

On Ideal Lattices and Learning with Errors Over Rings

Vadim Lyubashevsky*, Chris Peikert**, and Oded Regev***

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. Finally, the algebraic structure of ring-LWE might lead to new cryptographic applications previously not known to be based on 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.

* Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. Supported by a European Research Council (ERC) Starting Grant. Part of this work was performed while visiting Georgia Tech.

** School of Computer Science, College of Computing, Georgia Institute of Technology. Email: cpeikert@cc.gatech.edu. This material is based upon work supported by the National Science Foundation under Grant CNS-0716786. 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.

*** Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. 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.

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 [2] to be at least as hard as approximating several lattice problems, such as the (gap) shortest vector problem, to within a polynomial factor in the lattice dimension. More recently, Regev [31] 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 an efficient quantum algorithm for approximate lattice problems.) Peikert [26] subsequently proved the hardness of LWE under certain lattice assumptions, via a classical reduction.

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 ones, where $\mathbf{s} \in \mathbb{Z}_q^n$ 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 relatively concentrated around 0 (modulo q).

In recent years, a multitude of cryptographic schemes have been proposed around the SIS and LWE problems. As a search problem (without unique solution), SIS has been the foundation for one-way [2] and collision-resistant hash functions [15], identification schemes [25, 18, 17], and digital signatures [13, 8]. The LWE problem has proved to be amazingly versatile for encryption schemes, serving as the basis for secure public-key encryption under both chosen-plaintext [31, 29] and chosen-ciphertext [30, 26] attacks, oblivious transfer [29], identity-based encryption [13, 8, 1], various forms of leakage-resilient encryption (e.g., [4, 6]), and more.

One 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 functions, have key sizes at least *quadratic* in the primary security parameter, which needs to be in the several hundreds for sufficient security against the best known attacks.

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 [16], Micciancio [23] proposed a “compact,” efficient one-way function using a ring-based variant of SIS that he showed to be at least as hard as worst-case problems on *cyclic lattices*. Later, Peikert and Rosen [27] and Lyubashevsky and Micciancio [20] independently constructed collision-resistant hash functions based on *ideal lattices* (a generalization of cyclic lattices), and provided a fast and practical implementation [22]. These results paved the way for other

efficient cryptographic constructions, including identification schemes [19] and signatures [21, 19]. (The recent fully homomorphic cryptosystem of Gentry [12] is also based on ideal lattices, but it relies on new assumptions that are not related to SIS or LWE.)

Despite its expected utility, a compact analogue of LWE with comparable security properties has not yet appeared in the literature (though see Section 1.5 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 an appropriate version of the learning with errors problem in a wide class of rings, and to prove its hardness under worst-case assumptions on ideal lattices in these rings.

1.1 Informal Description of Results

Here we give an informal overview of the ring-LWE problem and our hardness results for it. For concreteness, this summary deals with one particular ‘nice’ ring, and deliberately omits the exact error distribution for which we can prove hardness. Our results actually apply much more generally to *rings of algebraic integers* in number fields, and the error distribution is defined precisely using concepts from algebraic number theory.

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. (This particular $f(x)$ comes from the family of *cyclotomic* polynomials, which play a special role in this work.) Let $R = \mathbb{Z}[x]/\langle f(x) \rangle$ be the ring of integer polynomials modulo $f(x)$. Elements of R (i.e., residues mod $f(x)$) are typically 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 . Elements of R_q may be represented by polynomials of degree less than n whose coefficients are from $\{0, \dots, q-1\}$.

In the above-described ring, the R -LWE problem may be described as follows. Let $s = s(x) \in R_q$ be a uniformly random ring element, which is kept secret. Analogously to standard LWE, the goal of the attacker is to distinguish arbitrarily many (independent) ‘random noisy ring equations’ from truly uniform ones. More specifically, the noisy equations are of the form $(a, b \approx a \cdot s) \in R_q \times R_q$, where a is uniformly random and the product $a \cdot s$ is perturbed by some ‘small’ random error term, chosen from a certain distribution over R .

Main Theorem 1 (Informal) *Suppose that it is hard for polynomial-time quantum algorithms to approximate the shortest vector problem (SVP) in the worst case on ideal lattices¹ in R to within a fixed $\text{poly}(n)$ factor. Then any $\text{poly}(n)$ number of samples drawn from the R -LWE distribution are pseudorandom to any polynomial-time (even quantum) attacker.*

¹ Briefly, an ideal lattice in R is just an ideal under some appropriate geometric embedding. See Section 1.3 for a precise definition and discussion.

Our main theorem follows from two component results, which are each of independent interest.

Worst-case hardness of the search problem. 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 follows the general outline of Regev’s iterative quantum reduction for general lattices [31], but ideal lattices introduce several new technical roadblocks in both the ‘algebraic’ and ‘geometric’ components of the reduction. We overcome these obstacles using perspectives and tools from algebraic number theory, in particular, the canonical embedding of a number field and the Chinese remainder theorem. Our result is stated formally as Theorem 1, and is proved throughout Section 3.

We point out that in contrast with standard LWE, the precise error distribution for which we can prove worst-case hardness is somewhat subtle: the distribution has up to n independent parameters (one for each direction in a certain orthogonal basis) which themselves are chosen at random and kept secret. Most cryptographic applications only require (for correctness) that the error distribution be relatively concentrated, so this form of noise generally presents no problem. (It is also possible to show hardness for a fixed spherical distribution, but for a slightly super-polynomial approximation factor, modulus q , and reduction runtime.) The non-spherical error distribution is an artifact of our proof technique, and can perhaps be avoided using additional ideas.

Search / decision equivalence. We then show that the R -LWE distribution is in fact *pseudorandom* if the search problem is hard (given arbitrarily many samples). This result is also inspired by analogous reductions for the standard LWE problem [7, 31], but again the ring context presents new obstacles, primarily related to proving that the *entire* n -dimensional quantity $b \approx a \cdot s$ is simultaneously pseudorandom. Here again, the solution seems to rely inherently on tools from algebraic number theory. The full result is stated as Theorem 2, and is proved throughout Section 4.

We stress that our search/decision equivalence works for a wide class of natural noise distributions, and is entirely classical (no quantum). Therefore, it is of value even without our worst-case reduction, and can be understood independently of it. For example, if one makes the plausible conjecture that the search version of R -LWE is hard for a fixed spherical error distribution and small modulus q , then our proof demonstrates that the same R -LWE distribution is also pseudorandom.

1.2 Discussion and Applications

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 function): each noisy product $b \approx a \cdot s$ is a pseudorandom n -dimensional vector over \mathbb{Z}_q , rather than just a

scalar, and we can generate as many of these values as we like. Yet the cost of generating them is quite small: polynomial multiplication can be performed in $O(n \log n)$ scalar operations using the Fast Fourier Transform (FFT). Moreover, the specific choice of polynomial $f(x) = x^n + 1$ and modulus $q = 1 \bmod 2n$ (among others) admits an optimized implementation that works entirely over the field \mathbb{Z}_q , and is very fast on modern architectures (see [22]). 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 $\Theta(n)$ factor. This is especially beneficial because key size has probably been the main barrier to practical lattice-based cryptosystems with rigorous security analysis.

