

A Toolkit for Ring-LWE Cryptography

Vadim Lyubashevsky^{1*}, Chris Peikert^{2**}, and Oded Regev^{3***}

¹ INRIA and École Normale Supérieure, Paris

² Georgia Institute of Technology

³ Courant Institute, New York University

Abstract. Recent advances in lattice cryptography, mainly stemming from the development of ring-based primitives such as ring-LWE, have made it possible to design cryptographic schemes whose efficiency is competitive with that of more traditional number-theoretic ones, along with entirely new applications like fully homomorphic encryption. Unfortunately, realizing the full potential of ring-based cryptography has so far been hindered by a lack of practical algorithms and analytical tools for working in this context. As a result, most previous works have focused on very special classes of rings such as power-of-two cyclotomics, which significantly restricts the possible applications.

We bridge this gap by introducing a toolkit of fast, modular algorithms and analytical techniques that can be used in a wide variety of ring-based cryptographic applications, particularly those built around ring-LWE. Our techniques yield applications that work in *arbitrary* cyclotomic rings, with *no loss* in their underlying worst-case hardness guarantees, and very little loss in computational efficiency, relative to power-of-two cyclotomics. To demonstrate the toolkit’s applicability, we develop two illustrative applications: a public-key cryptosystem and a “somewhat homomorphic” symmetric encryption scheme. Both apply to arbitrary cyclotomics, have tight parameters, and very efficient implementations.

1 Introduction

The past few years have seen many exciting developments in lattice-based cryptography. Two such trends are the development of schemes whose efficiency is competitive with traditional number-theoretic ones (e.g., [27] and follow-ups), and

* Part of this work was performed while at Tel Aviv University and also while visiting Georgia Tech. Partially supported by a European Research Council Starting Grant.

** This material is based upon work supported by the National Science Foundation under CAREER Award CCF-1054495, by DARPA under agreement number FA8750-11-C-0096, 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, DARPA or the U.S. Government, or the Sloan Foundation.

*** Supported by a European Research Council Starting Grant. Part of the work done while the author was with the CNRS, DI, ENS, Paris.

the breakthrough work of Gentry [14, 13] (followed by others) on fully homomorphic encryption. While these two research threads currently occupy opposite ends of the efficiency spectrum, they are united by their use of algebraically structured *ideal lattices* arising from polynomial rings. The most efficient and advanced systems in both categories rely on the ring-LWE problem [26], an analogue of the standard *learning with errors* problem [31]. Informally (and a bit inaccurately), in a ring $R = \mathbb{Z}[X]/(f(X))$ for monic irreducible $f(X)$ of degree n , and for an integer modulus q defining the quotient ring $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$, the ring-LWE problem is to distinguish pairs $(a_i, b_i = a_i \cdot s + e_i) \in R_q \times R_q$ from uniformly random pairs, where $s \in R_q$ is a random secret (which stays fixed over all pairs), the $a_i \in R_q$ are uniformly random and independent, and the error (or “noise”) terms $e_i \in R$ are independent and “short.”

In all applications of ring-LWE, and particularly those related to homomorphic encryption, a main technical challenge is to control the sizes of the noise terms when manipulating ring-LWE samples under addition, multiplication, and other operations. For correct decryption, q must be chosen large enough so that the final accumulated error terms do not “wrap around” modulo q and cause decryption error. On the other hand, the *error rate* (roughly, the ratio of the noise magnitude to the modulus q) of the original published ring-LWE samples and the dimension n trade off to determine the theoretical and concrete hardness of the ring-LWE problem. Tighter control of the noise growth therefore allows for a larger initial error rate, which permits a smaller modulus q and dimension n , which leads to smaller keys and ciphertexts, and faster operations for a given level of security.

