (r_1, \ldots, r_n) \in (\mathbb{R}^+)^n$ be a vector of positive real numbers such that $r_{j+s_1+s_2} = r_{j+s_1}$ for each $j \in [s_2]$. Then a sample from $D_{\mathbf{r}}$ is given by $\sum_{i \in [n]} x_i \mathbf{h}_i$, where the x_i are chosen independently from the (one-dimensional) Gaussian distribution D_{r_i} over $\mathbb{R}$.

Micciancio and Regev [MR04] introduced a lattice quantity called the *smoothing parameter*, and related it to various lattice quantities.

Definition 2.1. *For a lattice Λ and positive real $\varepsilon > 0$, the* smoothing parameter *$\eta_\varepsilon(\Lambda)$ is defined to be the smallest r such that $\rho_{1/r}(\Lambda^* \backslash \{\mathbf{0}\}) \leq \varepsilon$.*

Lemma 2.2 ([MR04, Lemmas 3.2, 3.3]). *For any n-dimensional lattice Λ, we have $\eta_{2^{-2n}}(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^*)$,*[1] *and $\eta_\varepsilon(\Lambda) \leq \sqrt{\ln(n/\varepsilon)}\, \lambda_n(\Lambda)$ for all $0 < \varepsilon < 1$.*

The following lemma explains the name "smoothing parameter."

Lemma 2.3 ([MR04, Lemma 4.1] and [Reg05, Claim 3.8]). *For any lattice Λ, $\varepsilon > 0$, $r \geq \eta_\varepsilon(\Lambda)$, and $\mathbf{c} \in H$, the statistical distance between $(D_r + \mathbf{c}) \bmod \Lambda$ and the uniform distribution modulo Λ is at most $\varepsilon/2$. Alternatively, we have $\rho_r(\Lambda + \mathbf{c}) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \rho_r(\Lambda)$.*

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

For a lattice Λ, point $\mathbf{u} \in H$, and real $r > 0$, define the *discrete Gaussian probability distribution over* $\Lambda + \mathbf{u}$ with parameter r as

$$D_{\Lambda+\mathbf{u},r}(\mathbf{x}) = \frac{\rho_r(\mathbf{x})}{\rho_r(\Lambda + \mathbf{u})} \quad \forall\, \mathbf{x} \in \Lambda + \mathbf{u}.$$

Lemma 2.4 ([Ban93, Lemma 1.5(i)]). *For any n-dimensional lattice Λ and $r > 0$, a point sampled from $D_{\Lambda,r}$ has Euclidean norm at most $r\sqrt{n}$, except with probability at most 2^{-2n}.*

We also need the following property of the smoothing parameter, which says that continuous noise 'smooths' the discrete structure of a discrete Gaussian distribution into a continuous one.

Lemma 2.5 ([Reg05]). *Let Λ be a lattice, let $\mathbf{u} \in H$ be any vector, and let $r, s > 0$ be reals. Assume that $1/\sqrt{1/r^2 + 1/s^2} \geq \eta_\varepsilon(\Lambda)$ for some $\varepsilon < \frac{1}{2}$. Consider the continuous distribution Y on H obtained by sampling from $D_{\Lambda+\mathbf{u},r}$ and then adding an element drawn independently from D_s. Then the statistical distance between Y and $D_{\sqrt{r^2+s^2}}$ is at most 4ε.*

2.3 Algebraic Number Theory Background

Algebraic number theory is the study of *number fields*. Here we review the necessary background, presenting for concreteness the special case of *cyclotomic* number fields as a running example. In this subsection we cover the relevant mathematical and computational background; in Section 2.4 we cover additional special properties of cyclotomic number fields. More background and complete proofs can be found in any introductory book on the subject, e.g., [Ste04].

2.3.1 Number Fields

A *number field* can be defined as a field extension $K = \mathbb{Q}(\zeta)$ obtained by adjoining an abstract element ζ to the field of rationals, where ζ satisfies the relation $f(\zeta) = 0$ for some irreducible polynomial $f(x) \in \mathbb{Q}[x]$, which is monic without loss of generality. The polynomial f is called the *minimal polynomial* of ζ, and the *degree* n of the number field is the degree of f. Because $f(\zeta) = 0$, the number field K can be seen as an n-dimensional vector space over $\mathbb{Q}$ with basis $\{1, \zeta, \ldots, \zeta^{n-1}\}$; this is called the *power basis* of K. Of course, associating ζ with indeterminate x yields a natural isomorphism between K and $\mathbb{Q}[x]/f(x)$.

Let m be a positive integer, and let $\zeta = \zeta_m$ denote an element of multiplicative order m, i.e., a primitive mth root of unity. The mth *cyclotomic* number field is $K = \mathbb{Q}(\zeta)$, and the minimal polynomial of ζ is the mth *cyclotomic polynomial*

$$\Phi_m(x) = \prod_{i \in \mathbb{Z}_m^*} (x - \omega_m^i) \in \mathbb{Z}[x],$$

where $\omega_m \in \mathbb{C}$ is any primitive mth complex root of unity, e.g., $\omega_m = \exp(2\pi\sqrt{-1}/m)$. Observe that the complex roots ω_m^i of $\Phi_m(x)$ are exactly the primitive mth roots of unity in $\mathbb{C}$, and that $\Phi_m(x)$ has degree $n = \varphi(m)$, the totient of m. The form of $\Phi_m(x)$ will not play any role in this work, aside from the fact that it is computable in polynomial time given m (in unary).

2.3.2 Embeddings and Geometry

Here we describe the *embeddings* of a number field, which induce a natural 'canonical' geometry on it.

A number field $K = \mathbb{Q}(\zeta)$ of degree n has exactly n ring embeddings (injective ring homomorphisms) $\sigma_i \colon K \to \mathbb{C}$. Concretely, these embeddings map ζ to each of the complex roots of its minimal polynomial f;

it is easy to see that these are the only ring embeddings from K to $\mathbb{C}$, because $f(\sigma_i(\zeta)) = \sigma_i(f(\zeta)) = 0$. 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 called a complex embedding. Because complex roots of $f(x)$ come in conjugate pairs, so too do the complex embeddings. The number of real embeddings is denoted s_1 and the number of *pairs* of complex embeddings is denoted s_2, so we have $n = s_1 + 2s_2$. By convention, we let $\{\sigma_j\}_{j\in[s_1]}$ be the real embeddings, and we order the complex embeddings so that $\sigma_{s_1+s_2+j} = \overline{\sigma_{s_1+j}}$ for $j \in [s_2]$. The *canonical embedding* $\sigma\colon K \to \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ is then defined as

$$\sigma(x) = (\sigma_1(x), \ldots, \sigma_n(x)).$$

Note that it is a ring homomorphism from K to $\mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$, where multiplication and addition in the latter are both component-wise. Due to the pairing of the complex embeddings, we have that σ maps into H (defined in Section 2.1 above).

By identifying elements of K with their canonical embeddings in H, we can speak of geometric norms (e.g., the Euclidean norm) on K. Recalling that we define norms on H as those induced from $\mathbb{C}^n$, we see that for any $x \in K$ and any $p \in [1, \infty]$, the ℓ_p norm of x is simply $\|x\|_p = \|\sigma(x)\|_p = (\sum_{i\in[n]}|\sigma_i(x)|^p)^{1/p}$ for $p < \infty$, and is $\max_{i\in[n]}|\sigma_i(x)|$ for $p = \infty$. (As always, we assume the ℓ_2 norm when p is omitted.) Because multiplication of embedded elements is component-wise (since σ is a ring homomorphism), we have

$$\|x \cdot y\|_p \leq \|x\|_\infty \cdot \|y\|_p$$

for any $x, y \in K$ and any $p \in [1, \infty]$. Thus the ℓ_∞ norm acts as an 'absolute value' for K that bounds how much an element 'expands' any other by multiplication.

Using the canonical embedding also allows us to think of the Gaussian distribution $D_{\mathbf{r}}$ for $\mathbf{r} \in (\mathbb{R}^+)^n$ over H, or its discrete analogue over a lattice in H, as a distribution over K. Strictly speaking, the distribution $D_{\mathbf{r}}$ is not over K, but rather over the field tensor product $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R}$, which is isomorphic to H. Since multiplication of elements in $K_\mathbb{R}$ is mapped to coordinate-wise multiplication in H, we have that for any element $x \in K_\mathbb{R}$, the distribution of $x \cdot D_{\mathbf{r}}$ is $D_{\mathbf{r}'}$, where $r'_i = r_i \cdot |\sigma_i(x)|$. (This uses the fact that our distributions have the same variance in the real and imaginary components of each complex embedding.)

Example 2.6. For the mth cyclotomic field where $\zeta = \zeta_m$ for $m \geq 3$, there are $2s_2 = n = \varphi(m)$ complex embeddings (and no real ones), which are given by $\sigma_i(\zeta) = \zeta^i$ for $i \in \mathbb{Z}_m^*$. (In this case it is convenient to index the embeddings σ_i by elements of $\mathbb{Z}_m^*$ instead of $[n]$.) For any power $\zeta^j \in K$, all the embeddings $\sigma_i(\zeta^j) \in \mathbb{C}$ are roots of unity and hence have magnitude 1, so $\|\zeta^j\|_2 = \sqrt{n}$ and $\|\zeta^j\|_\infty = 1$.

2.3.3 Trace and Norm

Abstractly, the (field) *trace* $\mathrm{Tr} = \mathrm{Tr}_{K/\mathbb{Q}}\colon K \to \mathbb{Q}$ and (field) *norm* $\mathrm{N} = \mathrm{N}_{K/\mathbb{Q}}\colon K \to \mathbb{Q}$ of $x \in K$ are the trace and determinant, respectively, of the linear transformation on K (viewed as a vector space over $\mathbb{Q}$) representing multiplication by x. Concretely, the trace and norm can be shown to be the sum and product, respectively, of the embeddings:

$$\mathrm{Tr}(x) = \sum_{i\in[n]} \sigma_i(x) \quad \text{and} \quad \mathrm{N}(x) = \prod_{i\in[n]} \sigma_i(x).$$

Using either definition, it is routine to verify that trace and norm are additive and multiplicative, respectively. Moreover, for all $x, y \in K$,

$$\mathrm{Tr}(x \cdot y) = \sum_{i\in[n]} \sigma_i(x) \cdot \sigma_i(y) = \langle \sigma(x), \overline{\sigma(y)} \rangle.$$

Thus, $\mathrm{Tr}(x \cdot y)$ is a symmetric bilinear form akin to the inner product of the embeddings of x and y.

Example 2.7. Let $\zeta = \zeta_5$ be a root of the cyclotomic polynomial $\Phi_5(x) = \prod_{i \in \mathbb{Z}_5^*}(x - \zeta^i) = x^4 + x^3 + x^2 + x + 1$, and consider the element $y = \frac{1}{2} - \zeta \in K = \mathbb{Q}(\zeta)$. Then $\mathrm{Tr}(y) = \sum_{i \in \mathbb{Z}_5^*}(\frac{1}{2} - \zeta^i) = 2 - (-1) = 3$, and $\mathrm{N}(y) = \prod_{i \in \mathbb{Z}_5^*}(\frac{1}{2} - \zeta^i) = \Phi_5(\frac{1}{2}) = \frac{31}{16}$.

2.3.4 Ring of Integers and Its Ideals

An *algebraic integer* is an element whose minimal polynomial over the rationals has integer coefficients. For a number field K, let $\mathcal{O}_K \subset K$ denote the set of all algebraic integers in K. This set forms a ring (under the usual addition and multiplication operations in K), called the *ring of integers* of the number field. The norm and trace of an algebraic integer are themselves rational integers (i.e., in $\mathbb{Z}$).

It happens that $\mathcal{O}_K$ is a free $\mathbb{Z}$-module of rank n (the degree of K), i.e., it is the set of all $\mathbb{Z}$-linear combinations of some basis $B = \{b_1, \ldots, b_n\} \subset \mathcal{O}_K$. Such a set is called an *integral basis*, and it is also a $\mathbb{Q}$-basis for K (and an $\mathbb{R}$-basis for $K_\mathbb{R}$). As usual, there are infinitely many such bases when $n > 1$.

Example 2.8. Continuing our example of the mth cyclotomic number field $K = \mathbb{Q}(\zeta_m)$ of degree $n = \varphi(m)$, the power basis $\{1, \zeta_m, \ldots, \zeta_m^{n-1}\}$ of K also happens to be an integral basis, i.e., $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$. (In general, it is unusual for the power basis of a number field to generate the entire ring of integers.)

An *(integral) ideal* $\mathcal{I} \subseteq \mathcal{O}_K$ is a nontrivial (i.e., $I \neq \emptyset$ and $\mathcal{I} \neq \{0\}$) additive subgroup that is closed under multiplication by $\mathcal{O}_K$, i.e., $r \cdot x \in \mathcal{I}$ for any $r \in \mathcal{O}_K$ and $x \in \mathcal{I}$.[2] An ideal $\mathcal{I}$ in $\mathcal{O}_K$ is finitely generated as the set of all $\mathcal{O}_K$-linear combinations of some $g_1, g_2, \ldots \in \mathcal{O}_K$, denoted $\mathcal{I} = \langle g_1, g_2, \ldots \rangle$. (In fact, it is known that two generators always suffice.) More useful to us is the fact that an ideal is also a free $\mathbb{Z}$-module of rank n, i.e., it is generated as the set of all $\mathbb{Z}$-linear combinations of some basis $\{u_1, \ldots, u_n\} \subset \mathcal{O}_K$.