Sample cryptosystem. As an example application, we exploit the pseudorandomness of the R -LWE distribution (e.g., over the ring $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ described above) to construct a simple semantically secure public-key cryptosystem. This scheme and its security proof are a direct translation of the ‘dual’ scheme from [13] based on the standard LWE problem, and similarly direct adaptations are possible for most other LWE-based schemes, including Regev’s original ‘primal’ cryptosystem [31], Peikert’s CCA-secure cryptosystem [26], and at least the identity-based encryption schemes of [13, 8].²

In our example cryptosystem, the key generation algorithm chooses $m \approx \lg q = O(\log n)$ uniformly random and independent elements $a_i \in R_q$, along with m random ‘small’ ring elements $r_i \in R$ (e.g., having uniformly random and independent 0-1 coefficients when viewed as polynomials). The element $a_{m+1} \in R_q$ is computed as $a_{m+1} = \sum_{i \in [m]} r_i \cdot a_i$. The public and secret keys, respectively, are the tuples

$$(a_1, \dots, a_{m+1}) \in R_q^{m+1} \quad \text{and} \quad (r_1, \dots, r_m, r_{m+1} = -1) \in R^{m+1}.$$

This key generation procedure has two main properties: first, the public key is essentially uniform (statistically) over R_q^{m+1} , which can be shown by a variant of the leftover hash lemma for the ring R_q [23]. Second, the public and secret keys satisfy $\sum_i r_i \cdot a_i = 0 \in R_q$.

To encrypt an n -bit message $z \in \{0, 1\}^n$, view it as an element of R by using its bits as the 0-1 coefficients of a polynomial. Choose a uniformly random $s \in R_q$, and for each $i \in [m + 1]$ compute $b_i \approx a_i \cdot s \in R_q$, where each product is perturbed by an independent ‘small’ error term $e_i \in R$ from the prescribed LWE error distribution. Lastly, subtract (modulo q) from b_{m+1} the ring element $z \cdot \lfloor q/2 \rfloor$. The ciphertext is the tuple $(b_1, \dots, b_{m+1}) \in R_q^{m+1}$. Note that semantic security is straightforward to prove, because the adversary’s view, i.e., the public key and ciphertext, simply consists of $m + 1$ samples from the pseudorandom R -LWE distribution, which hide the message.

To decrypt the ciphertext, simply compute

$$\sum r_i \cdot b_i \approx z \cdot \lfloor q/2 \rfloor + \left(\sum r_i \cdot a_i \right) \cdot s = z \cdot \lfloor q/2 \rfloor + 0 \cdot s \in R_q,$$

² Some of these constructions also require an adaptation of the basis-generation procedure of [5] to the ring setting, which was done in [32].

where the \approx symbol hides $\sum_i r_i \cdot e_i \in R$, the error terms accumulated by the short elements from the secret key. For appropriate choices of parameters, the coefficients of this sum have magnitudes much smaller than $q/2$, so the bits of z can be recovered by rounding each coefficient back to either 0 or $\lfloor q/2 \rfloor$, whichever is closest (mod q).

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 much effort in this vein, we have been unable to make any significant progress in attacking these problems. The best known algorithms for ideal lattices perform essentially no better than their generic counterparts, which require exponential time and space to achieve a $\text{poly}(n)$ approximation factor [3].

We also gain some confidence in the inherent hardness of ideal lattices from the fact that they arise (under a suitable definition; see Section 1.3 below) from a deep and well-studied branch of mathematics, which has also been investigated reasonably thoroughly from a computational point of view (see, e.g., [9]). 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.

1.3 Ideal Lattices

Here we give a brief description of *ideal lattices*, survey their use in previous work, and compare to our work. All of the definitions of ideal lattices from prior work are instances of the following general notion: let R be a ring whose additive group is isomorphic to \mathbb{Z}^n (i.e., it is a free \mathbb{Z} -module of rank n), and let σ be an additive isomorphism mapping R to some lattice $\sigma(R)$ in an n -dimensional real vector space (e.g., \mathbb{R}^n). The family of ideal lattices for the ring R under the embedding σ is the set of all lattices $\sigma(\mathcal{I})$, where \mathcal{I} is an ideal in R .³ For instance, taking $R = \mathbb{Z}[x]/\langle x^n - 1 \rangle$ and the naïve “coefficient embedding” σ , i.e., the one that views the coefficients of a polynomial residue (modulo $x^n - 1$) as an integer vector in \mathbb{Z}^n , leads exactly to the family of (integer) cyclic lattices. Note that under the coefficient embedding, addition of ring elements simply corresponds to (coordinate-wise) addition of their vectors in \mathbb{Z}^n , but multiplication does not have such a nice geometrical interpretation, due to the reduction modulo $x^n - 1$.

The main difference between this work and almost all previous work is in the choice of embedding σ . Prior works [23, 27, 20, 21, 12, 19, 32] used rings of the form $\mathbb{Z}[x]/\langle f(x) \rangle$ with the coefficient embedding described above. In this work, following Peikert and Rosen [28], we instead consider the so-called *canonical embedding* from algebraic number theory. Strictly speaking, the coefficient and

³ An ideal \mathcal{I} in a ring R is an additive subgroup of R that is closed under multiplication by R .

canonical embeddings are equivalent up to a fixed linear transformation that introduces some distortion. (In fact, this is true of *any* two fixed embeddings, under our definition above.) Moreover, in many cases the distortion is small; for example, in the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$ for n a power of 2, the transformation is even an isometry (i.e., a scaled rotation). In such cases, lattice problems are essentially equivalent under either embedding. Yet due to its central role in the study of number fields and useful geometric properties (explained below), we contend that the canonical embedding is the ‘right’ notion to use in the study of ideal lattices.

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 leading to tight bounds, and probability distributions such as Gaussians behave very nicely under multiplication by fixed elements. In contrast, understanding the behavior of multiplication under the coefficient embedding required previous work to introduce notions like the *expansion factor*, which implicitly measures the distortion involved in going between the coefficient and canonical embeddings, but is not of much help for analyzing probability distributions. Second, although for many rings the two embeddings are nearly isometric, in many other rings of interest the distortion can be quite large — even *super-polynomial* in the dimension for some cyclotomic polynomial rings [11]. This may explain why we can prove tight hardness results for *all* such cyclotomic rings (as explained below), whereas previous work was mostly restricted to $\mathbb{Z}[x]/\langle x^n + 1 \rangle$ for n a power of 2 (and a few others). A third point in favor of the canonical embedding is that it also behaves very nicely under the automorphisms that we use in our search-to-decision reductions for ring-LWE.