Regarding the choice of ring, the class of *cyclotomic* rings $R \cong \mathbb{Z}[X]/\Phi_m(X)$, where $\Phi_m(X)$ is the m th cyclotomic polynomial (which has degree $n = \varphi(m)$ and is monic and irreducible over the rationals), has many attractive features that have proved very useful in cryptography. For example, the search/decision equivalence for ring-LWE in arbitrary cyclotomics [26] relies on their special algebraic properties, as do many recent works that aim for more efficient fully homomorphic encryption schemes (e.g., [32, 8, 17, 18, 16]). In particular, *power-of-two* cyclotomics, i.e., where the index $m = 2^k$ for some $k \geq 1$, are especially nice to work with, because (among other reasons) $n = m/2$ is also a power of two, $\Phi_m(X) = X^n + 1$ is maximally sparse, and polynomial arithmetic modulo $\Phi_m(X)$ can be performed very efficiently using just a slight tweak of the classical n -dimensional FFT (see, e.g., [25]). Indeed, power-of-two cyclotomics have become the dominant and preferred class of rings in almost all recent ring-based cryptographic schemes (e.g., [25, 24, 21, 14, 15, 26, 33, 9, 8, 17, 18, 22, 5, 28, 20, 16]), often to the exclusion of all other rings.

While power-of-two cyclotomic rings are very convenient to use, there are several reasons why it is essential to consider other cyclotomics as well. The most obvious, practical reason is that powers of two are sparsely distributed, and the desired concrete security level for an application may call for a ring dimension much smaller than the next-largest power of two. So restricting to powers of two could lead to key sizes and runtimes that are at least twice as large as necessary. A more fundamental reason is that certain applications, such as the above-

mentioned works that aim for more efficient (fully) homomorphic encryption, *require* the use of non-power-of-two cyclotomic rings. This is because power-of-two cyclotomics lack the requisite algebraic properties needed to implement features like SIMD operations on “packed” ciphertexts, or plaintext spaces isomorphic to finite fields of characteristic two (other than \mathbb{F}_2 itself). A final important reason is diversification of security assumptions. While some results are known [16] that relate ring-LWE in cyclotomic rings when one index m divides the other, no other connections appear to be known. So while we might conjecture that ring-LWE and ideal lattice problems are hard in *every* cyclotomic ring (of sufficiently high dimension), some rings might turn out to be significantly easier than others.

Unfortunately, working in non-power-of-two cyclotomics is rather delicate, and the current state of affairs is unsatisfactory in several ways. Unlike the special case where m is a power of two, in general the cyclotomic polynomial $\Phi_m(X)$ can be quite “irregular” and dense, with large coefficients. While in principle, polynomial arithmetic modulo $\Phi_m(X)$ can still be done in $O(n \log n)$ scalar operations (on high-precision complex numbers), the generic algorithms for achieving this are rather complex and hard to implement, with large constants hidden by the $O(\cdot)$ notation.

Geometrically, the non-power-of-two case is even more problematic. If one views $\mathbb{Z}[X]/(\Phi_m(X))$ as the set of polynomial residues of the form $a_0 + a_1X + \dots + a_{n-1}X^{n-1}$, and uses the naïve “coefficient embedding” that views them as vectors $(a_0, a_1, \dots, a_{n-1}) \in \mathbb{Z}^n$ to define geometric quantities like the ℓ_2 norm, then both the concrete and theoretical security of cryptographic schemes depend heavily on the form of $\Phi_m(X)$. This stems directly from the fact that multiplying two polynomials with small norms can result in a polynomial residue having a *much* larger norm. The growth can be quantified by the “expansion factor” [23] of $\Phi_m(X)$, which unfortunately can be very large, up to $n^{\Omega(\log n)}$ in the case of highly composite m [12]. Later works [17] circumvented such large expansion by using tricks like lifting to the larger-dimensional ring $\mathbb{Z}[X]/(X^m - 1)$, but this still involves a significant loss in the tolerable noise rates as compared with the power-of-two case.