The *norm* of an ideal $\mathcal{I}$ is its index as an additive subgroup of $\mathcal{O}_K$, i.e., $\mathrm{N}(\mathcal{I}) = |\mathcal{O}_K/\mathcal{I}|$. The sum $\mathcal{I} + \mathcal{J}$ of two ideals is the set of all $x + y$ for $x \in \mathcal{I}$, $y \in \mathcal{J}$, and the product ideal $\mathcal{I}\mathcal{J}$ is the set of all finite sums of terms xy for $x \in \mathcal{I}$, $y \in \mathcal{J}$. This notion of norm for ideals generalizes the field norm defined above, in the sense that $\mathrm{N}(\langle x \rangle) = |\mathrm{N}(x)|$ for any $x \in \mathcal{O}_K$, and $\mathrm{N}(\mathcal{I}\mathcal{J}) = \mathrm{N}(\mathcal{I})\,\mathrm{N}(\mathcal{J})$.

Two ideals $\mathcal{I}, \mathcal{J} \subseteq \mathcal{O}_K$ are said to be *coprime* (or *relatively prime*) if $\mathcal{I} + \mathcal{J} = \mathcal{O}_K$. An ideal $\mathfrak{p} \subsetneq \mathcal{O}_K$ is *prime* if whenever $ab \in \mathfrak{p}$ for some $a, b \in \mathcal{O}_K$, then $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$ (or both). In $\mathcal{O}_K$, an ideal $\mathfrak{p}$ is prime if and only if it is *maximal*, i.e., if the only proper superideal of $\mathfrak{p}$ is $\mathcal{O}_K$ itself, which implies that the quotient ring $\mathcal{O}_K/\mathfrak{p}$ is the finite field of order $\mathrm{N}(\mathfrak{p})$. The ring $\mathcal{O}_K$ has *unique factorization of ideals*, that is, every ideal $\mathcal{I} \subseteq \mathcal{O}_K$ can be expressed uniquely as a product of powers of prime ideals.

A *fractional ideal* $\mathcal{I} \subset K$ is a set such that $d\mathcal{I} \subseteq \mathcal{O}_K$ is an integral ideal for some $d \in \mathcal{O}_K$. Its norm is defined as $\mathrm{N}(\mathcal{I}) = \mathrm{N}(d\mathcal{I})/|\mathrm{N}(d)|$. The set of fractional ideals form a group under multiplication, and the norm is a multiplicative homomorphism on this group.

2.3.5 Ideal Lattices

Here we recall how (fractional) ideals in K yield lattices under the canonical embedding, and describe some of their properties. Recall that a fractional ideal $\mathcal{I}$ has a $\mathbb{Z}$-basis $U = \{u_1, \ldots, u_n\}$. Therefore, under the canonical embedding σ, the ideal yields a rank-n *ideal lattice* $\sigma(\mathcal{I})$ having basis $\{\sigma(u_1), \ldots, \sigma(u_n)\} \subset H$. For convenience, we often identify an ideal with its embedded lattice, and speak of, e.g., the minimum distance $\lambda_1(\mathcal{I})$ of an ideal, etc.

[2] Some texts also define the trivial set $\{0\}$ as an ideal, but in this work it is more convenient to exclude it.

The (absolute) *discriminant* Δ_K of a number field K is defined to be the square of the fundamental volume of $\sigma(\mathcal{O}_K)$, the embedded ring of integers. Equivalently, $\Delta_K = |\det(\mathrm{Tr}(b_i \cdot b_j))|$ where $b_1, \ldots, b_n$ is any integral basis of $\mathcal{O}_K$.[3] Consequently, the fundamental volume of any ideal lattice $\sigma(\mathcal{I})$ is $\mathrm{N}(\mathcal{I}) \cdot \sqrt{\Delta_K}$. For example, the discriminant of the mth cyclotomic number field $K = \mathbb{Q}(\zeta_m)$ of degree $n = \varphi(m)$ is known to be

$$\Delta_K = \frac{m^n}{\displaystyle\prod_{\text{prime } p|m} p^{n/(p-1)}} \leq n^n,$$

where the product in the denominator runs over all primes p dividing m. The inequality above is tight in the special case when m is a power of two, where $n = m/2$.

The following classical lemma gives upper and lower bounds on the minimum distance of an ideal lattice. The upper bound is an immediate consequence of Minkowski's first theorem; the lower bound follows from the arithmetic mean/geometric mean inequality, and the fact that $|\mathrm{N}(x)| \geq \mathrm{N}(\mathcal{I})$ for any nonzero $x \in \mathcal{I}$. (For a detailed proof, see, e.g., [PR07].)

Lemma 2.9. *For any fractional ideal $\mathcal{I}$ in a number field K of degree n, and in any ℓ_p norm for $p \in [1, \infty]$,*

$$n^{1/p} \cdot \mathrm{N}(\mathcal{I})^{1/n} \;\leq\; \lambda_1(\mathcal{I}) \;\leq\; n^{1/p} \cdot \mathrm{N}(\mathcal{I})^{1/n} \cdot \sqrt{\Delta_K^{1/n}}.$$

2.3.6 Duality

Here we recall the notion of a dual ideal and explain its close connection to both the inverse ideal and the dual lattice. For more details, see [Con09] as an accessible reference.

For any lattice $\mathcal{L}$ in K (i.e., for the $\mathbb{Z}$-span of any $\mathbb{Q}$-basis of K), its *dual* is defined as

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

It is not difficult to see that, under the canonical embedding, $\mathcal{L}^\vee$ embeds as the complex conjugate of the dual lattice, i.e., $\sigma(\mathcal{L}^\vee) = \overline{\sigma(\mathcal{L})^*}$. This is due to the fact that $\mathrm{Tr}(xy) = \sum_i \sigma_i(x)\sigma_i(y) = \langle \sigma(x), \overline{\sigma(y)} \rangle$. It is also easy to check that $(\mathcal{L}^\vee)^\vee = \mathcal{L}$, and that if $\mathcal{L}$ is a fractional *ideal*, then $\mathcal{L}^\vee$ is one as well.

Except in the trivial number field $K = \mathbb{Q}$, the ring of integers $R = \mathcal{O}_K$ is not self-dual, nor are an ideal and its inverse dual to each other. Fortunately, a useful and important fact is that an ideal and its inverse *are* related by multiplication with the dual ideal of the ring: for any fractional ideal $\mathcal{I}$, its dual ideal is $\mathcal{I}^\vee = \mathcal{I}^{-1} \cdot R^\vee$. (Notice that for $\mathcal{I} = R$ this holds trivially, since $R^{-1} = R$.) The factor $R^\vee$ is a fractional ideal whose inverse $(R^\vee)^{-1}$, called the *different ideal*, is integral and of norm $\mathrm{N}((R^\vee)^{-1}) = \Delta_K$. The fractional ideal $R^\vee$ itself is often called the *codifferent*. One especially nice case is the mth cyclotomic number field for $m = 2^k$ of degree $n = \varphi(m) = m/2$, for which $R^\vee = \langle n^{-1} \rangle$ is just a scaling of R.

2.3.7 Computation in Number Fields

We now recall how objects over K and $\mathcal{O}_K$ are represented and operated upon by algorithms, in the general case. For more details, see, e.g., [Coh93]. (Significantly faster algorithms exist for cyclotomic number fields; see the companion paper [LPR13] for details.) When quantifying computational complexity in the context of a number field K, "polynomial" is taken to mean some polynomial in n, $\log \Delta_K$, and the total bit length of any inputs. (In all the concrete families of number fields we use, $\log \Delta_K$ is itself a small polynomial in n.)

[3] In some texts the discriminant is defined as a signed quantity, but in this work we only care about its magnitude.

Because $\mathcal{O}_K$ is a free $\mathbb{Z}$-module of rank n, the ring $\mathcal{O}_K$ can be represented relative to some integral basis $B = \{b_1, \ldots, b_n\} \subset \mathcal{O}_K$, which is also a $\mathbb{Q}$-basis for K. I.e., every element $x \in K$ is represented uniquely by a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Q}^n$ where $x = \sum_{i \in [n]} x_i \cdot b_i$, and $x \in \mathcal{O}_K$ if and only if every $x_i \in \mathbb{Z}$.

Addition in K (and $\mathcal{O}_K$) is computed simply by component-wise addition of the representation vectors of the addends. For computing the multiplication operation in K (and $\mathcal{O}_K$), it suffices by linearity to know each product $b_i \cdot b_j \in \mathcal{O}_K$ for $i, j \in [n]$. The representations of these terms (as usual, with respect to B) are integral vectors of polynomial size, and constitute the entire description of K and $\mathcal{O}_K$. Given this description, it is possible to compute multiplicative inverses in polynomial time.

An integral ideal $\mathcal{I} \subseteq \mathcal{O}_K$ is represented by some $\mathbb{Z}$-basis for $\mathcal{I}$, that is, a set $U_{\mathcal{I}} = \{u_1, \ldots, u_n\} \subset \mathcal{O}_K$ such that $\mathcal{I} = \sum_{i \in [n]} \mathbb{Z} \cdot u_i$ (where as always, each u_i is represented relative to a fixed integral basis B). A fractional ideal $\mathcal{I}$ is represented by also including a denominator $d \in \mathcal{O}_K$ such that $d \cdot \mathcal{I}$ is an integral ideal. With these representations, in deterministic polynomial time it is possible to check that a given basis generates an ideal $\mathcal{I}$, to compute the norm of $\mathcal{I}$, to compute the inverse ideal $\mathcal{I}^{-1}$ and dual ideal $\mathcal{I}^{\vee}$, to reduce a given element in K modulo a given basis of $\mathcal{I}$, and to compute the Hermite normal form (HNF) of $\mathcal{I}$ along with the unimodular integer matrix relating the HNF to the input basis. Given two ideals $\mathcal{I}, \mathcal{J}$, it is possible to compute the product ideal $\mathcal{I}\mathcal{J}$ in deterministic polynomial time, and if $\mathcal{J} \subseteq \mathcal{I}$, to select a uniformly random element from the quotient group $\mathcal{I}/\mathcal{J}$ in polynomial time, and to enumerate $\mathcal{I}/\mathcal{J}$ in deterministic polynomial time per element.

Recall from Section 2.2.1 that the elliptical Gaussian distribution $D_{\mathbf{r}}$ corresponds (under the canonical embedding) to the sum of independent one-dimensional Gaussian multiples of the orthonormal basis vectors $\mathbf{h}_i$ for H. Therefore, it is possible to sample in polynomial time from $D_{\mathbf{r}}$ over $K_{\mathbb{R}}$ (up to any desired precision), given $\mathbf{r}$ and the representations of each $\mathbf{h}_i$ relative to B. (Equivalently, it is enough to know $\sigma(B)$, the embedding of the power basis into H.)

2.3.8 Ideal Lattice Problems

The following are the three main (seemingly hard) computational problems on ideal lattices that we deal with in this work: the *Shortest Vector Problem* (SVP), *Shortest Independent Vectors Problem* (SIVP), and the *Bounded-Distance Decoding* (BDD) problem. Without loss of generality, they may be restricted to *integral* ideals in $\mathcal{O}_K$, by the following scaling argument: if $\mathcal{I}$ is a fractional ideal with denominator $d \in \mathcal{O}_K$ (such that $d\mathcal{I} \subseteq \mathcal{O}_K$ is an integral ideal), then the scaled ideal $\mathrm{N}(d) \cdot \mathcal{I} \subseteq \mathcal{O}_K$, because $\mathrm{N}(d) \in \langle d \rangle$.

Definition 2.10 (SVP **and** SIVP)**.** *Let K be a number field endowed with some geometric norm (e.g., the ℓ_2 norm), and let $\gamma \geq 1$. The K-SVP_γ problem in the given norm is: given a fractional ideal $\mathcal{I}$ in K, find some nonzero $x \in \mathcal{I}$ such that $\|x\| \leq \gamma \cdot \lambda_1(\mathcal{I})$. The K-SIVP_γ problem is defined similarly, where the goal is to find n linearly independent elements in $\mathcal{I}$ whose norms are all at most $\gamma \cdot \lambda_n(\mathcal{I})$.*

Definition 2.11 (BDD)**.** *Let K be a number field endowed with some geometric norm (e.g., the ℓ_∞ norm), let $\mathcal{I}$ be a fractional ideal in K, and let $d < \lambda_1(\mathcal{I})/2$. The K-$\mathsf{BDD}_{\mathcal{I},d}$ problem in the given norm is: given $\mathcal{I}$ and y of the form $y = x + e$ for some $x \in \mathcal{I}$ and $\|e\| \leq d$, find x.*

2.3.9 Chinese Remainder Theorem

Here we recall the Chinese Remainder Theorem (CRT) for the ring of integers $R = \mathcal{O}_K$ in a number field K, and some of its important consequences for our work.

Lemma 2.12 (Chinese Remainder Theorem). *Let $\mathcal{I}_1, \ldots, \mathcal{I}_r$ be pairwise coprime ideals in R, and let $\mathcal{I} = \prod_{i \in [r]} \mathcal{I}_i$. The natural ring homomorphism $R \to \bigoplus_{i \in [r]} (R/\mathcal{I}_i)$ induces a ring isomorphism $R/\mathcal{I} \to \bigoplus_{i \in [r]} (R/\mathcal{I}_i)$.*

The following lemma shows that we can efficiently compute a "CRT basis" C for any set of pairwise coprime ideals $\mathcal{I}_1, \ldots, \mathcal{I}_r$, i.e., elements $c_1, \ldots, c_r \in R$ such that $c_i = 1 \bmod \mathcal{I}_i$ and $c_i = 0 \bmod \mathcal{I}_j$ for all $i \neq j$. Such a basis allows us to invert the isomorphism described in Lemma 2.12, as follows: for any given $w = (w_1, \ldots, w_r) \in \bigoplus_i (R/\mathcal{I}_i)$, the value $v = \sum_i w_i \cdot c_i \bmod \mathcal{I}$ is the unique element in $R/\mathcal{I}$ that maps to w under the ring isomorphism.