Moving now to the choice of ring R , in this work our main focus is on the rings of integer polynomials modulo a *cyclotomic* polynomial.⁴ From an algebraic point of view, it is more natural to view these rings as the rings of (algebraic) integers in cyclotomic number fields, and this is indeed the perspective we adopt. Moreover, our main theorem’s first component (hardness of the search version of ring-LWE) applies generically to the ring of integers in *any* number field. Almost all previous work applied to rings of the form $\mathbb{Z}[x]/\langle f(x) \rangle$ for a monic irreducible $f(x)$ having small “expansion” (under the coefficient embedding mentioned above). This set of rings is incomparable to the set used in our work, although for some important examples like cyclotomics, our set is larger.

Rings of integers in number fields have some nice algebraic properties that are useful for our results. For instance, they have unique factorization of ideals, and their fractional ideals form a multiplicative group; in general, neither property holds in $\mathbb{Z}[x]/\langle f(x) \rangle$, even for monic and irreducible $f(x)$ (as demonstrated by the ring $\mathbb{Z}[x]/\langle x^2 + 3 \rangle$). Another useful property is that certain number fields, such as the cyclotomic number fields used in our search/decision reduction, have *automorphisms* that ‘shuffle’ groups of related prime ideals while still preserving

⁴ The m th cyclotomic polynomial in $\mathbb{Z}[x]$ is the polynomial of degree $n = \varphi(m)$ whose roots are the primitive m th roots of unity ζ_m^i for $i \in \mathbb{Z}_m^*$, where $\zeta_m = \exp(2\pi i/m)$.

the LWE error distribution (when appropriately defined using the canonical embedding).

To summarize, while the number-theoretic perspective on ideal lattices 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.

1.4 Techniques

We introduce several new techniques for working with rings of integers and their ideal lattices, which fall into two broad categories: first, those that work in *general* number fields for reducing worst-case problems on ideal lattices to ring-LWE (and related problems); second, those that work in *cyclotomic* number fields, where we demonstrate a search/decision equivalence for ring-LWE and construct cryptographic schemes. All of the new techniques are entirely classical, i.e., non-quantum. (Our main reduction uses existing quantum technology essentially as a black box.)

In the category of worst-case reductions for ideal lattices, we show how to use the Chinese remainder theorem (CRT) for ‘clearing the ideal’ \mathcal{I} from an arbitrary ideal lattice instance. This involves mapping the quotient ring $\mathcal{I}/q\mathcal{I}$ to the fixed quotient ring R/qR in an ‘algebraically consistent’ way. Our CRT techniques are also compatible with the ‘discrete Gaussian’ style of worst-to-average-case reduction from [13], which implies simpler and slightly tighter hardness proofs for ring-SIS. We remark that prior reductions following [23] work by restricting to a *principal subideal* of \mathcal{I} with known generator; however, this technique does not seem to be compatible with the approaches of [31, 13], where the reduction must deal with Gaussian samples from the full ideal \mathcal{I} .

In our search/decision equivalence for ring-LWE, we also 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 — namely, that it ‘splits completely’ into n prime ideals \mathfrak{q}_i each of norm $q = \text{poly}(n)$, which are permuted by the automorphisms. (Interestingly, this complete splitting of q is also useful for performing the ring operations very efficiently in practice; see [22].)

The basic layout of our pseudorandomness proof is as follows: first, a hybrid argument shows that any distinguisher between the ring-LWE distribution $A_{s,\psi}$ and the uniform distribution 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, an efficient search-to-decision reduction finds the value of s modulo \mathfrak{q}_i , using the fact that the ring modulo \mathfrak{q}_i is a field of order $q = \text{poly}(n)$. Then, because the automorphisms of the number field permute the \mathfrak{q}_i s, we can find s modulo *each* \mathfrak{q}_j by applying an appropriate automorphism to the distribution $A_{s,\psi}$. (Crucially, the error distribution ψ also remains legal under this transformation). Finally, we recover all of $s \bmod q$ using the Chinese remainder theorem.

1.5 Related Work

In a concurrent and independent work, Stehlé, Steinfeld, Tanaka, and Xagawa [32] also formulated a variant of LWE over certain polynomial rings and proved its hardness under a worst-case (quantum) assumption. Their main application is a public-key cryptosystem with $\log^{O(1)} n$ encryption and decryption time per message bit. Due to the close similarities between our works, we wish to give a detailed comparison of the approaches and final outcomes.

Stehlé *et al.* [32] give a quantum reduction from the (average-case) ring-SIS problem to (average-case) ring-LWE, by exploiting the duality between the two problems and making new observations about Regev’s quantum machinery [31]. Then, invoking prior worst-case hardness results for ring-SIS [20], they conclude that ring-LWE is hard under a worst-case quantum assumption. More precisely, they show that the *search* version of ring-LWE is hard for an *a priori bounded* number of samples with *spherical* Gaussian noise; however, the approach does not seem to extend to the *decision* version (i.e., pseudorandomness), nor to an unbounded number of samples.

The lack of pseudorandomness has some important drawbacks. For example, a primary motivation for the use of ring-LWE is to encrypt and decrypt faster than the most efficient cryptosystems based on standard LWE. In [32], achieving this goal requires *many* simultaneous hard-core bits for the search variant of ring-LWE, which are obtained via the efficient Goldreich-Levin construction using Toeplitz matrices [14, Section 2.5]. This approach, however, induces a security reduction that runs in time *exponential* in the number of hard bits. Therefore, to encrypt in amortized $\tilde{O}(1)$ time per message bit induces the assumption that ideal-SVP is hard for $2^{\omega(n)}$ -time quantum algorithms. In contrast, our scheme has the same (actually somewhat better) running times under a fully polynomial assumption.

It is also worth noting that the main proof technique from [32], while quite transparent and modular, requires an *a priori* bound on the number of LWE samples consumed, and the modulus q and underlying approximation factor for ideal-SVP grow with this bound. This is suboptimal for cryptographic schemes (such as those in [30, 26, 6, 8]) that use a large (or even unbounded) number of samples in their security proofs. Moreover, having an unbounded number of samples seems essential for proving a search/decision equivalence for any type of LWE problem, because at the very least, the reduction needs to amplify the adversary’s success probability.

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 $p = \infty$.

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

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

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 be 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 j th (complex) coordinate, and 0 elsewhere. Then for $j \in [s_1]$, the basis vector $\mathbf{h}_j = \mathbf{e}_j \in \mathbb{C}^n$; for $s_1 < j \leq s_1 + s_2$, the vector $\mathbf{h}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_j + \mathbf{e}_{j+s_2})$ 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 $\{s_1 + s_2 + 1, \dots, n\}$.