In [30, 26] a different geometric approach was used, which avoided any dependence on the form of the polynomial modulus $\Phi_m(X)$. In these works, the norm of a ring element is instead defined according to its *canonical embedding* into \mathbb{C}^n , a classical concept from algebraic number theory. This gives a much better way of analyzing expansion, since both addition and multiplication in the canonical embedding are simply coordinate-wise. Working with the canonical embedding, however, introduces a variety of practical issues, such as how to efficiently generate short noise terms having appropriate distributions over the ring. More generally, the focus of [26] was on giving an abstract mathematical definition of ring-LWE and proving its hardness under worst-case ideal lattice assumptions; in particular, it did not deal with issues related to practical efficiency, bounding noise growth, or designing applications in non-power-of-two cyclotomics.

1.1 Contributions

Our main contribution is a toolkit of modular algorithms and analytical techniques that can be used in a wide variety of ring-based cryptographic applications, particularly those built around ring-LWE. The high-level summary is that using our techniques, one can design applications to work in *arbitrary* cyclotomic rings, with *no loss* in their underlying worst-case hardness guarantees, and very little loss in computational efficiency, relative to the best known techniques in power-of-two cyclotomics. In fact, our analytical techniques even improve the state of the art for the power-of-two case.

In more detail, our toolkit includes fast, specialized algorithms for all the main cryptographic operations in arbitrary cyclotomic rings. Among others, these include: addition, multiplication, and conversions among various useful representations of rings elements; generation of noise terms under probability distributions that guarantee both worst-case and concrete hardness; and decoding of noise terms as needed in decryption and related operations. Our algorithms' efficiency and quality guarantees stem primarily from our use of simple but non-obvious representations of ring elements, which differ from their naïve representations as polynomial residues modulo $\Phi_m(X)$. (See the second part of Section 1.2 for more details.) On the analytical side, we give tools for tightly bounding noise growth under operations like addition, multiplication, and round-off/discretization. (Recall that noise growth is the main factor determining an application's parameters and noise rates, and hence its key sizes, efficiency, and concrete security.)

Some attractive features of the toolkit include:

- All the algorithms for arbitrary cyclotomics are simple, modular, and highly parallel, and work by elementary reductions to the (very simple) prime-index case. In particular, they do not require any polynomial reductions modulo $\Phi_m(X)$ – in fact, they never need to compute $\Phi_m(X)$ at all! The algorithms work entirely on vectors of dimension $n = \varphi(m)$, and run in $O(n \log n)$ or even $O(nd)$ scalar operations (with small hidden constants), where d is the number of distinct primes dividing m . With the exception of continuous noise generation, all scalar operations are low precision, i.e., they involve small integers. In summary, the algorithms are very amenable to practical implementation. (Indeed, we have implemented all the algorithms from scratch, which will be described in a separate work.)
- Our algorithm for decoding noise, used primarily in decryption, is fast (requiring $O(n \log n)$ or fewer small-integer operations) and correctly recovers from optimally large noise rates. (See the last part of Section 1.2 for details.) This improves upon prior techniques, which in general have worse noise tolerance by anywhere between an $m/2$ and super-polynomial $n^{\omega(1)}$ factor, and are computationally slower and more complex due to polynomial reduction modulo $\Phi_m(X)$, among other operations.
- Our bounds on noise growth under ring addition and multiplication are exactly the same in *all* cyclotomic rings; no ring-dependent “expansion factor” is incurred. (For discretizing continuous noise distributions, our bounds are the same up to very small $1 + o(1)$ factors, depending on the primes dividing m .)

This allows applications to use essentially the same underlying noise rate as a function of the ring dimension n , and hence be based on the same worst-case approximation factors, for all cyclotomics. Moreover, our bounds improve upon the state of the art even for power-of-two cyclotomics: e.g., our (average-case, high probability) expansion bound for ring multiplication improves upon the (worst-case) expansion-factor bound by almost a \sqrt{n} factor.