Lemma 2.13. *There is a deterministic polynomial-time algorithm that, given coprime ideals $\mathcal{I}, \mathcal{J} \subseteq R$ (represented by $\mathbb{Z}$-bases), outputs some $c \in \mathcal{J}$ such that $c = 1 \bmod \mathcal{I}$. More generally, there is a deterministic polynomial-time algorithm that, given pairwise coprime ideals $\mathcal{I}_1, \ldots, \mathcal{I}_r$, outputs a CRT basis $c_1, \ldots, c_r \in R$ for those ideals.*

Proof. The algorithm is a generalization of the extended Euclidean algorithm for the integers $\mathbb{Z}$. It works as follows: given arbitrary $\mathbb{Z}$-bases $B_{\mathcal{I}}$ and $B_{\mathcal{J}}$ for $\mathcal{I}$ and $\mathcal{J}$ respectively, let $B = B_{\mathcal{I}} \cup B_{\mathcal{J}}$ be the (possibly overdetermined) basis for $\mathcal{I} + \mathcal{J} = R$. Compute from B the Hermite normal form basis $H = \{h_1, \ldots, h_n\}$ for R, yielding an expression of each h_i as a $\mathbb{Z}$-combination of elements in B. Write the element $1 \in R$ as a $\mathbb{Z}$-combination of elements in H, and hence as a $\mathbb{Z}$-combination of elements in $B_{\mathcal{I}} \cup B_{\mathcal{J}}$. From this we get elements $x \in \mathcal{I}$, $y \in \mathcal{J}$ such that $x + y = 1$. Finally, output the element $c = y = 1 - x \in \mathcal{J}$, which is 1 modulo $\mathcal{I}$. For the second part of the claim, to compute each c_i we simply let $\mathcal{I} = \mathcal{I}_i$ and $\mathcal{J} = \prod_{j \neq i} \mathcal{I}_j$. □

The next two lemmas combine to give an efficiently computable bijection (and moreover, an isomorphism of R-modules) between the quotient groups $\mathcal{I}/q\mathcal{I}$ and $\mathcal{J}/q\mathcal{J}$ for any fractional ideals $\mathcal{I}, \mathcal{J}$. This will be an important tool for 'clearing out' the arbitrary ideal $\mathcal{I}$ in our BDD-to-LWE reduction in Section 4.2 (and more generally, in other worst-case to average-case reductions for ideal lattices). We note that these lemmas are probably standard for experts in computational number theory, but we believe their application in our context is new.

Lemma 2.14. *Let $\mathcal{I}$ and $\mathcal{J}$ be ideals in R. There exists $t \in \mathcal{I}$ such that the ideal $t \cdot \mathcal{I}^{-1} \subseteq R$ is coprime to $\mathcal{J}$. Moreover, such t can be found efficiently given $\mathcal{I}$ and the prime ideal factorization of $\mathcal{J}$.*

Proof. Let $\mathfrak{p}_1, \ldots, \mathfrak{p}_r$ be the given prime divisors of $\mathcal{J}$. For each $i \in [r]$, let $e_i \geq 0$ be the largest power of $\mathfrak{p}_i$ that divides $\mathcal{I}$. (Such e_i can be computed efficiently by trial division and binary search, because e_i cannot exceed $\log \mathrm{N}(\mathcal{I}) / \log \mathrm{N}(\mathfrak{p}_i)$.) For each $i \in [r]$, choose an arbitrary $t_i \in \mathfrak{p}_i^{e_i}$ that is not in $\mathfrak{p}_i^{e_i+1}$. By the Chinese Remainder Theorem, there exists $t \in R$ such that

$$t = 0 \bmod \Big(\mathcal{I} / \prod\nolimits_{i \in [r]} \mathfrak{p}_i^{e_i}\Big) \quad \text{and} \quad \forall i \in [r],\ t = t_i \bmod \mathfrak{p}_i^{e_i+1}.$$

(Note that the ideals in question are pairwise coprime.) Moreover, such t can be found efficiently using a CRT basis for the ideals $\mathfrak{p}_i^{e_i+1}$ and $\mathcal{I}/\prod_i \mathfrak{p}_i^{e_i}$. Now because t is 0 modulo every $\mathfrak{p}_i^{e_i}$, it follows that $t \in \mathcal{I}$.

To finish, we need to show that $t \cdot \mathcal{I}^{-1}$ is not divisible by any $\mathfrak{p}_i$. Supposing to the contrary implies that $\mathfrak{p}_i \mathcal{I} | \langle t \rangle$, and since $\mathfrak{p}_i^{e_i+1} | \mathfrak{p}_i \mathcal{I}$ we have $t \in \mathfrak{p}_i^{e_i+1}$. But $t = t_i \neq 0 \bmod \mathfrak{p}_i^{e_i+1}$, a contradiction. □

Upon first reading, in the following lemma the reader may wish to think of the ideal $\mathcal{M}$ as the multiplicative identity (i.e., the entire ring R).

Lemma 2.15. *Let $\mathcal{I}$ and $\mathcal{J}$ be ideals in R, let $t \in \mathcal{I}$ be such that $t \cdot \mathcal{I}^{-1}$ is coprime with $\mathcal{J}$, and let $\mathcal{M}$ be any fractional ideal in K. Then the function $\theta_t \colon K \to K$ defined as $\theta_t(u) = t \cdot u$ induces an isomorphism from $\mathcal{M}/\mathcal{J}\mathcal{M}$ to $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$, as R-modules. Moreover, this isomorphism may be efficiently inverted given $\mathcal{I}$, $\mathcal{J}$, $\mathcal{M}$, and t.*

Proof. That θ_t induces a homomorphism of R-modules follows immediately from the fact that it represents multiplication by a fixed $t \in R$.

Now consider the function induced by θ_t having domain $\mathcal{M}$ and range $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$. Its kernel is $\mathcal{J}\mathcal{M}$, which may be seen as follows: first, $\theta_t(\mathcal{J}\mathcal{M}) = t \cdot \mathcal{J}\mathcal{M} \subseteq \mathcal{I}\mathcal{J}\mathcal{M}$. Second, if $\theta_t(u) = 0$ for some $u \in \mathcal{M}$, then $t \cdot u \in \mathcal{I}\mathcal{J}\mathcal{M}$ which implies $(t \cdot \mathcal{I}^{-1}) \cdot (u \cdot \mathcal{M}^{-1}) \subseteq \mathcal{J}$. Because $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$ are coprime ideals in R, we have $u \cdot \mathcal{M}^{-1} \subseteq \mathcal{J} \Rightarrow u \in \mathcal{J}\mathcal{M}$. So the function from $\mathcal{M}/\mathcal{J}\mathcal{M}$ to $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$ induced by θ_t is injective. It remains to show that it is surjective. (Actually, surjectivity follows immediately from the fact that both quotient groups have cardinality $\mathrm{N}(\mathcal{J})$, but below we give a constructive proof that also demonstrates efficient invertibility.)

Let $v \in \mathcal{I}\mathcal{M}$ be arbitrary. By hypothesis, $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$ are coprime, so we can use the algorithm from Lemma 2.13 to compute some $c \in t \cdot \mathcal{I}^{-1}$ such that $c = 1 \bmod \mathcal{J}$. Then let $a = c \cdot v \in t \cdot \mathcal{M}$, and observe that $a - v = v \cdot (c - 1) \in \mathcal{I}\mathcal{J}\mathcal{M}$. Let $w = a/t \in \mathcal{M}$; then $\theta_t(w) = t \cdot (a/t) = v \bmod \mathcal{I}\mathcal{J}\mathcal{M}$, so $w \bmod \mathcal{J}\mathcal{M}$ is the preimage of $v \bmod \mathcal{I}\mathcal{J}\mathcal{M}$. □

2.4 Special Properties of Cyclotomic Number Fields

Here we recall a few important facts about cyclotomic number fields. We use these in our search-to-decision reductions of Section 5. Let $K = \mathbb{Q}(\zeta)$ for $\zeta = \zeta_m$ be the mth cyclotomic number field, which has minimal polynomial $\Phi_m(x)$ of degree $n = \varphi(m)$, and let $R = \mathcal{O}_K = \mathbb{Z}[\zeta]$.

The number field K has n automorphisms $\tau_k : K \to K$, which are defined by $\tau_k(\zeta) = \zeta^k$ for $k \in \mathbb{Z}_m^*$. It is easy to verify that τ_k is a ring homomorphism and is invertible, hence it is an automorphism. An important fact is that each automorphism permutes the coordinates of the canonical embedding. More precisely, $\sigma_i(\tau_k(\zeta)) = \sigma_{ik}(\zeta)$ for any $i, k \in \mathbb{Z}_m^*$, so $\sigma \circ \tau_k$ is just σ under a fixed permutation of its coordinates.

For an integer prime $q \in \mathbb{Z}$, the factorization of the ideal $\langle q \rangle = qR$ is as follows. Let q' be the largest power of q that divides m, let $e = \varphi(q')$, and let f be the multiplicative order of q modulo m/q'. Then $\langle q \rangle = \prod_i \mathfrak{q}_i^e$, where the $\mathfrak{q}_i$ are $n/(ef)$ distinct prime ideals, each of norm q^f. Concretely, these ideals are given by $\mathfrak{q}_i = \langle q, F_i(\zeta) \rangle$, where $\Phi_m(x) = \prod_i (F_i(x))^e$ is the factorization of the cyclotomic polynomial $\Phi_m(x)$ modulo q (i.e., in $\mathbb{Z}_q[x]$) into monic irreducible polynomials $F_i(x)$. (Note that this factorization can be computed efficiently [Sho09, Chapter 20].)

In particular, for an integer prime q congruent to 1 modulo m, we have $e = f = 1$, the field $\mathbb{Z}_q$ has a primitive mth root of unity ω (because the multiplicative group of $\mathbb{Z}_q$ is cyclic with order $q - 1$), and so $\Phi_m(x)$ factors in $\mathbb{Z}_q[x]$ as $\Phi(x) = \prod_{i \in \mathbb{Z}_m^*} (x - \omega^i)$. The ideal $\langle q \rangle$ then "splits completely" into n distinct prime ideals, as $\langle q \rangle = \prod_{i \in \mathbb{Z}_m^*} \mathfrak{q}_i$ where $\mathfrak{q}_i = \langle q, \zeta - \omega^i \rangle$ is prime and has norm q.

The following lemma says that the automorphisms τ_k "act transitively" on the prime ideals $\mathfrak{q}_i$, i.e., each $\mathfrak{q}_i$ is sent to each $\mathfrak{q}_j$ by some automorphism τ_k. We note that the lemma follows directly from the fact that cyclotomic number fields are *Galois extensions* of $\mathbb{Q}$ (see, e.g., [Ste04, Chapter 13]); here we give an elementary proof for completeness.

Lemma 2.16. *Using the notation above, for any $i, j \in \mathbb{Z}_m^*$ we have $\tau_j(\mathfrak{q}_i) = \mathfrak{q}_{i/j}$.*

Proof. By definition, $\tau_j(\mathfrak{q}_i) = \tau_j(\langle q, \zeta - \omega^i \rangle) = \langle q, \zeta^j - \omega^i \rangle$. Now observe that

$$\zeta^j - \omega^i = \zeta^j - (\omega^{i/j})^j = (\zeta - \omega^{i/j}) \cdot \left(\zeta^{j-1} + \omega^{i/j} \cdot \zeta^{j-2} + \cdots + (\omega^{i/j})^{j-1}\right) \bmod q,$$

where the summation on the right is in $\mathcal{O}_K$. Thus $\tau_j(\mathfrak{q}_i) \subseteq \langle q, \zeta - \omega^{i/j} \rangle = \mathfrak{q}_{i/j}$.

For the reverse inclusion, note that $\zeta - \omega^{i/j} = (\zeta^j)^{1/j} - (\omega^i)^{1/j}$ factors similarly (modulo q) as a multiple of $\zeta^j - \omega^i$, and therefore $\mathfrak{q}_{i/j} \subseteq \langle q, \zeta^j - \omega^i \rangle = \tau_j(\mathfrak{q}_i)$. □

3 The Ring-LWE Problem and Main Results

Here we define the ring-LWE distribution (actually, family of distributions) and the main computational problems associated with it. Ring-LWE is parameterized by a number field K with ring of integers $R = \mathcal{O}_K$ and a (rational) integer modulus $q \geq 2$. For any fractional ideal $\mathcal{J}$ in K, we let $\mathcal{J}_q$ denote $\mathcal{J}/q\mathcal{J}$. Recall that $R^\vee$ is the dual (or "codifferent") fractional ideal of R, and let $\mathbb{T} = K_\mathbb{R}/R^\vee$.

Definition 3.1 (Ring-LWE Distribution). *For $s \in R_q^\vee$ (the "secret") and an error distribution ψ over $K_\mathbb{R}$, a sample from the ring-*LWE *distribution $A_{s,\psi}$ over $R_q \times \mathbb{T}$ is generated by choosing $a \leftarrow R_q$ uniformly at random, choosing $e \leftarrow \psi$, and outputting $(a, b = (a \cdot s)/q + e \bmod R^\vee)$.*

Note that $(a \cdot s)/q \in \frac{1}{q}R^\vee/R^\vee$, so the reduction modulo $R^\vee$ in the second component of the sample is well defined.