We will also equip H with the ℓ_p norm induced on it from \mathbb{C}^n . We note that for any $p \in [1, \infty]$, this norm is equal within a factor of $\sqrt{2}$ to the ℓ_p norm induced on H from the isomorphism with \mathbb{R}^n described above, and that 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 caveat being the $\sqrt{2}$ factor when dealing with ℓ_p norms for $p \neq 2$.

2.1 Lattice Background

We define a *lattice* as a discrete additive subgroup of H . We deal here 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 $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\} \subset H$.

The *minimum distance* $\lambda_1(A)$ of a lattice A in some norm $\|\cdot\|$ is the length of a shortest nonzero lattice vector: $\lambda_1(A) = \min_{\mathbf{x} \in A, \mathbf{x} \neq \mathbf{0}} \|\mathbf{x}\|$. When left unspecified, the norm is taken to be the Euclidean norm; for the minimum distance of A in the ℓ_p norm, we write $\lambda_1^{(p)}(A)$.

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

Gaussian Measures. For $r > 0$, define the Gaussian function $\rho_r : H \rightarrow (0, 1]$ 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 $\mathbf{r} = (r_1, \dots, r_n) \in (\mathbb{R}^+)^n$ 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 $\sum_{i \in [n]} x_i \cdot \mathbf{h}_i$, where the x_i are chosen independently from the (one-dimensional) Gaussian distribution D_{r_i} over \mathbb{R} .

Micciancio and Regev [24] introduced a lattice quantity called the *smoothing parameter*, and related it to various lattice quantities.

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

Lemma 1 ([24, Lemma 4.1] and [31, Claim 3.8]). For any lattice A , $\epsilon > 0$, $r \geq \eta_\epsilon(A)$, and $\mathbf{c} \in H$, we have $\rho_r(A + \mathbf{c}) \in [\frac{1-\epsilon}{1+\epsilon}, 1] \cdot \rho_r(A)$.

For a lattice A , point $\mathbf{u} \in H$, and real $r > 0$, define the *discrete Gaussian probability distribution* over $A + \mathbf{u}$ with parameter r as the distribution assigning probability proportional to $\rho_r(\mathbf{x})$ to each $\mathbf{x} \in A + \mathbf{u}$.

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 ([31]). *Let Λ be a lattice, let $\mathbf{u} \in H$ be any vector, and let $r, s > 0$ be reals. Assume that $1/\sqrt{1/r^2 + 1/s^2} \geq \eta_\epsilon(\Lambda)$ for some $\epsilon < \frac{1}{2}$. 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 .*

2.2 Algebraic Number Theory Background

Due to space constraints, we assume familiarity with the standard concepts of a number field K , its field trace Tr and norm N , and its ring of integers \mathcal{O}_K , discriminant Δ_K , and group of (fractional) ideals. Details may be found in the full version or in any introductory book on the subject, e.g., [33].

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 field homomorphisms $\sigma_i : K \rightarrow \mathbb{C}$ that fix every element of \mathbb{Q} . Concretely, each embedding takes ζ to a different one of its conjugates; it can be verified that these are the only such field homomorphisms because the conjugates are the only roots of ζ ’s minimal polynomial $f(x)$. 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 and complex *pairs* of embeddings are denoted s_1 and s_2 respectively, so we have $n = s_1 + 2s_2$. The pair (s_1, s_2) is called the *signature* of K . 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* $\sigma : K \rightarrow \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ is defined as $\sigma(x) = (\sigma_1(x), \dots, \sigma_n(x))$. The canonical embedding σ is a field homomorphism from K to $\mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$, where multiplication and addition in $\mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ are both component-wise. Due to the pairing of the complex embeddings, we have that σ maps into H .

By identifying elements 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 distribution $D_{\mathbf{r}}$ (for $\mathbf{r} \in (\mathbb{R}^+)^n$) over H as a distribution over K . Strictly speaking, the

distribution $D_{\mathbf{r}}$ is not quite over K , but rather over the field $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$, which, roughly speaking, is to K as \mathbb{R} is to \mathbb{Q} . Since multiplication of elements in the number field is mapped to coordinate-wise multiplication, we get that for any element $x \in K$, 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 complex and real components of each embedding).

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 $U = \{u_1, \dots, u_n\}$. Therefore, under the canonical embedding σ , the ideal yields a rank- n *ideal lattice* having basis $\{\sigma(u_1), \dots, \sigma(u_n)\} \subset H$. The fundamental volume of the ideal lattice $\sigma(\mathcal{I})$ is $|\det(\sigma(U))| = N(\mathcal{I}) \cdot \sqrt{\Delta_K}$; as expected, this quantity is basis-invariant. 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.

We now 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, e.g., [10].

For a (fractional) ideal \mathcal{I} , its (fractional) *dual ideal* is defined as $\mathcal{I}^{\vee} = \{x \in K : \text{Tr}(x\mathcal{I}) \subset \mathbb{Z}\}$. It is not difficult to see that, under the canonical embedding into H , the dual ideal embeds exactly as the complex conjugate of the dual lattice, i.e., $\sigma(\mathcal{I}^{\vee}) = \overline{\sigma(\mathcal{I})}^*$. This is due to the fact that $\text{Tr}(xy) = \sum_i \sigma_i(x)\sigma_i(y) = \langle \sigma(x), \overline{\sigma(y)} \rangle$.

Except in the trivial number field $K = \mathbb{Q}$, the ring of integers \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* equivalent up to multiplication by the dual ideal of the entire ring. That is, for any fractional ideal \mathcal{I} , its dual ideal is $\mathcal{I}^{\vee} = \mathcal{I}^{-1} \cdot \mathcal{O}_K^{\vee}$. (Notice that for $\mathcal{I} = \mathcal{O}_K$ this holds by definition, since $\mathcal{O}_K^{-1} = \mathcal{O}_K$.) The dual ideal \mathcal{O}_K^{\vee} is itself sometimes called the *codifferent ideal*.

Chinese Remainder Theorem. Here we recall the Chinese remainder theorem (CRT) for the ring of integers in a number field, and some of its important consequences for this work. Let K be an arbitrary fixed number field and let $R = \mathcal{O}_K$ be its ring of integers.

Lemma 3 (Chinese remainder theorem). *Let $\mathcal{I}_1, \dots, \mathcal{I}_r$ be pairwise coprime ideals in R , and let $\mathcal{I} = \prod_{i \in [r]} \mathcal{I}_i$. The natural ring homomorphism $C : R \rightarrow \bigoplus_{i \in [r]} (R/\mathcal{I}_i)$ induces a ring isomorphism $R/\mathcal{I} \rightarrow \bigoplus_{i \in [r]} (R/\mathcal{I}_i)$.*

We state the following important consequences of the CRT; proofs are given in the full version.

Lemma 4. *Let \mathcal{I} and \mathcal{J} be ideals in R . Then 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} .*

Lemma 5. *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 $\theta_t : K \rightarrow K$*

defined as $\theta_i(u) = t \cdot u$ 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 inverted efficiently given \mathcal{I} , \mathcal{J} , \mathcal{M} , and t .