To illustrate the toolkit’s applicability, in Section 5 we construct an efficient and compact public-key cryptosystem, which is essentially the “two element” system outlined in [26], but generalized to arbitrary cyclotomics, and with tight parameters. Further applications are given in the full version of the paper.

A final contribution of independent interest is a new “regularity lemma” for arbitrary cyclotomics, i.e., a bound on the smoothing parameter of random q -ary lattices over the ring. Such a lemma is needed for porting many applications of standard SIS and LWE to the ring setting, including SIS-based signature schemes [19, 10, 7, 28], the “primal” [31] and “dual” [19] LWE cryptosystems, chosen ciphertext-secure encryption schemes [29, 28], and (hierarchical) identity-based encryption schemes [19, 10, 1]. In terms of generality and parameters, our lemma essentially subsumes a prior one of Micciancio [27] for the ring $\mathbb{Z}[X]/(X^n - 1)$, and an independent one of Stehlé et al. [34] for power-of-two cyclotomics. (See Section 4 for further discussion.)

1.2 Techniques

The tools we develop in this work involve several novel applications of classical notions from algebraic number theory. In summary, our results make central use of: (1) the *canonical embedding* of a number field, which endows the field (and its subrings) with a nice and easy-to-analyze geometry; (2) the decomposition of arbitrary cyclotomics into the *tensor product* of prime-power cyclotomics, which yields both simpler and faster algorithms for computing in the field, as well as geometrically nicer bases; and (3) the “*dual*” ideal R^\vee and its “*decoding*” basis \vec{d} , for fast noise generation and optimal noise tolerance in decryption and related operations. We elaborate on each of these next.

The canonical embedding. As in the previous works [30, 26], our analysis relies heavily on using the *canonical embedding* $\sigma: K \rightarrow \mathbb{C}^n$ (rather than, say, the naïve coefficient embedding) for defining all geometric quantities, such as Euclidean norms and inner products. For example, under the canonical embedding, the “expansion” incurred when multiplying by an element $a \in K$ is characterized exactly by $\|\sigma(a)\|_\infty$, its ℓ_∞ norm under the canonical embedding; no (worst-case) ring-dependent “expansion factor” is needed. So in the average-case setting, where the multiplicands are random elements from natural noise distributions, for each multiplication we get at least a $\tilde{\Omega}(\sqrt{n})$ factor improvement over using the expansion factor in *all* cyclotomics (including those with power-of-two index), and up to a super-polynomial $n^{\omega(1)}$ factor improvement in cyclotomics having

highly composite indices. In our analysis of the noise tolerance of decryption, we also get an additional $\tilde{\Omega}(\sqrt{n})$ factor savings over more simplistic analyses that only use norm information, by using the notion of *subgaussian* random variables. These behave under linear transformations in essentially the same way as Gaussians, and have Gaussian tails. (This builds upon prior works that use subgaussianity in lattice cryptography, e.g., [2, 28].)

Tensorial decomposition. An important fact at the heart of this work is that the m th cyclotomic number field $K = \mathbb{Q}(\zeta_m) \cong \mathbb{Q}[X]/(\Phi_m(X))$ may instead be viewed as (i.e., is isomorphic to) the *tensor product* of prime-power cyclotomics:

$$K \cong \bigotimes_{\ell} K_{\ell} = \mathbb{Q}(\zeta_{m_1}, \zeta_{m_2}, \dots),$$

where $m = \prod_{\ell} m_{\ell}$ is the prime-power factorization of m and $K_{\ell} = \mathbb{Q}(\zeta_{m_{\ell}})$. Equivalently, in terms of polynomials we may view K as the multivariate field

$$K \cong \mathbb{Q}[X_1, X_2, \dots]/(\Phi_{m_1}(X_1), \Phi_{m_2}(X_2), \dots), \quad (1)$$

where there is one indeterminant X_{ℓ} and modulus $\Phi_{m_{\ell}}(X_{\ell})$ per prime-power divisor of m . Similar decompositions hold for the ring of integers $R \cong \mathbb{Z}[X]/\Phi_m(X)$ and other important objects in K , such as the dual ideal R^{\vee} (described below).