Definition 3.2 (Ring-LWE, Search). *Let Ψ be a family of distributions over $K_\mathbb{R}$. The* search *version of the ring-*LWE *problem, denoted R-$\mathsf{LWE}_{q,\Psi}$, is defined as follows: given access to arbitrarily many independent samples from $A_{s,\psi}$ for some arbitrary $s \in R_q^\vee$ and $\psi \in \Psi$, find s.*

The following *decision* form of the problem, whose hardness means (informally) that the ring-LWE distribution is pseudorandom, is usually more suitable for cryptographic applications. In Section 5, we show that it is in fact equivalent to the search version, under certain conditions on the parameters.

Definition 3.3 (Ring-LWE, Average-Case Decision). *Let Υ be a distribution over a family of error distributions, each over $K_\mathbb{R}$. The* average-case decision *version of the ring-*LWE *problem, denoted R-$\mathsf{DLWE}_{q,\Upsilon}$, is to distinguish with non-negligible advantage between arbitrarily many independent samples from $A_{s,\psi}$, for a* random *choice of $(s, \psi) \leftarrow U(R_q^\vee) \times \Upsilon$, and the same number of uniformly random and independent samples from $R_q \times \mathbb{T}$.*

For an asymptotic treatment of the ring-LWE problems, we let K come from an infinite sequence of number fields $\mathcal{K} = \{K_n\}$ of increasing dimension n, and let q, Ψ, and Υ depend on n as well.

Recall that when informally describing ring-LWE in the introduction, we said that the secret s belongs to R_q (and so the products $a \cdot s$ are also in R_q), whereas in the formal definition above, s is in $R_q^\vee$ (and so $a \cdot s \in R_q^\vee$). Since the description in the introduction was specialized to the mth cyclotomic ring for $m = 2^k$ (i.e., $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ for $n = 2^{k-1}$), these two variants are actually equivalent. This follows from the fact that for this ring the codifferent $R^\vee = n^{-1}R$ is simply a scaling of the ring, and so one can transform samples $(a, b = (a \cdot s)/q + e)$ to $(a, b' = b \cdot n = (a \cdot s')/q + e')$, where $s' = s \cdot n \in R_q$ and $e' = e \cdot n \in R$. More generally, it is usually more appropriate to leave the secret in $R_q^\vee$, as in our formal definition above; see Section 3.3 below for further discussion.

3.1 Error Distributions

We first define the family of LWE error distributions for which our reduction to the search version of ring-LWE (in Section 4) applies.

Definition 3.4. *For a positive real $\alpha > 0$, the family $\Psi_{\leq \alpha}$ is the set of all elliptical Gaussian distributions $D_{\mathbf{r}}$ (over $K_{\mathbb{R}}$) where each parameter $r_i \leq \alpha$.*

Our hardness results for the average-case decision problem (in Section 5) apply to cyclotomic number fields $K = \mathbb{Q}(\zeta_m)$; recall that these have zero real embeddings and $n/2 = \varphi(m)/2$ pairs of complex embeddings, $\sigma_i = \overline{\sigma_{i+n/2}}$ for $i \in [n/2]$. Our results use a distribution Υ over error distributions, defined as follows. The gamma distribution $\Gamma(2,1)$ with shape parameter 2 and scale parameter 1 has density given by $x\exp(-x)$ for $x \geq 0$, and zero for $x < 0$. Sampling from this gamma distribution can be done efficiently by, e.g., sampling two uniform variables U_1, U_2 in $[0,1]$ and outputting $-\ln U_1 - \ln U_2$. Other equally good choices are possible (e.g., a Gaussian distribution) and we make this particular choice for convenience.

Definition 3.5. *Let K be the mth cyclotomic number field having degree $n = \varphi(m)$. For a positive real $\alpha > 0$, a distribution sampled from Υ_α is given by an elliptical Gaussian distribution $D_{\mathbf{r}}$ (over $K_{\mathbb{R}}$) whose parameters are $r_i^2 = r_{i+n/2}^2 = \alpha^2(1 + \sqrt{n}x_i)$, where $x_1, \ldots, x_{n/2}$ are chosen independently from the distribution $\Gamma(2,1)$.*

Notice that error distributions drawn from Υ_α typically have parameters of size roughly $O(\alpha \cdot n^{1/4})$.

It is important to keep in mind that in our definition of ring-LWE, the error distribution is added modulo $R^\vee$. As a result, in order for the problem not to be trivially impossible to solve, the error must not exceed the smoothing parameter of $R^\vee$, or else the ring-LWE distribution will be statistically indistinguishable from uniform (for any value of the secret s). For example, in the case of a cyclotomic $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$ for $n = 2^k$, we have $\lambda_n(R^\vee) = 1/\sqrt{n}$ and so by Lemma 2.2 we obtain an upper bound of $O(\sqrt{\log n/n})$ on the error parameters. This is in contrast to standard LWE, where the error is added modulo $\mathbb{Z}$ and hence can be as large as $O(\sqrt{\log n})$.

3.2 Main Theorem

We can now finally state our main theorem, obtained by combining Theorem 4.1 with Theorems 5.1 and 5.2. We note that each of these component theorems, as well as Theorem 5.3, should be of independent interest.

Theorem 3.6. *Let K be the mth cyclotomic number field having dimension $n = \varphi(m)$ and $R = \mathcal{O}_K$ be its ring of integers. Let $\alpha < \sqrt{\log n/n}$, and let $q = q(n) \geq 2$, $q = 1 \bmod m$ be a $\mathrm{poly}(n)$-bounded prime such that $\alpha q \geq \omega(\sqrt{\log n})$. Then there is a polynomial-time quantum reduction from $\tilde{O}(\sqrt{n}/\alpha)$-approximate* SIVP *(or* SVP*) on ideal lattices in K to R-$\mathsf{DLWE}_{q,\Upsilon_\alpha}$. Alternatively, for any $\ell \geq 1$, we can replace the target problem by the problem of solving R-DLWE_{q,D_ξ} given only ℓ samples, where $\xi = \alpha \cdot (n\ell/\log(n\ell))^{1/4}$.*

Notice that in the latter reduction we use DLWE with a fixed (spherical) error distribution D_ξ, instead of a distribution over error distributions. Also note that when ℓ is small, which is often the case in applications, both reductions lead to essentially the same error parameters. See Section 5 for further discussion.

For applications, it is typically more useful to work with a 'discrete' variant of ring-LWE in which the b component of each sample is taken from a finite set, instead of the continuous domain $\mathbb{T}$. (Indeed, this is how we described the ring-LWE problem in the introduction.) As in standard LWE, this is easily achieved by discretizing the samples. See [LPR13, Section 2.6] for details.

3.3 Why This is the Right Definition of Ring-LWE

Our definition of ring-LWE and hardness results make three, seemingly arbitrary, choices: the public value a is drawn from R_q, the secret s belongs to $R_q^\vee$, and the error distributions are "spherically bounded," namely, the bound α in the definition of $\Psi_{\leq\alpha}$ is the same for all coordinates of the embedding (and similarly for the error distributions used in the average-case problem). Indeed, there is some arbitrariness in the definition: for instance, taking $a \in \mathcal{J}_q$ and $s \in \mathcal{J}_q^\vee$ for any fixed fractional ideal $\mathcal{J}$ (and keeping the same error distribution) leads to a computationally equivalent problem.[4] For example, we could take a from $R_q^\vee$ and s from R_q.

This brings us to the following natural question: modulo these equivalences, why are our choices the "right" ones? In particular, why not take both a and s in R_q, and use a spherically bounded error distribution? As described earlier, in certain cases (e.g., the mth cyclotomic for $m = 2^k$) this definition is equivalent to the original one, but it turns out that in general it leads to a pure loss in two ways: once in the provable hardness of the problem for a given spherical error bound (due to the inherent distortion in mapping $R^\vee$ to R), and again in the amount of error that can be used in typical applications. By contrast, our choices turn out to be the most mathematically natural and computationally effective ones, for a variety of reasons that we now explain.

The first reason is that our choices emerge naturally in the core BDD-to-LWE reduction of Section 4, due to the $R^\vee$ ratio between the dual and inverse ideals (recall that $\mathcal{I}^\vee = R^\vee \cdot \mathcal{I}^{-1}$ for any fractional ideal $\mathcal{I}$). In more detail, in our reduction the secret s corresponds to the unknown closest point $x \in \mathcal{I}^\vee$ of a BDD instance, and the public elements a are obtained from Gaussian samples over $\mathcal{I}$. After 'clearing the ideal' $\mathcal{I}$ (using Lemma 2.15), we end up with secret $s \in R_q^\vee$, public elements $a \in R_q$, and spherically bounded error from the products of the spherical Gaussian samples with the BDD offset vector.

Another important reason for our choices relates to cryptographic applications in cyclotomic rings. Fixing encryption as an example, we wish to have as much security as possible, which by our results means using as much spherically bounded error as possible, yet still decrypt correctly, which requires removing the error from noisy products using knowledge of the secret. For our choices, this means solving BDD for spherically bounded error on the ideal $R^\vee$. The amount of spherical error that can be efficiently decoded from an ideal essentially depends inversely on the λ_n parameter of its dual ideal. The dual of $R^\vee$ is R, which, for cyclotomic rings of degree n, has n elements $1, \zeta, \ldots, \zeta^{n-1}$ of Euclidean norm $\sqrt{n}$, and hence $\lambda_n(R) = \sqrt{n}$. This turns out to be the smallest possible (relative to its norm), by Lemma 2.9. Therefore, among all ideals of the same norm, the codifferent $R^\vee$ is decodable under the most spherical Gaussian error. By contrast, for most cyclotomics this is not true of R itself, because its dual ideal $R^\vee$ typically does not have optimally short vectors.[5]

A final reason for our choice is by analogy to the standard LWE and SIS (short integer solution) problems, which may be seen as problems on dual random lattices (see, e.g., [GPV08]). In the ring setting, an instance of ring-SIS is given by a tuple $\mathbf{a} = (a_1, \ldots, a_m) \in R_q^m$, which defines the "$q$-ary" lattice $\Lambda^\perp(\mathbf{a})$:

$$qR^m \quad \subseteq \quad \Lambda^\perp(\mathbf{a}) = \Big\{ \mathbf{z} = (z_1, \ldots, z_m) \in R^m \ : \ \textstyle\sum_{i\in[m]} a_i z_i = 0 \in R_q \Big\} \quad \subseteq \quad R^m.$$

Analogously to standard LWE, we wish to view ring-LWE as a bounded-distance decoding problem on the

[4] This follows by the 'clearing ideals' technique of Lemma 2.15. We note that standard LWE also admits an analogous formulation: for any n-dimensional lattice L and its dual L^*, we can take $\mathbf{a} \in L/qL$, $\mathbf{s} \in L^*/qL^*$, and $b \approx \langle \mathbf{a}, \mathbf{s} \rangle / q \bmod 1$. In that setting, the self-dual lattice $L = \mathbb{Z}^n$ is of course the most natural choice.

[5] In some cryptographic applications, several ring-LWE samples are combined, requiring decoding on ideals derived from the one used in the original ring-LWE samples (e.g., starting from our $R^\vee$, one needs to decode $(R^\vee)^k$ for some $k > 1$). The same considerations discussed here reveal that $R^\vee$ is essentially optimal also in these applications. See [LPR13] for further details.

dual lattice of $\Lambda^\perp(\mathbf{a}) \subseteq R^m$, under the canonical embedding.[6] It can be seen that this dual lattice is

$$(\Lambda^\perp(\mathbf{a}))^\vee = (R^\vee)^m + \big\{(\mathbf{a}\cdot s)/q \;:\; s \in R_q^\vee\big\} \quad \subseteq \quad (1/q)\cdot(R^\vee)^m,$$

where the product $\mathbf{a}\cdot s = (a_1 \cdot s, \ldots, a_m \cdot s)$. Therefore, it is natural that the ring-LWE distribution should produce noisy products of the form $(a_i \cdot s)/q$, for random values $a_i \in R_q$ and some fixed $s \in R_q^\vee$.

Finally, we remark that working with $R^\vee$ and $R_q^\vee$ is computationally just as efficient as working with R and R_q, by the Chinese Remainder Theorem and its consequences. (See the companion paper [LPR13] for details.)

4 Hardness of Search-LWE

Throughout this section, let K denote an arbitrary number field of degree n with ring of integers $R = \mathcal{O}_K$. For concreteness, the reader may wish to keep in mind the particular case of a cyclotomic number field, although the results in this section apply to arbitrary number fields. The following is the main theorem of this section. (Throughout this section, $\omega(\sqrt{\log n})$ denotes some fixed, arbitrary function that grows asymptotically faster than $\sqrt{\log n}$.)