Other Properties of Cyclotomic Number Fields. Here we state a few more useful facts about cyclotomic number fields, which are used only in our search-to-decision reductions of Section 4.

Letting $K = \mathbb{Q}(\zeta)$ for $\zeta = \zeta_m$ be the m th cyclotomic number field, recall that $\mathcal{O}_K = \mathbb{Z}[\zeta]$. For an integer prime $q \in \mathbb{Z}$, the factorization of the ideal $\langle q \rangle$ is as follows. Let the factorization of $\Phi_m(x)$ modulo q (i.e., in $\mathbb{Z}_q[x]$) into monic irreducible polynomials be $\Phi_m(x) = \prod_i (f_i(x))^{e_i}$. Then in \mathcal{O}_K , the prime ideal factorization of $\langle q \rangle$ is $\langle q \rangle = \prod_i \mathfrak{q}_i^{e_i}$, where each $\mathfrak{q}_i = \langle q, f_i(\zeta) \rangle$ is a prime ideal with norm $q^{\deg(f_i)}$.⁵

For an integer prime $q = 1 \pmod m$, the field \mathbb{Z}_q has a primitive m th root of unity r , because the multiplicative group of \mathbb{Z}_q is cyclic with order $q - 1$. Indeed, there are $n = \varphi(m)$ distinct such roots of unity $r^i \in \mathbb{Z}_q$, for $i \in \mathbb{Z}_m^*$. Therefore, the cyclotomic polynomial $\Phi_m(x)$ factors in $\mathbb{Z}_q[x]$ as $\Phi(x) = \prod_{i \in \mathbb{Z}_m^*} (x - r^i)$. The ideal $\langle q \rangle \subset \mathcal{O}_K$ then “splits completely” into n distinct prime ideals, as $\langle q \rangle = \prod_{i \in \mathbb{Z}_m^*} \mathfrak{q}_i$ where $\mathfrak{q}_i = \langle q, \zeta - r^i \rangle$ is prime and has norm q . (The fact that the ideal $\langle q \rangle$ splits into *distinct* prime ideals with *small* norm will be crucial in our search-to-decision for ring-LWE.)

We also need the fact that K has $n = \varphi(m)$ automorphisms $\tau_k : K \rightarrow K$ indexed by $k \in \mathbb{Z}_m^*$, which are defined by $\tau_k(\zeta) = \zeta^k$. The automorphisms τ_j “act transitively” on the prime ideals \mathfrak{q}_i , i.e., $\tau_j(\mathfrak{q}_i) = \mathfrak{q}_{i/j}$. This fact follows directly from the fact that cyclotomic number fields are *Galois extensions* of \mathbb{Q} .

Computation in Number Fields. All of the operations required by our reductions can be performed in polynomial time given a suitable representation of the number field and its ring of integers. Due to space constraints, we defer the details to the full version.

We now define some seemingly hard computational problems on ideal lattices in number fields.

Definition 2. 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 $_\gamma$ problem (in the given norm) is: given a (fractional) ideal \mathcal{I} in K , find some nonzero $x \in \mathcal{I}$ such that $\|x\| \leq \gamma \cdot \lambda_1(\mathcal{I})$.

Definition 3. 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 -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 .

Without loss of generality, both of the above problems may be restricted to *integral* ideals $\mathcal{I} \subseteq \mathcal{O}_K$, 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 $N(d) \cdot \mathcal{I} \subseteq \mathcal{O}_K$, because $N(d) \in \langle d \rangle$.

⁵ In fact, this factorization holds in any ‘monogenic’ ring of integers $\mathcal{O}_K = \mathbb{Z}[\zeta]$, with $\Phi_m(x)$ replaced by the minimal polynomial of ζ .

2.3 The Ring-LWE Problem

Let K be a number field, let $R = \mathcal{O}_K$ be its ring of integers, let R^\vee be its dual (codifferent) ideal, let $q \geq 2$ be a (rational) integer modulus, and let $R_q = R/qR$ and $R_q^\vee = R^\vee/qR^\vee$. Let $\mathbb{T} = K_{\mathbb{R}}/R^\vee$.

For an $s \in R_q^\vee$ and a distribution ψ over $K_{\mathbb{R}}$, the 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 $(a, (a \cdot s)/q + e)$, where addition in the second component is in \mathbb{T} (i.e., modulo R^\vee).

Definition 4 (Learning with Errors in a Ring of Integers). *Let $q \geq 2$ be a (rational) integer and let Ψ be a family of distributions over $K_{\mathbb{R}}$. The ring-LWE problem in $R = \mathcal{O}_K$, denoted $R\text{-LWE}_{q,\Psi}$, 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 .*

For an asymptotic treatment of the ring-LWE problem, we let K come from an infinite sequence of number fields $\mathcal{K} = \{K_n\}$ of increasing dimension n .

A natural question at this point is, why is the secret s chosen from the domain R_q^\vee rather than R_q , as the values a are? From a purely *algebraic* perspective, it is possible to transform the ring-LWE distribution to make s come from the quotient ring $\mathcal{I}/q\mathcal{I}$ for any desired fractional ideal \mathcal{I} , making the choice of domain appear arbitrary. However, from a *geometric* perspective, such a transformation can in general introduce some distortion in the noise distribution. Upon close inspection, there are several reasons why R_q^\vee is the most natural “canonical” domain for s ; due to space constraints, we defer an explanation to the full version.

We now define the exact LWE error distributions for which our results apply. Informally, they are elliptical Gaussians whose widths along each axis (in the canonical embedding) are bounded by some parameter α .

Definition 5. *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 $r_i \leq \alpha$.*

In Section 4.1, we exploit a particular closure property for the family $\Psi_{\leq \alpha}$ over the m th cyclotomic number field $K = \mathbb{Q}(\zeta)$, where $\zeta = \zeta_m$. Let $\tau_j : K \rightarrow K$ be an automorphism of K , which is of the form $\tau_j(\zeta) = \zeta^j$ for some $j \in \mathbb{Z}_m^*$. Then $\Psi_{\leq \alpha}$ is closed under τ_j , i.e., for any $\psi = D_{\mathbf{r}} \in \Psi_{\leq \alpha}$, we have $\tau_j(D_{\mathbf{r}}) = D_{\mathbf{r}'}$ where the entries of \mathbf{r}' are merely a rearrangement of the entries of \mathbf{r} and hence are at most α .

3 Main Reduction

Since the results in this section apply to arbitrary number fields, we choose to present them in their most general form. For concreteness, the reader may wish to keep in mind the particular case of a cyclotomic number field.

Throughout this section, let K denote an arbitrary number field of degree n . We prove that solving the search problem $\mathcal{O}_K\text{-LWE}$ (for a certain family of error distributions) is at least as hard as quantumly solving $K\text{-SVP}_\gamma$, where the approximation factor γ depends on the parameters of the error distributions.