Adopting the polynomial interpretation of K from Equation (1) for concreteness, notice that a natural \mathbb{Q} -basis is the set of multinomials $\prod_{\ell} X_{\ell}^{j_{\ell}}$ for each choice of $0 \leq j_{\ell} < \varphi(m_{\ell})$. We call this set the “powerful” basis of K (and of R). For non-prime-power m , under the field isomorphism with $\mathbb{Q}[X]/(\Phi_m(X))$ that maps each $X_{\ell} \rightarrow X^{m/m_{\ell}}$, the powerful basis does *not* coincide with the standard “power” basis $1, X, X^2, \dots, X^{\varphi(m)-1}$ usually used to represent the univariate field. It turns out that in general, the powerful basis has much nicer computational and geometric properties than the power basis, as we outline next.

Computationally, the tensorial decomposition of K (with the powerful basis) allows us to modularly reduce essentially all operations in K (or R , or powers of R^{\vee}) to their counterparts in much simpler prime-power cyclotomics (which themselves easily reduce to the prime-index case). We can therefore completely avoid all the many algorithmic complications associated with working with polynomials modulo $\Phi_m(X)$. In particular, we obtain novel, simple and fast algorithms, similar to the FFT, for converting between the multivariate “polynomial” representation (i.e., the powerful basis) and the “evaluation” or “Chinese remainder” representation, in which addition and multiplication are essentially linear time. Similarly, we obtain linear-time (or nearly so) algorithms for switching between the polynomial representation and “decoding” representation used in decryption (described below), and for generating noise terms in the decoding representation. A final advantage of the tensorial representation is that it yields trivial linear-time algorithms for computing the *trace* function to subfields of K , which is at the heart of the “ring-switching” technique from [16].

The tensorial representation also comes with important geometrical advantages. In particular, under the canonical embedding the powerful basis is better-conditioned than the power basis, i.e., the ratio of its maximal and minimal

singular values can be much smaller. This turns out to be important when bounding the additional error introduced when discretizing (rounding off) field elements in noise-generation and modulus-reduction algorithms, among others.

The dual ideal R^\vee and its decoding basis. Under the canonical embedding, the cyclotomic ring R of index m embeds as a lattice which, unlike \mathbb{Z}^n , is in general not self-dual. Instead, its dual lattice corresponds to a fractional ideal $R^\vee \subset K$ satisfying $R \subseteq R^\vee \subseteq m^{-1}R$, where the latter inclusion is nearly an equality. (In fact, R^\vee is a scaling of R exactly when m is a power of two, in which case $R = (m/2)R^\vee$.) In [26] it is shown that the “right” definition of the ring-LWE distribution, which arises naturally from the worst-case to average-case reduction, involves the dual ideal R^\vee : the secret belongs to the quotient $R_q^\vee = R^\vee/qR^\vee$, and ring-LWE samples are of the form $(a, b = a \cdot s + e \bmod qR^\vee)$ for uniformly random $a \in R_q$ and error e which is essentially spherical in the canonical embedding.

While it is possible [11] to simplify the ring-LWE definition by replacing every instance of R^\vee with R , while retaining essentially spherical error (but scaled up by about m , corresponding to the approximate ratio of R to R^\vee), in this work we show that *it is actually advantageous to retain R^\vee and expose it in applications.*⁴ The reason is that in general, R^\vee supports correct bounded-distance decoding—which is the main operation performed in decryption—under a larger error rate than R does.⁵ In fact, R^\vee ’s error tolerance is *optimal* for the simple, fast lattice decoding algorithm used implicitly in essentially all decryption procedures, namely Babai’s “round-off” algorithm [4]. The reason is that when decoding a lattice Λ using some basis $\{\mathbf{b}_i\}$, the error tolerance depends inversely on the Euclidean lengths of the vectors dual to $\{\mathbf{b}_i\}$. For R^\vee , there is a particular “*decoding*” basis whose dual basis is optimally short (relative to the determinant of R), whereas for R no such basis exists in general.⁶ In fact, the decoding basis of R^\vee is simply the dual of the powerful basis described above!