Theorem 4.1. *Let K be an arbitrary number field of degree n and $R = \mathcal{O}_K$. Let $\alpha = \alpha(n) > 0$, and let $q = q(n) \geq 2$ be such that $\alpha q \geq 2 \cdot \omega(\sqrt{\log n})$. For some negligible $\varepsilon = \varepsilon(n)$, there is a probabilistic polynomial-time quantum reduction from K-DGS_γ to R-$\mathsf{LWE}_{q,\Psi_{\leq\alpha}}$, where*

$$\gamma = \max\Big\{\eta_\varepsilon(\mathcal{I}) \cdot (\sqrt{2}/\alpha)\cdot \omega(\sqrt{\log n}),\ \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)\Big\}. \tag{4.1}$$

Here, K-DGS_γ denotes the *discrete Gaussian sampling* problem [Reg05], which asks, given an ideal $\mathcal{I}$ in K and a number $s \geq \gamma = \gamma(\mathcal{I})$, to produce samples from the distribution $D_{\mathcal{I},s}$. Using the easy inequality $\eta_\varepsilon(\mathcal{I}) > 1/\lambda_1(\mathcal{I}^\vee)$ [Reg05, Claim 2.13], we get that as long as $\alpha < \sqrt{\log n/n}$ (which is virtually always the case in applications), the first term in the maximum in Equation (4.1) dominates.

As shown in [Reg05, Section 3.3], there are easy reductions from standard lattice problems to DGS. Namely, using the facts that $\eta_\varepsilon(\mathcal{I}) \leq \lambda_n(\mathcal{I}) \cdot \omega(\sqrt{\log n})$ for any fractional ideal $\mathcal{I}$ and some negligible $\varepsilon(n)$ (Lemma 2.2), and that a sample from $D_{\mathcal{I},\gamma}$ has length at most $\gamma\sqrt{n}$ with overwhelming probability (Lemma 2.4), an oracle for K-DGS_γ with $\gamma = \eta_\varepsilon(\mathcal{I}) \cdot \tilde{O}(1/\alpha)$ immediately implies an oracle for $\tilde{O}(\sqrt{n}/\alpha)$-approximate SIVP on ideal lattices in K. In cyclotomic number fields, where $\lambda_n(\mathcal{I}) = \lambda_1(\mathcal{I})$ for any fractional ideal $\mathcal{I}$ (because multiplying a shortest nonzero element $v \in \mathcal{I}$ by $1, \zeta, \ldots, \zeta^{n-1}$ gives n linearly independent elements of the same length), this also implies an oracle for $\tilde{O}(\sqrt{n}/\alpha)$-approximate SVP.

Meaningful error rates and approximation factors. As mentioned in Section 3.1, in order for ring-LWE to be information theoretically solvable, it is necessary that $\alpha < \eta_\varepsilon(R^\vee)$ for all negligible $\varepsilon(n)$ (otherwise, ring-LWE samples will just be essentially uniform). Moreover, it can be shown that a sufficient condition for solvability is that $\alpha \leq \lambda_1(R^\vee)/(C\sqrt{n})$ for some universal constant $C > 0$; this is based on the fact that a sample from D_α has norm at most $\alpha\sqrt{n}$ with overwhelming probability.

For example, consider the mth cyclotomic number field, which has degree $n = \varphi(m)$. When m is a power of two, we have $\lambda_n(R^\vee) = 1/\sqrt{n}$, so by Lemma 2.2 we need $\alpha = O(\sqrt{\log n/n})$, which also turns out to be sufficient for information-theoretic solvability. This corresponds to superlinear $\omega(n)$ approximation

[6] As expected, the dual $\Lambda^\vee$ of a lattice $\Lambda \subset K^m$ is the set of all $\mathbf{x} \in K^m$ such that $\sum_{i\in[m]} \operatorname{Tr}(x_i \cdot v_i) \in \mathbb{Z}$ for all $\mathbf{v} \in \Lambda$.

factors for SVP and SIVP. For arbitrary m, one can show the slightly looser bound $\lambda_n(R^\vee) \leq 2^d \sqrt{n}/m$, where d is the number of distinct prime divisors of m.

As a second example, consider any family of number fields K of increasing dimension n that have discriminants $\Delta_K = 2^{\Theta(n)}$ (see, e.g., [Roq67]). Then $\lambda_1(R^\vee) = \Theta(\sqrt{n})$ by Lemma 2.9, so we may take $\alpha > 0$ to be as large as a small constant while still ensuring solvability. Now for any fractional ideal $\mathcal{I}$ in K, we have $\lambda_1(\mathcal{I}) \cdot \lambda_1(\mathcal{I}^\vee) = \Theta(\sqrt{n})$ by Lemma 2.9, so $\eta_\varepsilon(\mathcal{I}) = \lambda_1(\mathcal{I})/\Theta(\sqrt{n})$ for $\varepsilon = 2^{-\Omega(n)}$ by Lemma 2.2. Therefore, we can take $\gamma = \omega(\sqrt{\log n})$ and obtain worst-case approximation factors as small as $\omega(\sqrt{\log n})$ for SVP and SIVP. This essentially matches the $O(\sqrt{\log n})$ approximation factors obtained in [PR07] for the ring-SIS problem in the same families of number fields.

4.1 Proof of Theorem 4.1

The proof of Theorem 4.1 follows Regev's proof in [Reg05] for general lattices, replacing its core component with an analogous statement for ideal lattices (Lemma 4.3). For completeness, we now describe the reduction in some detail, focusing on the necessary modifications. The reduction works by repeated applications of the following *iterative step*.

Lemma 4.2. *Let $\alpha > 0$ and $q \geq 2$ be an integer. There exists an efficient quantum algorithm that, given a fractional ideal $\mathcal{I}$ in K, a number $r \geq \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I})$ for some negligible $\varepsilon = \varepsilon(n)$ such that $r' := r \cdot \omega(\sqrt{\log n})/(\alpha q) > \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)$, an oracle to R-$\mathsf{LWE}_{q,\Psi_{\leq\alpha}}$, and a list of samples from the discrete Gaussian distribution $D_{\mathcal{I},r}$ (as many as required by the R-LWE oracle), outputs an independent sample from $D_{\mathcal{I},r'}$.*

Theorem 4.1 follows easily from this iterative step, as we now sketch; see [Reg05] for more details. We start with a very large value of r, say $r \geq 2^{2n}\lambda_n(\mathcal{I})$, so that any polynomial number of samples from $D_{\mathcal{I},r}$ can be generated classically (see [Reg05, Lemma 3.2]). Then, given the samples from $D_{\mathcal{I},r}$, we apply the iterative step of Lemma 4.2 a polynomial number of times (using the same samples) to obtain a polynomial number of independent samples from $D_{\mathcal{I},r'}$ for $r' = r/2$. Repeating this, we obtain samples from progressively narrower and narrower distributions, until we get samples with the desired Gaussian parameter $s \geq \gamma$. Note that the γ given in Equation (4.1) corresponds to values of r, r' satisfying the hypotheses of Lemma 4.2.

The iterative step of Lemma 4.2 is obtained by combining two reductions. The first, whose proof is given in Section 4.2, is a reduction from BDD (on $\mathcal{I}^\vee$) to LWE, which uses Gaussian samples over $\mathcal{I}$.

Lemma 4.3. *Let $\alpha > 0$, let $q \geq 2$ be an integer with known factorization, let $\mathcal{I}$ be a fractional ideal in K, and let $r \geq \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I})$ for some negligible $\varepsilon = \varepsilon(n)$. Given an oracle for the discrete Gaussian distribution $D_{\mathcal{I},r}$, there is a probabilistic polynomial-time (classical) reduction from $\mathsf{BDD}_{\mathcal{I}^\vee,d}$ in the ℓ_∞ norm to R-$\mathsf{LWE}_{q,\Psi_{\leq\alpha}}$, where $d = \alpha q/(\sqrt{2}r)$.*

The second part is quantum, and is nearly identical to the one in Regev's reduction [Reg05, Lemma 3.14]. The only difference is that we allow the BDD oracle to err with some negligible probability.[7] This stronger statement follows from the proof in [Reg05] by noticing that the algorithm calls the BDD oracle on points whose offset is sampled from the continuous Gaussian distribution $D_{d'/(\sqrt{2n})}$.

Lemma 4.4. *There is an efficient quantum algorithm that, given any n-dimensional lattice Λ, a number $d' < \lambda_1(\Lambda^\vee)/2$ (where λ_1 is with respect to the ℓ_2 norm), and an oracle that solves BDD on $\Lambda^\vee$ with all but negligible probability for points whose offset from $\Lambda^\vee$ is sampled from $D_{d'/\sqrt{2n}}$, outputs a sample from*

[7] A similar but more extensive strengthening is used by Stehlé et al. [SSTX09].

$D_{\Lambda,\sqrt{n}/(\sqrt{2}d')}$. *In particular, since a sample from $D_{d'/\sqrt{2n}}$ has ℓ_∞ norm at most $d' \cdot \omega(\sqrt{\log n})/\sqrt{n}$ except with negligible probability, it suffices if the oracle solves $\mathsf{BDD}_{\mathcal{I}^\vee,d}$ in the ℓ_∞ norm, where $d = d' \cdot \omega(\sqrt{\log n})/\sqrt{n}$.*

Proof of Lemma 4.2. Using Lemma 4.3 with samples from $D_{\mathcal{I},r}$ and the given oracle for R-$\mathsf{LWE}_{q,\Psi_{\leq\alpha}}$, we obtain an algorithm for BDD on $\mathcal{I}^\vee$ to within distance $d = \alpha q/(\sqrt{2}r)$ in the ℓ_∞ norm. Using Lemma 4.4 with $d' = d\sqrt{n}/\omega(\sqrt{\log n}) = \sqrt{n/2}/r' < \lambda_1(\mathcal{I}^\vee)/2$, we obtain a quantum procedure that produces samples from the discrete Gaussian distribution $D_{\mathcal{I},r'}$, as promised. □

We note that one can also interpret Lemma 4.3 as a standalone *classical* (non-quantum) reduction from a seemingly hard (but non-standard) worst-case lattice problem to ring-LWE, similar to what was done in [Pei09]. In particular, since one can efficiently generate Gaussian samples from Λ using any set of n sufficiently short linearly independent lattice vectors [GPV08], the worst-case problem can defined as follows: given an ideal $\mathcal{I}$ together with a set of n linearly independent elements (or in the case of cyclotomics, even just one nonzero element) of $\mathcal{I}$ of Euclidean length at most $r/\omega(\sqrt{\log n})$, solve BDD on the dual $\mathcal{I}^\vee$ to within ℓ_∞ distance $\alpha q/(\sqrt{2}r)$.

4.2 The BDD to LWE Reduction

Our goal in this section is to prove Lemma 4.3. We first observe that to solve BDD on an ideal $\mathcal{I}$, it suffices to find the solution modulo $q\mathcal{I}$. This is actually a special case of a lemma from [Reg05], which gives a *lattice-preserving* reduction for BDD in general lattices. Because the reduction is lattice-preserving, it also applies to ideal lattices.

Definition 4.5. *The q-$\mathsf{BDD}_{\mathcal{I},d}$ problem (in any norm) is: given an instance y of $\mathsf{BDD}_{\mathcal{I},d}$ that has solution $x \in \mathcal{I}$, find $x \bmod q\mathcal{I}$.*

Lemma 4.6 (Special case of [Reg05, Lemma 3.5]). *For any $q \geq 2$, there is a deterministic polynomial-time reduction from $\mathsf{BDD}_{\mathcal{I},d}$ (in any ℓ_p norm) to q-$\mathsf{BDD}_{\mathcal{I},d}$ (in the same norm).*

Therefore, it suffices in the following to present a reduction as in Lemma 4.3 but from q-BDD. Notice that by the scaling argument in Section 2.3.8 we can assume without loss of generality that $\mathcal{I}$ is an integral ideal (in R). Finally, recall the notation $\mathbb{T} = K_\mathbb{R}/R^\vee$ and $\mathcal{J}_q = \mathcal{J}/q\mathcal{J}$ for any ideal $\mathcal{J}$.

The high-level description of the reduction is as follows. Its input is a q-$\mathsf{BDD}_{\mathcal{I}^\vee,d}$ instance $y = x + e$ (where $x \in \mathcal{I}^\vee$ and $\|e\|_\infty \leq d$), and it is given access to an oracle that generates independent samples from the discrete Gaussian distribution $D_{\mathcal{I},r}$, and an oracle $\mathcal{L}$ that solves R-LWE. The reduction produces samples from the LWE distribution $A_{s,\psi}$, where the secret s and the error distribution ψ are related to x and e, respectively. Finally, given the solution s output by $\mathcal{L}$, the reduction recovers $x \bmod q\mathcal{I}^\vee$ from s.

In detail, the reduction does the following, given a q-$\mathsf{BDD}_{\mathcal{I}^\vee,d}$ instance y:

1. Compute an element $t \in \mathcal{I}$ such that $t \cdot \mathcal{I}^{-1}$ and $\langle q \rangle$ are coprime.

   (By Lemma 2.14, such t exists and can be found efficiently using the factorization of $\langle q \rangle$.)

2. For each sample requested by $\mathcal{L}$, get a fresh $z \leftarrow D_{\mathcal{I},r}$ from the Gaussian oracle and provide to $\mathcal{L}$ the pair $(a, b) \in R_q \times \mathbb{T}$, computed as follows: let $e' \leftarrow D_{\alpha/\sqrt{2}}$, and

$$a = \theta_t^{-1}(z \bmod q\mathcal{I}) \in R_q \quad \text{and} \quad b = (z \cdot y)/q + e' \bmod R^\vee.$$

   (Recall that by Lemma 2.15 with $\mathcal{J} = \langle q \rangle$ and $\mathcal{M} = R$, the function $\theta_t(u) = t \cdot u$ induces a bijection from R_q to $\mathcal{I}_q$, which can be efficiently inverted given $\mathcal{I}$, q, and t.)