3.1 Main Theorem and Proof Overview

The following is the main theorem of this section. Here, $K\text{-DGS}_\gamma$ denotes the *discrete Gaussian sampling* problem [31], which asks, given an ideal \mathcal{I} and a number $r \geq \gamma$, to produce samples from the distribution $D_{\mathcal{I},r}$. It is easy to show reductions from other more standard lattice problems such as **SVP** to **DGS** (see [31] for some examples).

Theorem 1. *Let K be an arbitrary number field of degree n . Let $\alpha = \alpha(n) \in (0, 1)$ be arbitrary, and let the (rational) integer modulus $q = q(n) \geq 2$ be such that $\alpha \cdot q \geq \omega(\sqrt{\log n})$. There is a probabilistic polynomial-time quantum reduction from $K\text{-DGS}_\gamma$ to $\mathcal{O}_K\text{-LWE}_{q,\Psi_{\leq \alpha}}$, where $\gamma = \eta_\epsilon(\mathcal{I}) \cdot \omega(\sqrt{\log n})/\alpha$.*

We prove the theorem by taking Regev’s iterative reduction for general lattices [31] and replacing its core component (namely, the reduction from the bounded-distance decoding (**BDD**) problem to **LWE**) with an analogous statement for the ideal case (Lemma 7). It is here that we crucially apply algebraic techniques such as the Chinese remainder theorem, and we view this as one of our main contributions.

For self-containment, we describe now the main steps of the iterative reduction of [31], making the necessary changes for our setting. The reduction works by repeated application of an *iterative step*, which consists of the following two components.

1. The first component, which forms the core of [31], is a reduction from **BDD** on the dual lattice to **LWE** that uses Gaussian samples over the primal lattice. In Section 3.2 we show how to perform an analogous reduction in the ring setting. Namely, we show that given an oracle that generates samples from the discrete Gaussian distribution $D_{\mathcal{I},r}$ for some (not too small) $r > 0$, using an $\mathcal{O}_K\text{-LWE}_{q,\Psi_{\leq \alpha}}$ oracle we can solve the **BDD** problem on the dual ideal \mathcal{I}^\vee to within distance $d = \alpha \cdot q/r$ in the ℓ_∞ norm. From this it follows that with probability negligibly close to 1, we can also solve **BDD** on \mathcal{I}^\vee where the unknown offset vector e is drawn from the distribution $D_{d'}$ for $d' = \alpha \cdot q/(r \cdot \omega(\sqrt{\log n}))$. The reason is that a sample from D_s has ℓ_∞ norm at most $s \cdot \omega(\sqrt{\log n})$, except with negligible probability.
2. The second step is quantum, and is essentially identical to the corresponding step in Regev’s reduction [31, Lemma 3.14]. This step uses an oracle that with all but negligible probability solves the **BDD** problem on \mathcal{I}^\vee , where the offset vector e is chosen from the distribution $D_{d'}$. Using a quantum procedure, it is shown in [31] how to use such an oracle to produce samples from the discrete Gaussian distribution $D_{\mathcal{I},r'}$ for $r' = 1/(2d')$. The exact statement of [31, Lemma 3.14] is more specialized; to get the above statement, one has to observe that the procedure used to prove the lemma calls the oracle with offset vectors e chosen from $D_{d'}$, and that correctness is maintained even if the oracle errs with negligible probability over the choice of e .

Notice that when $\alpha \cdot q \geq \omega(\sqrt{\log n})$, we can choose $r' \leq r/2$ so that the output distribution $D_{\mathcal{I},r'}$ of Step 2 is half as wide as the input distribution $D_{\mathcal{I},r}$

of Step 1. The value of r starts out exponentially large so that the samples for the first execution of Step 1 can be generated classically (see [31, Lemma 3.2]), then in later phases of the iteration they are produced by the quantum part. By iterating back and forth between the two procedures, we can sample from a progressively tighter distribution until we obtain a sample from $D_{\mathcal{I},r}$ for the (typically small) r given as input to the DGS problem.

3.2 Core Step: The BDD to LWE Reduction

We first observe that to solve BDD on an ideal \mathcal{J} , it suffices to find the solution modulo $q\mathcal{J}$. This is actually a special case of a lemma from [31], which gives a *lattice-preserving* reduction for BDD in general lattices. Because the reduction is lattice-preserving, it also applies to ideal lattices.

Definition 6. *The q -BDD $_{\mathcal{J},d}$ problem (in any norm) is: given an instance y of BDD $_{\mathcal{J},d}$ that has solution $x \in \mathcal{J}$, find $x \bmod q\mathcal{J}$.*

Lemma 6 (Special case of [31, Lemma 3.5]). *For any $q \geq 2$, there is a deterministic polynomial-time reduction from BDD $_{\mathcal{J},d}$ (in any ℓ_p norm) to q -BDD $_{\mathcal{J},d}$ (in the same norm).*

Lemma 7. *Let $\alpha \in (0, 1)$, let $q \geq 2$ be a (rational) integer with known factorization, let \mathcal{I} be an ideal in $R = \mathcal{O}_K$, and let $r \geq \sqrt{2}q \cdot \eta_\epsilon(\mathcal{I})$ for some negligible $\epsilon = \epsilon(n)$. Given an oracle for the discrete Gaussian distribution $D_{\mathcal{I},r}$, there is a probabilistic polynomial-time (classical) reduction from q -BDD $_{\mathcal{I}^\vee,d}$ (in the ℓ_∞ norm) to R -LWE $_{q,\Psi_{\leq \alpha}}$, where $d = \alpha q / (\sqrt{2}r)$.*

Note that the hypothesis that \mathcal{I} is an integral ideal (in \mathcal{O}_K) is without loss of generality, by the scaling argument at the end of Section 2.2.

Proof. The high-level description of the reduction is as follows. Its input is a q -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 -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 4, 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) , 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 5 with $\mathcal{J} = \langle q \rangle$ and $\mathcal{M} = R$, the function $\theta_t(u) = t \cdot u$ induces a bijection from $R_q = R/qR$ to $\mathcal{I}/q\mathcal{I}$ which may be inverted efficiently given \mathcal{I} , q , and t .)

3. When \mathcal{L} produces a solution $s \in R_q^\vee$, output $\theta_t^{-1}(s) \in \mathcal{I}^\vee/q\mathcal{I}^\vee$.
 (Again, by Lemma 5 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 $\mathcal{I}^\vee/q\mathcal{I}^\vee$ to $R_q^\vee = R^\vee/qR^\vee$, which may be inverted efficiently).

The correctness of the reduction follows from Lemma 8 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 -BDD $_{\mathcal{I}^\vee, d}$ input instance.