3. When $\mathcal{L}$ produces a solution $s \in R_q^\vee$, output $\theta_t^{-1}(s) \in \mathcal{I}_q^\vee$.

   (Again, by Lemma 2.15 with $\mathcal{J} = \langle q \rangle$ and $\mathcal{M} = \mathcal{I}^\vee = \mathcal{I}^{-1} \cdot R^\vee$, the function θ_t induces a bijection from $\mathcal{I}_q^\vee$ to $R_q^\vee$, which can be efficiently inverted).

The correctness of the reduction follows from Lemma 4.7 below, which says that the samples (a, b) are distributed according to $A_{s,\psi}$ for $s = \theta_t(x \bmod q\mathcal{I}^\vee) \in R_q^\vee$ and some $\psi \in \Psi_{\leq \alpha}$. By hypothesis, $\mathcal{L}$ returns s, so the reduction outputs $\theta_t^{-1}(s) = x \bmod q\mathcal{I}^\vee$, which is the correct solution to its q-$\mathsf{BDD}_{\mathcal{I}^\vee,d}$ input instance.

Lemma 4.7. *Let y be the $\mathsf{BDD}_{\mathcal{I}^\vee,d}$ instance given to the reduction above, where $y = x + e$ for some $x \in \mathcal{I}^\vee$ and $\|e\|_\infty \leq d$. Each pair (a, b) produced by the reduction has distribution $A_{s,\psi}$ (up to negligible statistical distance), for $s = \theta_t(x \bmod q\mathcal{I}^\vee) = t \cdot x \in R_q^\vee$ and some $\psi \in \Psi_{\leq \alpha}$.*

Proof. We first show that in each output pair (a, b), the component $a \in R_q$ is within negligible distance of uniform. Because $r \geq q \cdot \eta_\varepsilon(\mathcal{I})$, the second statement in Lemma 2.3 implies that all possible values of $z \bmod q\mathcal{I}$ (when z is chosen from $D_{\mathcal{I},r}$) are obtained with probabilities that are in some interval $[\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \beta$ for some $\beta > 0$, from which it follows easily that $z \bmod q\mathcal{I}$ is within distance (say) 2ε of the uniform distribution on $\mathcal{I}_q$. Finally, because θ_t induces a bijection from R_q to $\mathcal{I}_q$ by Lemma 2.15, $a = \theta_t^{-1}(z \bmod q\mathcal{I})$ is within statistical distance 2ε of uniform over R_q.

Now condition on any fixed value of a. We next analyze the component

$$b = (z \cdot y)/q + e' = (z \cdot x)/q + (z/q) \cdot e + e' \bmod R^\vee,$$

starting with $(z \cdot x)/q$. By definition of a, we have $z = \theta_t(a) = a \cdot t \in \mathcal{I}_q$. Because $x \in \mathcal{I}^\vee = \mathcal{I}^{-1} \cdot R^\vee$, we have

$$z \cdot x = \theta_t(a) \cdot x = a \cdot (t \cdot x) \bmod R_q^\vee.$$

Then because $s = t \cdot x \bmod R_q^\vee$, we have $z \cdot x = a \cdot s \bmod R_q^\vee$, which implies $(z \cdot x)/q = (a \cdot s)/q \bmod R^\vee$.

To analyze the remaining $(z/q) \cdot e + e'$ term, note that conditioned on the value of a, the random variable z/q has distribution $D_{\mathcal{I}+u/q,r/q}$, where $\mathcal{I} + u/q$ is some coset of $\mathcal{I}$ (specifically, $u = \theta_t(a) \bmod q\mathcal{I}$) and $r/q \geq \sqrt{2} \cdot \eta_\varepsilon(\mathcal{I})$. Note that

$$(r/q) \cdot \|e\|_\infty \leq (r/q) \cdot d = \alpha/\sqrt{2},$$

so we may apply Lemma 4.8 below; it implies that the distribution of $(z/q) \cdot e + e'$ is within negligible statistical distance of the elliptical Gaussian $D_{\mathbf{r}}$, where each

$$r_i^2 = (r/q)^2 \cdot |\sigma_i(e)|^2 + (\alpha/\sqrt{2})^2 \leq (r/q)^2 \cdot d^2 + \alpha^2/2 = \alpha^2.$$

We conclude that each (a, b) is distributed as $A_{s,\psi}$ for some $\psi \in \Psi_{\leq \alpha}$, as desired. □

Lemma 4.8. *Let $\mathcal{I}$ be a (fractional) ideal in K, and let $r \geq \sqrt{2} \cdot \eta_\varepsilon(\mathcal{I})$ for some $\varepsilon = \operatorname{negl}(n)$. Let $e \in K$ be fixed, let z be distributed as $D_{\mathcal{I}+v,r}$ for arbitrary $v \in K$, and let e' be distributed as $D_{r'}$ for some $r' \geq r \cdot \|e\|_\infty$. Then the distribution of $z \cdot e + e'$ is within negligible statistical distance of the elliptical Gaussian distribution $D_{\mathbf{r}}$ over $K_{\mathbb{R}}$, where $r_i^2 = r^2 \cdot |\sigma_i(e)|^2 + (r')^2$.*

Proof. Assume that $e \neq 0$, and hence every $\sigma_i(e) \neq 0$, otherwise the result holds trivially. We can write $z \cdot e + e'$ as $(z + e'/e) \cdot e$. The distribution of e'/e is the elliptical Gaussian $D_{\mathbf{t}}$, where each $t_i = r'/|\sigma_i(e)| \geq r'/\|e\|_\infty \geq r$. Thus e'/e can be written as the sum $f + g$ of independent f and g, where f has distribution D_r, and g has distribution $D_{\mathbf{t}'}$ where $(t_i')^2 = t_i^2 - r^2$.

Now by Lemma 2.5, the distribution of $z + f$ is negligibly far from $D_{\sqrt{2}r}$, so $(z + e'/e) = (z + f + g)$ has distribution negligibly far from $D_{\mathbf{t}''}$, where

$$(t_i'')^2 = 2r^2 + t_i^2 - r^2 = r^2 + (r')^2/|\sigma_i(e)|^2.$$

We conclude that $(z + e'/e) \cdot e$ has distribution negligibly far from $D_{\mathbf{r}}$, as desired. □

5 Pseudorandomness of Ring-LWE

In this section we show that for appropriate choices of ring, modulus, and error distribution, the average-case decision version of the ring-LWE problem is hard, i.e., the ring-LWE distribution is pseudorandom. For concreteness and simplicity, we specialize the discussion to cyclotomic fields (though it seems likely that our techniques can be extended to deal with other number fields). So throughout this section we assume that $\zeta = \zeta_m$ is a primitive mth root of unity, $K = \mathbb{Q}(\zeta)$ is the mth cyclotomic number field having dimension $n = \varphi(m)$, $R = \mathcal{O}_K = \mathbb{Z}[\zeta]$ is its ring of integers, $R^\vee = \mathcal{O}_K^\vee$ is its dual (codifferent) ideal, and $q = 1 \bmod m$ is a $\mathrm{poly}(n)$-bounded prime. Recall from Section 2.4 that q splits completely as $\langle q \rangle = \prod_{i \in \mathbb{Z}_m^*} \mathfrak{q}_i$, for distinct prime ideals $\mathfrak{q}_i$ each of norm $\mathrm{N}(\mathfrak{q}_i) = q$. Also recall that the automorphisms of K are $\tau_k(\zeta) = \zeta^k$ for $k \in \mathbb{Z}_m^*$, and that $\tau_k(\mathfrak{q}_i) = \mathfrak{q}_{i \cdot k^{-1}}$ and $\tau_k^{-1} = \tau_{k^{-1}}$.[8]

The following is the main theorem of this section. It gives a reduction from the search variant of ring-LWE (which by Theorem 4.1 is as hard as a worst-case lattice problem) to the average-case decision problem ring-DLWE (see Definition 3.3). This establishes the hardness of the average-case problem, which means that the LWE distribution $A_{s,\psi}$ is itself pseudorandom when both s and the error distribution ψ are chosen at random from appropriate distributions (and kept secret).

Theorem 5.1. *Let R and q be as above and let $\alpha q \geq \eta_\varepsilon(R^\vee)$ for some negligible $\varepsilon = \varepsilon(n)$. Then there is a randomized polynomial-time reduction from R-$\mathsf{LWE}_{q,\Psi_{\leq\alpha}}$ to R-$\mathsf{DLWE}_{q,\Upsilon_\alpha}$.*

Note that $\eta_{2^{-n}}(R^\vee) \leq \sqrt{n}/\lambda_1(R) = 1$, where the inequality follows by Lemma 2.2, and the equality $\lambda_1(R) = \sqrt{n}$ holds because $\|\sigma(1)\| = \sqrt{n}$ and $\lambda_1(R) \geq \sqrt{n}$ by Lemma 2.9. So in the above theorem it suffices to take $\alpha q \geq 1$, which is a slightly weaker hypothesis than that of Theorem 4.1.

The proof of Theorem 5.1 is obtained by combining four reductions, as summarized in the following diagram; the numbers refer to lemma numbers, and the definitions of all intermediate problems are given later. We note that in order to apply the last reduction (Lemma 5.14), we need a certain property of our family of noise distribution; this property is proved in Lemma 5.13.

$$\mathsf{LWE}_{q,\Psi} \xrightarrow[\text{Automorphisms}]{5.5} \mathfrak{q}_i\text{-}\mathsf{LWE}_{q,\Psi} \xrightarrow[\text{Search/Decision}]{5.9} \mathsf{WDLWE}^i_{q,\Psi} \xrightarrow[\text{Worst/Average}]{5.12} \mathsf{DLWE}^i_{q,\Upsilon} \xrightarrow[\text{Hybrid}]{5.14} \mathsf{DLWE}_{q,\Upsilon}.$$

This sequence of reductions is similar in spirit to the one given in previous work on the standard LWE problem [Reg05]. However, there are a few important differences, requiring the introduction of new tools. One fundamental issue arising in the ring setting is that an oracle for DLWE might only let us deduce the value of the secret s relative to *one* ideal factor $\mathfrak{q}_i$ of $\langle q \rangle$. In order to recover the entire secret, we 'shuffle' the $\mathfrak{q}_i$ factors using the field's automorphisms to recover s relative to *every* $\mathfrak{q}_j$ (see Lemma 5.5).

Another challenge arises from the fact that the reduction in Section 4 establishes the hardness of $\mathsf{LWE}_{q,\Psi}$ for *non-spherical* Gaussian error distributions $\psi \in \Psi$, which individually are not necessarily invariant under

[8] In fact, for any constant c dividing n, our results generalize easily to the case where q splits only into n/c distinct prime ideals $\mathfrak{q}_i$, each of norm $q^c = \mathrm{poly}(n)$, because the automorphisms still act transitively upon the $\mathfrak{q}_i$s.

the field's automorphisms. (However, the whole family Ψ of distributions is.) As a result, our reduction to an average-case problem (obtained in Lemma 5.12) needs to randomize the error distribution itself, which leads to a distribution Υ over Gaussian noise distributions that are both non-spherical and wider by a factor of about $n^{1/4}$. Although this is somewhat undesirable, we do not see any way to avoid it completely. Fortunately, this has only a minor effect on the resulting applications, i.e., adding an extra step of choosing the noise parameters.

Alternatively, there are two ways to avoid randomizing the error distribution in certain contexts. First, in many cryptographic applications there is a natural bound on the number of LWE samples available to the adversary. In such cases, the following theorem establishes pseudorandomness with a fixed spherical noise distribution.

Theorem 5.2. *Let R, q, and α be as in Theorem 5.1, and let $\ell \geq 1$. There is a randomized polynomial-time reduction from solving $R\text{-}\mathsf{LWE}_{q,\Psi_{\leq\alpha}}$ to solving $R\text{-}\mathsf{DLWE}_{q,D_\xi}$ given only ℓ samples, where $\xi = \alpha \cdot (n\ell/\log(n\ell))^{1/4}$.*

The proof of Theorem 5.2 uses the same sequence of reductions as in Theorem 5.1, except that Lemma 5.12 is replaced with Lemma 5.16.

Second, any fixed *spherical* Gaussian distribution D_α is invariant under all the automorphisms; therefore, if one assumes that the search problem $\mathsf{LWE}_{q,D_\alpha}$ is hard (which seems very plausible, though we do not have a worst-case hardness proof), then one can simplify our chain of reductions to use error distribution D_α in all the average-case problems. In this case, there is no need to use a distribution Υ over error distributions, and we do not need to lose the factor $n^{1/4}$. The proof again uses the same sequence of reductions as that of Theorem 5.1, except that Lemma 5.12 is modified so as not to randomize the error distribution, only the secret s (resulting in a considerably simpler proof).

Theorem 5.3. *Let R, q, and α be as in Theorem 5.1. There is a randomized polynomial-time reduction from solving $R\text{-}\mathsf{LWE}_{q,D_\alpha}$ to solving $R\text{-}\mathsf{DLWE}_{q,D_\alpha}$.*

5.1 Search to Worst-Case Decision

Here we reduce the search version of $\mathsf{LWE}_{q,\Psi}$ to a certain decision problem relative to just *one arbitrary* prime ideal $\mathfrak{q}_i$. All of the problems considered here are *worst-case* over the choice of $s \in R_q^\vee$ and error distribution $\psi \in \Psi$, where Ψ is the family of allowed error distributions (though the actual error terms drawn from ψ are still random), and their solutions must be found with overwhelming probability (over all the randomness of the experiment).

Our first reduction is to the following intermediate problem. Note that by Lemmas 2.12 and 2.15, there is an efficiently computable and invertible R-module isomorphism between $R_q^\vee$ and $\bigoplus_{i\in\mathbb{Z}_m^*}(R^\vee/\mathfrak{q}_i R^\vee)$.

Definition 5.4 (LWE **over** $\mathfrak{q}_i$)**.** *The $\mathfrak{q}_i\text{-}\mathsf{LWE}_{q,\Psi}$ problem is: given access to $A_{s,\psi}$ for some arbitrary $s \in R_q^\vee$ and $\psi \in \Psi$, find $s \bmod \mathfrak{q}_i R^\vee$.*

Lemma 5.5 (LWE **to** $\mathfrak{q}_i$**-**LWE)**.** *Suppose that the family Ψ is closed under all the automorphisms of K, i.e., $\psi \in \Psi \Rightarrow \tau_k(\psi) \in \Psi$ for every $k \in \mathbb{Z}_m^*$. Then for every $i \in \mathbb{Z}_m^*$, there is a deterministic polynomial-time reduction from $\mathsf{LWE}_{q,\Psi}$ to $\mathfrak{q}_i\text{-}\mathsf{LWE}_{q,\Psi}$.*

Proof. We use the oracle for $\mathfrak{q}_i$-LWE along with the field automorphisms τ_k to recover the value $s \bmod \mathfrak{q}_j R^\vee$ for *every* $j \in \mathbb{Z}_m^*$. We can then efficiently reconstruct $s \in R_q^\vee$ using the Chinese Remainder Theorem.

The reduction that finds $s \bmod \mathfrak{q}_j R^\vee$ works as follows: transform each sample $(a, b) \leftarrow A_{s,\psi}$ into the sample $(\tau_k(a), \tau_k(b)) \in R_q \times \mathbb{T}$, where $k = j/i \in \mathbb{Z}_m^*$ and hence $\tau_k(\mathfrak{q}_j) = \mathfrak{q}_i$. (Note that R and $R^\vee$ are fixed by every automorphism τ_k, so $\tau_k(R_q) = R_q$ and $\tau_k(\mathbb{T}) = \mathbb{T} = K_\mathbb{R}/R^\vee$.) Give the transformed samples to the $\mathfrak{q}_i$-$\mathsf{LWE}_{q,\Psi}$ oracle, and when the oracle returns its answer $t \in R^\vee/\mathfrak{q}_i R^\vee$, return $\tau_k^{-1}(t) \in R^\vee/\mathfrak{q}_j R^\vee$.

We now prove that $\tau_k^{-1}(t) = s \bmod \mathfrak{q}_j R^\vee$. For each sample (a, b) from $A_{s,\psi}$, notice that because $b = as/q + e \bmod R^\vee$ and $\tau_k(q) = q$, we have

$$\tau_k(b) = \tau_k(a) \cdot \tau_k(s)/q + \tau_k(e) \bmod R^\vee.$$

Because τ_k is an automorphism on R, $\tau_k(a)$ is uniformly random in $\tau_k(R_q) = R_q$, and the pairs $(\tau_k(a), \tau_k(b))$ are distributed according to $A_{\tau_k(s),\psi'}$ where $\psi' = \tau_k(\psi) \in \Psi$. The oracle must therefore return $t = \tau_k(s) \bmod \mathfrak{q}_i R^\vee$, and $\tau_k^{-1}(t) = s \bmod \tau_k^{-1}(\mathfrak{q}_i R^\vee) = s \bmod \mathfrak{q}_j R^\vee$, as desired. □

We now observe that $\Psi_{\leq\alpha}$ satisfies the closure property required by the above lemma.

Lemma 5.6. *For any $\alpha > 0$, the family $\Psi_{\leq\alpha}$ is closed under every automorphism τ of K, i.e., $\psi \in \Psi_{\leq\alpha} \Rightarrow \tau(\psi) \in \Psi_{\leq\alpha}$.*

Proof. Let $\tau_k : K \to K$ be any automorphism of K, which is of the form $\tau_k(\zeta) = \zeta^k$ for some $k \in \mathbb{Z}_m^*$. Then by the fact that τ_k simply permutes the coordinates of the canonical embedding (see Section 2.4), for any $\psi = D_\mathbf{r} \in \Psi_{\leq\alpha}$, we have $\tau_k(D_\mathbf{r}) = D_{\mathbf{r}'} \in \Psi_{\leq\alpha}$, where the entries of $\mathbf{r}'$ are merely a rearrangement of the entries of $\mathbf{r}$ and hence are all at most α. □

For our second reduction, we need to introduce a few more definitions. For notational convenience, we identify the elements of $\mathbb{Z}_m^*$ with their integer representatives from the set $\{1, \ldots, m-1\}$, with the usual ordering. For $i \in \mathbb{Z}_m^*$ we let $i-$ denote the largest element in $\mathbb{Z}_m^*$ less than i, defining $1-$ to be 0.

Definition 5.7 (Hybrid LWE distribution). *For $i \in \mathbb{Z}_m^*$, $s \in R_q^\vee$, and a distribution ψ over $K_\mathbb{R}$, the distribution $A_{s,\psi}^i$ over $R_q \times \mathbb{T}$ is defined as follows: choose $(a, b) \leftarrow A_{s,\psi}$ and output $(a, b + h/q)$ where $h \in R_q^\vee$ is uniformly random and independent mod $\mathfrak{q}_j R^\vee$ for all $j \leq i$, and is 0 mod all the remaining $\mathfrak{q}_j R^\vee$. Also define $A_{s,\psi}^0$ simply as $A_{s,\psi}$.*

Definition 5.8 (Worst-case decision LWE relative to $\mathfrak{q}_i$). *For $i \in \mathbb{Z}_m^*$ and a family of distributions Ψ, the $\mathsf{WDLWE}_{q,\Psi}^i$ problem is defined as follows: given access to $A_{s,\psi}^j$ for arbitrary $s \in R_q^\vee$, $\psi \in \Psi$, and $j \in \{i-, i\}$, find j.*

Lemma 5.9 (Search to Decision). *For any $i \in \mathbb{Z}_m^*$, there is a probabilistic polynomial-time reduction from $\mathfrak{q}_i$-$\mathsf{LWE}_{q,\Psi}$ to $\mathsf{WDLWE}_{q,\Psi}^i$.*

Proof. The idea for recovering $s \bmod \mathfrak{q}_i R^\vee$ is to try each of its possible values, modifying the samples we receive from $A_{s,\psi}$ so that on the correct value the modified samples are distributed according to $A_{s,\psi}^{i-}$, whereas on all the other values the modified samples are distributed according to $A_{s,\psi}^i$. We can then use the $\mathsf{WDLWE}_{q,\Psi}^i$ oracle to tell us which distribution was generated. Because there are only $\mathrm{N}(\mathfrak{q}_i) = q = \mathrm{poly}(n)$ possible values for $s \bmod \mathfrak{q}_i R^\vee$, we can enumerate over all of them efficiently and discover the correct value.

We now give the transformation that takes some $g \in R_q^\vee$ and maps $A_{s,\psi}$ to either $A_{s,\psi}^{i-}$ or $A_{s,\psi}^i$, depending on whether or not $g = s \bmod \mathfrak{q}_i R^\vee$ (its values modulo the other $\mathfrak{q}_j R^\vee$ are irrelevant). Given a sample $(a, b) \leftarrow A_{s,\psi}$, the transformation produces a sample

$$(a', b') = (a + v, b + (h + vg)/q) \in R_q \times \mathbb{T},$$

where $v \in R_q$ is uniformly random mod $\mathfrak{q}_i$ and is 0 mod the other $\mathfrak{q}_j$, and $h \in R_q^\vee$ is uniformly random and independent mod $\mathfrak{q}_j R^\vee$ for all $j < i$, and is 0 mod all the remaining $\mathfrak{q}_j R^\vee$. First, notice that since a is uniformly distributed in R_q, so is a'. Next, condition on any fixed value of a'. Then b' can be written as

$$\begin{aligned} b' = b + (h + vg)/q &= (as + h + vg)/q + e \\ &= (a's + h + v(g - s))/q + e, \end{aligned}$$

where e is drawn from ψ.

We consider two cases. First, assume that $g = s \bmod \mathfrak{q}_i R^\vee$. Then by the Chinese Remainder Theorem (Lemma 2.12), $v(g - s) = 0 \in R_q^\vee$, and hence the distribution of (a', b') is exactly $A_{s,\psi}^{i-}$. Next, assume that $g \neq s \bmod \mathfrak{q}_i R^\vee$. Then since $\mathfrak{q}_i$ is a maximal ideal (which in R is equivalent to being a prime ideal), $R/\mathfrak{q}_i$ is a field, and hence $v(g - s) \in R_q^\vee$ is distributed uniformly mod $\mathfrak{q}_i R^\vee$, and is zero mod all other $\mathfrak{q}_j R^\vee$. From this it follows that $v(g - s) + h$ is uniformly random and independent mod $\mathfrak{q}_j R^\vee$ for all $j \leq i$, and is 0 mod all the remaining $\mathfrak{q}_j R^\vee$. Hence, the distribution of (a', b') is exactly $A_{s,\psi}^i$, as claimed. □

5.2 Worst-Case Decision to Average-Case Decision

We now reduce the worst-case decision problem $\mathsf{WDLWE}_{q,\Psi}^i$ (where Ψ is a family of Gaussian noise distributions) to an entirely average-case problem, namely, distinguishing $A_{s,\psi}$ from the uniform distribution (with any non-negligible advantage) for a *random* choice of both s and ψ, where the parameters of the error distribution ψ themselves are drawn at random from a certain distribution Υ and kept secret.

We first define the following variant of average-case decision LWE.

Definition 5.10 (Average-case decision LWE relative to $\mathfrak{q}_i$). *For $i \in \mathbb{Z}_m^*$ and a distribution Υ over error distributions, we say that an algorithm solves the $\mathsf{DLWE}_{q,\Upsilon}^i$ problem if with a non-negligible probability over the choice of a random $(s, \psi) \leftarrow U(R_q^\vee) \times \Upsilon$, it has a non-negligible difference in acceptance probability on inputs from $A_{s,\psi}^i$ versus inputs from $A_{s,\psi}^{i-}$.*

We will need the following technical claim.

Claim 5.11. *Let P be the distribution $\Gamma(2,1)^n$ and Q be the distribution $(\Gamma(2,1) - z_1) \times \cdots \times (\Gamma(2,1) - z_n)$ for some $0 \leq z_1, \ldots, z_n \leq 1/\sqrt{n}$. Then any set $A \subseteq \mathbb{R}^n$ whose measure under P is non-negligible also has non-negligible measure under Q.*

We remark that the claim is sharp in the sense that if we take $z_1 = \cdots = z_n = \omega(1/\sqrt{n})$ then the positive quadrant has measure 1 under P, but negligible measure under Q.

Proof. For any two probability density functions $P, Q : \mathbb{R}^n \to \mathbb{R}^{\geq 0}$ where $P(x) = 0$ whenever $Q(x) = 0$ (which is the case for the P and Q in the lemma statement), define

$$R(P||Q) = \int_{\mathbb{R}^n} \frac{P(x)^2}{Q(x)} \, \mathrm{d}x,$$

with the convention that the fraction is zero when both numerator and denominator are zero. (The logarithm of this quantity is known as the Rényi divergence of order 2.) By Cauchy-Schwarz, for any set $A \subseteq \mathbb{R}^n$,

$$\frac{(\int_A P(x) \, \mathrm{d}x)^2}{\int_A Q(x) \, \mathrm{d}x} \leq \int_A \frac{P(x)^2}{Q(x)} \, \mathrm{d}x \leq R(P||Q).$$

Hence if a set A has non-negligible measure under P and $R(P||Q) \leq \mathrm{poly}(n)$, then A also has non-negligible measure under Q.[9]

We now apply this in our setting. A straightforward calculation shows that for all $z > 0$,

$$R(\Gamma(2,1) \,||\, \Gamma(2,1) - z) = e^z \left(1 - z + z^2 e^z \int_z^\infty x^{-1} e^{-x} \mathrm{d}x\right),$$

which for small z is easily seen to be $1 + z^2 \log(1/z) + O(z^2)$. Hence,

$$\begin{aligned} &R(\Gamma(2,1)^n \,||\, (\Gamma(2,1) - z_1) \times \cdots \times (\Gamma(2,1) - z_n)) \\ &= R(\Gamma(2,1) \,||\, \Gamma(2,1) - z_1) \cdots R(\Gamma(2,1) \,||\, \Gamma(2,1) - z_n) \end{aligned}$$

is polynomial in n. □

Lemma 5.12 (Worst-case to average-case). *For any $\alpha > 0$ and every $i \in \mathbb{Z}_m^*$, there is a randomized polynomial-time reduction from $\mathsf{WDLWE}^i_{q,\Psi_{\leq\alpha}}$ to $\mathsf{DLWE}^i_{q,\Upsilon_\alpha}$.*

Proof. For some $s' \in R_q^\vee$, $\mathbf{r}' \in (\mathbb{R}^+)^n$, and $k \in \mathbb{Z}_m^*$, consider the transformation mapping each (a, b) to $(a, b + (a \cdot s' + h)/q + e')$ where e' is chosen from $D_{\mathbf{r}'}$, and $h \in R_q^\vee$ is uniformly random and independent mod $\mathfrak{q}_j R^\vee$ for all $j \leq k$, and 0 mod all the remaining $\mathfrak{q}_j R^\vee$. Then it is easy to see that for all $s \in R_q^\vee$ and $i \in \mathbb{Z}_m^*$, this transformation maps $A^i_{s,\psi}$ to $A^{\max\{k,i\}}_{s+s',\psi+D_{\mathbf{r}'}}$.