Lemma 8. *Let y be the 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 2ϵ -uniform. Because $r \geq q \cdot \eta_\epsilon(\mathcal{I})$, each sample z from $D_{\mathcal{I}, r}$ is 2ϵ -uniform in $\mathcal{I}/q\mathcal{I}$ by Lemma 1. Then because θ_t induces a bijection from $R_q = R/qR$ to $\mathcal{I}/q\mathcal{I}$ by Lemma 5, $a = \theta_t^{-1}(z \bmod q\mathcal{I})$ is 2ϵ -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 $z = \theta_t(a) = a \cdot t \in \mathcal{I}/q\mathcal{I}$. 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 a coset of \mathcal{I} (specifically, $u = \theta_t(a) \bmod q\mathcal{I}$) and $r/q \geq \sqrt{2} \cdot \eta_\epsilon(\mathcal{I})$. Note that

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

so we may apply Lemma 9 below, which 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 9. *Let \mathcal{I} be a (fractional) ideal in K , and let $r \geq \sqrt{2} \cdot \eta_\epsilon(\mathcal{I})$ for some $\epsilon = \text{negl}(n)$. 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. 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, the distribution of $z + f$ is negligibly close to $D_{\sqrt{2}r}$, so $(z + e'/e) = (z + f + g)$ has distribution $D_{\mathbf{t}''}$, 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 $D_{\mathbf{r}}$, as desired.

4 Pseudorandomness of Ring-LWE

In this section we show that for an appropriate choice of ring, modulus, and error distribution, the ring-LWE distribution is pseudorandom. For concreteness and simplicity, we specialize the discussion to cyclotomic fields (though our techniques generalize somewhat to others). So throughout this section we assume that $\zeta = \zeta_m \in \mathbb{C}$ is a primitive m th root of unity, $K = \mathbb{Q}(\zeta)$ is the m th cyclotomic number field, $R = \mathcal{O}_K$ is its ring of integers, $R^\vee = \mathcal{O}_K^\vee$ is its dual (codifferent) ideal, and $q = 1 \bmod m$ is a poly(n)-bounded prime.

The main goal of this section is to show that the following average-case problem is hard. (Recall that $R_q = R/qR$, $R_q^\vee = R^\vee/qR^\vee$, and $\mathbb{T} = K_{\mathbb{R}}/R^\vee$.)

Definition 7 (Distinguishing LWE). *For a distribution Υ over a family of noise distributions (each over $K_{\mathbb{R}}$), we say that an algorithm solves the $\text{DLWE}_{q,\Upsilon}$ problem if its acceptance probability given samples from $A_{s,\psi}$, over the random choice of $(s, \psi) \leftarrow U(R_q^\vee) \times \Upsilon$ and all other randomness of the experiment, differs by a non-negligible amount from its acceptance probability on uniformly random samples from $R_q \times \mathbb{T}$.*

The following is the main theorem of this section. It shows a reduction from the worst-case search variant of LWE (which by Theorem 1 is as hard as a worst-case lattice problem) to the above average-case problem. 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 both chosen at random from appropriate distributions (and kept secret).

Theorem 2. *Let R, m, q be as above, and let $\alpha \cdot q \geq 1 \geq \eta_{2^{-n}}(R^\vee)$. Then there is a randomized polynomial-time reduction from $\text{LWE}_{q,\psi_{\leq \alpha}}$ to $\text{DLWE}_{q,\Upsilon_\alpha}$.*

The proof of Theorem 2 goes through a chain of reductions, summarized in the following diagram (the numbers refer to lemma numbers, and the definitions of all intermediate problems are given later).

$$\begin{array}{ccccccc} \text{LWE}_{q,\psi} & \xrightarrow[10]{\text{Automorphisms}} & \mathfrak{q}_i\text{-LWE}_{q,\psi} & \xrightarrow[11]{\text{Search to Decision}} & \text{DecLWE}_{q,\psi}^i & & \\ & \xrightarrow{\text{Worst to Average}} & \text{DecLWE}_{q,\Upsilon}^i & \xrightarrow{\text{Amplification}} & \text{DLWE}_{q,\Upsilon}^i & \xrightarrow{\text{Hybrid}} & \text{DLWE}_{q,\Upsilon} \end{array}$$

This sequence of reductions is similar in spirit to the one given in previous work on the (non-ideal) LWE problem [31]. 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 ideal factors using the field’s automorphisms (see Lemma 10) to recover s relative to *every* factor \mathfrak{q}_j .

Another issue arises from the fact that the reduction in Section 3 only establishes the hardness of $\text{LWE}_{q,\Psi}$ for a family of distributions Ψ that contains *non-spherical* Gaussian distributions. As a result, the average-case problem requires a distribution \mathcal{T} over Gaussian noise distributions that are both non-spherical and wider by a factor of \sqrt{n} . Although this is somewhat undesirable, we do not yet see any way to avoid it; luckily, this only has a minor effect on the resulting cryptographic applications, i.e., adding an extra step of choosing the noise parameters. Let us mention, though, that if one is willing to assume that (the search problem) $\text{LWE}_{q,\Psi}$ is hard with *spherical* Gaussian noise distributions, then we can fix the noise distribution in all the average-case problems (so there is no need to use a distribution over noise distributions \mathcal{T}) and we do not need to lose the factor \sqrt{n} .

Due to space constraints, here we present only the first two steps of the chain of reductions, which contain the bulk of the novel ideas. The rest can be found in the full version.

4.1 Worst-Case Search to Worst-Case Decision

In this subsection we reduce the search version of $\text{LWE}_{q,\Psi}$ to a certain decision problem over 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 a family of allowed error distributions (though the actual error terms drawn from ψ are still random), and require their solutions to be found with overwhelming probability (over all the randomness of the experiment).

We first define some intermediate computational problems and probability distributions, then present two reductions. For notational convenience, we identify the elements of \mathbb{Z}_m^* with their integer representatives from the set $\{1, \dots, m-1\}$, 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.

We define the notation $R_{\mathfrak{q}_i}^\vee = R^\vee / \mathfrak{q}_i R^\vee$, and note that by Lemmas 3 and 5, there is an efficiently computable R -module isomorphism between R_q^\vee and $\bigoplus_i R_{\mathfrak{q}_i}^\vee$.

Definition 8 (LWE over \mathfrak{q}_i). *The \mathfrak{q}_i -LWE $_{q,\Psi}$ problem is: given access to $A_{s,\psi}$ for some arbitrary $s \in R_q^\vee$ and $\psi \in \Psi$, find $s \in R_{\mathfrak{q}_i}^\vee$.*

Definition 9 (Hybrid LWE distribution). *For $i \in \mathbb{Z}_m^*$, $s \in R_q^\vee$, and a distribution ψ over K , 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 $(a, b + r/q)$ where $r \in R_q^\vee$ is uniformly random and independent in $R_{\mathfrak{q}_j}^\vee$ for all $j \leq i$, and is 0 in all the remaining $R_{\mathfrak{q}_j}^\vee$. Also define $A_{s,\psi}^0$ simply as $A_{s,\psi}$.*

Definition 10 (Decision LWE relative to \mathfrak{q}_i). For $i \in \mathbb{Z}_m^*$ and a family of distributions Ψ , the $\text{DeclWE}_{q,\Psi}^i$ 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 .

Claim. For any $i \in \mathbb{Z}_m^*$ there exists an efficient procedure that transforms $A_{s,\psi}^k$ (for any unknown $k \in \mathbb{Z}_m^* \cup \{0\}$, s , and ψ) into $A_{s,\psi}^{\max\{i,k\}}$.

Lemma 10 (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 $\text{LWE}_{q,\Psi}$ to \mathfrak{q}_i - $\text{LWE}_{q,\Psi}$.

Proof. To compute $s \in R_{\mathfrak{q}_j}^\vee$, we use the oracle for \mathfrak{q}_i -LWE along with the field automorphisms τ_k to recover the value $s \in R_{\mathfrak{q}_j}^\vee$ for all $j \in \mathbb{Z}_m^*$. We can then efficiently reconstruct $s \in R_q^\vee$.

The reduction that finds $s \in R_{\mathfrak{q}_j}^\vee$ works as follows: transform each sample $(a, b) \leftarrow A_{s,\psi}$ into the sample $(\tau_k(a), \tau_k(b))$, where $k = j/i \in \mathbb{Z}_m^*$ and hence $\tau_k(\mathfrak{q}_j) = \mathfrak{q}_i$. (Also note that because τ_k is an automorphism on K and R is the set of all its algebraic integers, $\tau_k(R) = R$ and $\tau_k(R^\vee) = R^\vee$.) Give the transformed samples to the \mathfrak{q}_i - $\text{LWE}_{q,\Psi}$ oracle, and when the oracle returns some element $t \in R_{\mathfrak{q}_i}^\vee$, return $\tau_k^{-1}(t) \in R_{\mathfrak{q}_j}^\vee$.

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

$$\tau_k(b) = \tau_k(a)\tau_k(s)/q + \tau_k(e).$$

Because τ_k is an automorphism on R , $\tau_k(a) \in R_q$ is uniformly random, and because $\psi' = \tau_k(\psi) \in \Psi$, the pairs $(\tau_k(a), \tau_k(b))$ are distributed according to $A_{\tau_k(s), \psi'}$. By hypothesis, the oracle returns $t = \tau_k(s) \in R_{\mathfrak{q}_i}^\vee$. Thus $\tau_k^{-1}(t) = s \in \tau_k^{-1}(R_{\mathfrak{q}_i}^\vee) = R_{\mathfrak{q}_j}^\vee$.

Lemma 11 (Search to Decision). For any $i \in \mathbb{Z}_m^*$, there is a probabilistic polynomial-time reduction from \mathfrak{q}_i - $\text{LWE}_{q,\Psi}$ to $\text{DeclWE}_{q,\Psi}^i$.

Proof. The idea for recovering $s \in R_{\mathfrak{q}_i}^\vee$ 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 $\text{DeclWE}_{q,\Psi}^i$ oracle to tell us which distribution was generated. Because there are only $N(\mathfrak{q}_i) = q = \text{poly}(n)$ possible values for $s \in R_{\mathfrak{q}_i}^\vee$, we can enumerate over all of them efficiently and discover the correct value.

First note that by Claim 4.1, we can transform our input distribution $A_{s,\psi}$ to $A_{s,\psi}^{i-}$. We now give the transformation that takes some $g \in R_q^\vee$ and maps $A_{s,\psi}^{i-}$ to either $A_{s,\psi}^{i-}$ or $A_{s,\psi}^i$, depending on whether or not $g = s \in R_{\mathfrak{q}_i}^\vee$ (its values in the other $R_{\mathfrak{q}_j}^\vee$ are irrelevant). Given a sample $(a, b) \leftarrow A_{s,\psi}^{i-}$, the transformation produces a sample

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

where $v \in R_q$ is uniformly random modulo \mathfrak{q}_i and is 0 modulo the other \mathfrak{q}_j . 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

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

where e is chosen from ψ , and r is distributed as in the definition of $A_{s,\psi}^{i-}$, i.e., it is uniformly random and independent modulo \mathfrak{q}_j for all $j < i$, and is 0 modulo all the remaining \mathfrak{q}_j .

We consider two cases. First, assume that $g = s \in R_{\mathfrak{q}_i}^\vee$. Then by the Chinese remainder theorem (Lemma 3), $v(g - s) = 0 \in R_q^\vee$, and hence the distribution of (a', b') is exactly $A_{s,\psi}^{i-}$. Next, assume that $g \neq s \pmod{\mathfrak{q}_i}$. Then since \mathfrak{q}_i is a maximal ideal (which in R is equivalent to being a prime ideal), $R_{\mathfrak{q}_i}^\vee$ is a field, and hence $v(g - s)$ is distributed uniformly in $R_{\mathfrak{q}_i}^\vee$ (and is zero in all other $R_{\mathfrak{q}_j}^\vee$). From this it follows that $v(g - s) + r$ is distributed uniformly random and independently in $R_{\mathfrak{q}_j}^\vee$ for all $j \leq i$, and is 0 in all the remaining $R_{\mathfrak{q}_j}^\vee$. Hence, the distribution of (a', b') is exactly $A_{s,\psi}^i$, as promised.

References

- [1] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, 2010. To appear.
- [2] Miklós Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in *STOC* 1996.
- [3] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610, 2001.
- [4] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *TCC*, pages 474–495, 2009.
- [5] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. In *STACS*, pages 75–86, 2009.
- [6] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618, 2009.
- [7] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In *CRYPTO*, pages 278–291, 1993.
- [8] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, 2010. To appear.
- [9] Henri Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1993.
- [10] Keith Conrad. The different ideal, 2009. Available at <http://www.math.uconn.edu/~kconrad/blurbs/>, last accessed 12 Oct 2009.
- [11] Paul Erdős. On the coefficients of the cyclotomic polynomial. *Bulletin of the American Mathematical Society*, 52(2):179–184, 1946.
- [12] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.

- [13] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [14] Oded Goldreich. *Foundations of Cryptography*, volume II. Cambridge University Press, 2004.
- [15] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(42), 1996.
- [16] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288, 1998.
- [17] Akinori Kawachi, Keisuke Tanaka, and Keita Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In *ASIACRYPT*, pages 372–389, 2008.
- [18] Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Public Key Cryptography*, pages 162–179, 2008.
- [19] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
- [20] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155, 2006.
- [21] Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, pages 37–54, 2008.
- [22] Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72, 2008.
- [23] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.
- [24] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.
- [25] Daniele Micciancio and Salil P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO*, pages 282–298, 2003.
- [26] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342, 2009.
- [27] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006.
- [28] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487, 2007.
- [29] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.
- [30] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196, 2008.
- [31] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009. Preliminary version in STOC 2005.
- [32] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635, 2009.
- [33] William 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.