The reduction repeats the following a polynomial number of times. Choose a uniform $s' \in R_q^\vee$ as well as reals $x_1, \ldots, x_{n/2}$ chosen independently from the distribution $\Gamma(2, 1)$ and let $\mathbf{r}' \in (\mathbb{R}^+)^n$ be defined by $r'^2_j = r'^2_{j+n/2} = \alpha^2 \sqrt{n} x_j$ for $j \in [n/2]$. Then estimate the acceptance probability of the oracle on the following two input distributions: the first is obtained from our input by applying the above transformation with parameters s', $\mathbf{r}'$, and $i-$; the second is obtained similarly using parameters s', $\mathbf{r}'$, and i. If in any of these polynomial number of attempts a non-negligible difference is observed between the two acceptance probabilities, output "$i-$"; otherwise output "i".

Notice that if our input distribution is $A^i_{s,\psi}$, then in each of the attempts, the two distributions on which we estimate the oracle's acceptance probability are exactly the same, hence we output "i" with overwhelming probability. So assume that our input distribution is $A^{i-}_{s,D_\mathbf{r}}$ for some $\mathbf{r}$ satisfying that all r_i are in $[0, \alpha]$. In this case we estimate the oracle's acceptance probability on $A^{i-}_{s+s',D_\mathbf{r}+D_{\mathbf{r}'}}$ and $A^i_{s+s',D_\mathbf{r}+D_{\mathbf{r}'}}$, and notice that $D_\mathbf{r} + D_{\mathbf{r}'} = D_{\mathbf{r}''}$ where $r''^2_j = r^2_j + r'^2_j$. Let S be the set of all pairs (s, ψ) for which the oracle has a non-negligible difference in acceptance probability on $A^{i-}_{s,\psi}$ and $A^i_{s,\psi}$. By assumption, the measure of S under $U(R_q^\vee) \times \Upsilon_\alpha$ is non-negligible. By Claim 5.11, $(s + s', D_\mathbf{r} + D_{\mathbf{r}'}) \in S$ with non-negligible probability and the lemma follows. □

Lemma 5.13. *Let $\alpha \geq \eta_\varepsilon(R^\vee)/q$ for some $\varepsilon > 0$. Then for any ψ in the support of Υ_α and $s \in R_q^\vee$, the distribution $A^{m-1}_{s,\psi}$ is within statistical distance $\varepsilon/2$ of the uniform distribution over $(R_q, \mathbb{T})$.*

Proof. By definition, a sample from the distribution $A^{m-1}_{s,\psi}$ is given by $(a, (a \cdot s + h)/q + e)$ where e is chosen from ψ, a is chosen uniformly from R_q, and h is chosen uniformly from $R_q^\vee$. It suffices to show that conditioned on any fixed value of a, the second element of the pair is within statistical distance ε of the uniform distribution over $\mathbb{T}$. So fix some value of a. Notice first that $(a \cdot s + h)/q$ is distributed like a uniform

[9] We note that other notions of relative entropy also have this property. The advantage of our choice of $R(\cdot||\cdot)$ is that it is easy to calculate, as we do below.

element of $(q^{-1}R^\vee)/R^\vee$. Moreover, any noise distribution ψ in the support of Υ_α can be written as the sum of two independent Gaussian noise distributions $D_{\mathbf{r}} + D_{\mathbf{r}'}$, the first with parameters $r_i = \alpha$ and the second with parameters $(r'_i)^2 = x_i \geq 0$. By Lemma 2.3 and our assumption on α, the sum of a uniform element of $(q^{-1}R^\vee)/R^\vee$ and noise chosen from $D_{\mathbf{r}}$ is within statistical distance $\varepsilon/2$ of the uniform distribution on $\mathbb{T}$, and clearly this remains the case after adding the independent noise $D_{\mathbf{r}'}$. □

Lemma 5.14 (Hybrid). *Let Υ be a distribution over noise distributions satisfying that for any ψ in the support of Υ and any $s \in R_q^\vee$, the distribution $A_{s,\psi}^{m-1}$ is within negligible statistical distance from uniform. Then for any oracle solving the $\mathsf{DLWE}_{q,\Upsilon}$ problem, there exists an $i \in \mathbb{Z}_m^*$ and an efficient algorithm that solves $\mathsf{DLWE}_{q,\Upsilon}^i$ using the oracle.*

Proof. We use a simple hybrid argument. Let (s, ψ) be any pair for which the oracle distinguishes between $A_{s,\psi}$ and uniform inputs with a non-negligible advantage. By Markov's inequality, the probability measure of such pairs is non-negligible. Since $A_{s,\psi}^0 = A_{s,\psi}$, and $A_{s,\psi}^{m-1}$ is negligibly far from the uniform distribution, we see that for each such (s, ψ) there must exist an $i \in \mathbb{Z}_m^*$ for which the oracle distinguishes between $A_{s,\psi}^i$ and $A_{s,\psi}^{i-}$ with non-negligible advantage. The lemma follows by taking the i that is associated to the set of pairs (s, ψ) of highest probability. □

This completes the proof of Theorem 5.1. To prove Theorem 5.2, we start with a technical claim and proceed with an alternative noise reduction.

Claim 5.15. *Let $r_1, \ldots, r_n \in \mathbb{R}^+$ and $s_1, \ldots, s_n \in \mathbb{R}^+$ be such that for all i, $|s_i/r_i - 1| < \sqrt{\log n/n}$. Then any set $A \subseteq \mathbb{R}^n$ whose measure under the Gaussian distribution $D_{r_1} \times \cdots \times D_{r_n}$ is non-negligible, also has non-negligible measure under $D_{s_1} \times \cdots \times D_{s_n}$.*

Proof. We use the same notation and technique as in Claim 5.11. An easy calculation shows that for all $r > 0$ and $\alpha > 1/\sqrt{2}$,

$$R(D_r \,||\, D_{\alpha r}) = \frac{\alpha^2}{\sqrt{2\alpha^2 - 1}},$$

which is smaller than, say, $1 + 3(\alpha - 1)^2$ for α sufficiently close to 1. Hence,

$$\begin{aligned} R(D_{r_1} \times \cdots \times D_{r_n} \,||\, D_{s_1} \times \cdots \times D_{s_n}) &= R(D_{r_1} \,||\, D_{s_1}) \cdots R(D_{r_n} \,||\, D_{s_n}) \\ &\leq (1 + 3 \log n/n)^n = \operatorname{poly}(n). \end{aligned}$$

□

Lemma 5.16 (Worst-case to average-case with spherical noise). *For any $\alpha > 0$, $\ell \geq 1$, and every $i \in \mathbb{Z}_m^*$, there is a randomized polynomial-time reduction from solving $\mathsf{WDLWE}_{q,\Psi_{\leq\alpha}}^i$ to solving $\mathsf{DLWE}_{q,D_\xi}^i$ given only ℓ samples, where $\xi = \alpha(n\ell/\log(n\ell))^{1/4}$.*

Proof. For some $s' \in R_q^\vee$, $k \in \mathbb{Z}_m^*$, and $e_1, \ldots, e_\ell \in \mathbb{T}$, consider the transformation mapping ℓ samples $(a_i, b_i)_{i=1}^\ell$ to $(a_i, b_i + (a_i \cdot s' + h_i)/q + e_i)_{i=1}^\ell$ where $h_1, \ldots, h_\ell \in R_q^\vee$ are chosen independently to be uniform mod $\mathfrak{q}_j R^\vee$ for all $j \leq k$, and 0 mod all the remaining $\mathfrak{q}_j R^\vee$. Then it is easy to see that for all $s \in R_q^\vee$, ψ, $\mathbf{r}'$, and $i \in \mathbb{Z}_m^*$, if we sample from $(A_{s,\psi}^i)^\ell$ (i.e., ℓ independent samples from $A_{s,\psi}^i$) and apply this transformation with $e_1, \ldots, e_\ell$ chosen independently from $D_{\mathbf{r}'}$, then the output distribution (averaged over the choice of $e_1, \ldots, e_\ell$) is $(A_{s+s',\psi+D_{\mathbf{r}'}}^{\max\{k,i\}})^\ell$.

The reduction repeats the following a polynomial number of times. Choose a uniform $s' \in R_q^\vee$ as well as $e_1, \ldots, e_\ell$ chosen independently from D_ξ. Then estimate the acceptance probability of the oracle on the following two input distributions: the first is obtained from our input by applying the above transformation with parameters s', $e_1, \ldots, e_\ell$, and $i-$; the second is obtained similarly using parameters s', $e_1, \ldots, e_\ell$, and i. If in any of these polynomial number of attempts a non-negligible difference is observed between the two acceptance probabilities, output "$i-$"; otherwise output "i".

Notice that if our input distribution is $A^i_{s,\psi}$, then in each of the attempts, the two distributions on which we estimate the oracle's acceptance probability are exactly the same, hence we output "i" with overwhelming probability. So assume that our input distribution is $A^{i-}_{s,D_{\mathbf{r}}}$ for some $\mathbf{r}$ satisfying that all r_i are in $[0, \alpha]$. Let $B^{i-}(s', e_1, \ldots, e_\ell)$ and $B^i(s', e_1, \ldots, e_\ell)$ be the two distributions on ℓ pairs which our reduction uses as input to the oracle. Define the vector $\mathbf{r}'$ with coordinates $r_j'^2 = \xi^2 - r_j^2$ so that $D_{\mathbf{r}} + D_{\mathbf{r}'} = D_\xi$. By our observation above, the average of $B^{i-}(s', e_1, \ldots, e_\ell)$ over $e_1, \ldots, e_\ell$ chosen independently from $D_{\mathbf{r}'}$ is $(A^{i-}_{s+s',D_\xi})^\ell$ and similarly with B^i and A^i. Let S be the set of all tuples $(s, e_1, \ldots, e_\ell)$ for which the oracle has a non-negligible difference in acceptance probability on $B^{i-}(s', e_1, \ldots, e_\ell)$ and $B^i(s', e_1, \ldots, e_\ell)$. By assumption and a Markov argument, the measure of S under $U(R_q^\vee) \times (D_{\mathbf{r}'})^\ell$ is non-negligible. Since

$$1 \leq \frac{\xi}{\sqrt{\xi^2 - r_i^2}} \leq \frac{\xi}{\sqrt{\xi^2 - \alpha^2}} \leq 1 + \sqrt{\frac{\log(n\ell)}{n\ell}},$$

it follows from Claim 5.15 that the measure of S under $U(R_q^\vee) \times (D_\xi)^\ell$ is also non-negligible, and we are done. □

References

[ABB10a] S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, pages 553–572. 2010.

[ABB10b] S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115. 2010.

[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.

[AGV09] A. Akavia, S. Goldwasser, and V. Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *TCC*, pages 474–495. 2009.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.

[AKS01] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610. 2001.

[Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.

[BFKL93] A. Blum, M. L. Furst, M. J. Kearns, and R. J. Lipton. Cryptographic primitives based on hard learning problems. In *CRYPTO*, pages 278–291. 1993.

[BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *ICTS*, pages 309–325. 2012.

[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.

[Boy10] X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Public Key Cryptography*, pages 499–517. 2010.

[BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, pages 97–106. 2011.

[CHKP10] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552. 2010.

[Coh93] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1993.

[Con09] K. Conrad. The different ideal, 2009. Available at `http://www.math.uconn.edu/~kconrad/blurbs/`, last accessed 12 Oct 2009.

[DH76] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.

[ElG84] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *CRYPTO*, pages 10–18. 1984.

[Erd46] P. Erdős. On the coefficients of the cyclotomic polynomial. *Bulletin of the American Mathematical Society*, 52(2):179–184, 1946.

[Gen09] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178. 2009.

[GGH96] O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(42), 1996.

[GKPV10] S. Goldwasser, Y. T. Kalai, C. Peikert, and V. Vaikuntanathan. Robustness of the learning with errors assumption. In *ICS*, pages 230–240. 2010.

[Gol04] O. Goldreich. *Foundations of Cryptography*, volume II. Cambridge University Press, 2004.

[GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008.

[HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288. 1998.

[KTX08] A. Kawachi, K. Tanaka, and K. Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In *ASIACRYPT*, pages 372–389. 2008.

[LM06] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155. 2006.

[LM08] V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, pages 37–54. 2008.

[LMPR08] V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72. 2008.

[LP11] R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In *CT-RSA*, pages 319–339. 2011.

[LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In *EUROCRYPT*, pages 35–54. 2013.

[LS13] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices, 2013. Submitted.

[LTV12] A. López-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, pages 1219–1234. 2012.

[Lyu08] V. Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Public Key Cryptography*, pages 162–179. 2008.

[Lyu09] V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616. 2009.

[Lyu12] V. Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755. 2012.

[Mic02] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.

[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.

[MV03] D. Micciancio and S. P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO*, pages 282–298. 2003.

[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.

[PR06] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166. 2006.

[PR07] C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487. 2007.

[PVW08] C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571. 2008.

[PW08] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196. 2008.

[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.

[Roq67] P. Roquette. On class field towers. In J. W. S. Cassels and A. Fröhlich, editors, *Algebraic Number Theory*, pages 231–249. Academic Press, 1967.

[Sho09] V. Shoup. *A Computational Introduction to Number Theory and Algebra*. Cambridge University Press, 2009.

[SS11] D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, pages 27–47. 2011.

[SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635. 2009.

[Ste04] W. Stein. A brief introduction to classical and adelic algebraic number theory, 2004. Available at `http://modular.math.washington.edu/papers/ant/`, last accessed 12 Oct 2009.