In addition to its optimal error tolerance, we also show that the decoding basis has good computational properties. In particular, there are linear-time (or nearly so) algorithms for converting to the decoding basis from the other bases of R^\vee or R_q^\vee that are more appropriate for other computational tasks. And Gaussian errors (especially spherical ones) can be sampled in (near-)linear time in the decoding basis.

Acknowledgments. We thank Markus Püschel for his help with the sparse decomposition of the “Chinese remainder transform,” and Damien Stehlé for useful discussions.

⁴ This is unless m is a power of two, in which case nothing is lost by simply scaling up by exactly $m/2$ to replace R^\vee with R .

⁵ By “error rate” here we mean the ratio of the error (in, say, ℓ_2 norm) to the dimension-normalized determinant $\det(\Lambda)^{1/n}$ of the lattice Λ , so exact scaling has no effect on the error rate.

⁶ We note that decoding by “lifting” R to the larger-dimensional ring $\mathbb{Z}[X]/(X^m - 1)$, as done in [17], still leads to an m or $m/2$ factor loss in error tolerance overall, because some inherent loss is already incurred when replacing R^\vee with R , and a bit more is lost in the lifting procedure.

2 Preliminaries

For a positive integer k , we denote by $[k]$ the set $\{0, \dots, k-1\}$. For a real $a \in \mathbb{R}$, define $\lfloor a \rfloor = \lfloor a + \frac{1}{2} \rfloor \in \mathbb{Z}$. For any $\bar{a} \in \mathbb{R}/\mathbb{Z}$, we let $\llbracket \bar{a} \rrbracket \in \mathbb{R}$ denote the unique representative $a \in (\bar{a} + \mathbb{Z}) \cap [-1/2, 1/2)$. Similarly, for $\bar{a} \in \mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ we let $\llbracket \bar{a} \rrbracket$ denote the unique representative $a \in (\bar{a} + q\mathbb{Z}) \cap [-q/2, q/2)$. We extend $\lfloor \cdot \rfloor$ and $\llbracket \cdot \rrbracket$ entrywise to vectors and matrices. The *radical* of a positive integer m , denoted $\text{rad}(m)$, is the product of all primes dividing m . We also define $\hat{m} = m/2$ whenever m is even, and $\hat{m} = m$ otherwise. For a vector \mathbf{x} over \mathbb{R} or \mathbb{C} , we define the ℓ_2 norm as $\|\mathbf{x}\|_2 = (\sum_i |x_i|^2)^{1/2}$, and the ℓ_∞ norm as $\|\mathbf{x}\|_\infty = \max_i |x_i|$. When the subscript is omitted, we mean the ℓ_2 norm.

Throughout this paper, the entries of a vector over a domain D are always indexed (in no particular order) by some finite set S , and we write D^S to denote the set of all such vectors. Similarly, the rows and columns of an “ R -by- C matrix” over D are indexed by some finite sets R and C , respectively. All the standard matrix and vector operations, including the Kronecker (or tensor) product, are defined in the natural way.

2.1 The Space H

When working with cyclotomic number fields and ideal lattices, it is convenient to work with the subspace $H \subseteq \mathbb{C}^{\mathbb{Z}_m^*}$ for integer $m \geq 2$, defined as

$$H = \{\mathbf{x} \in \mathbb{C}^{\mathbb{Z}_m^*} : x_i = \overline{x_{m-i}}, \forall i \in \mathbb{Z}_m^*\}.$$

Letting $n = \varphi(m)$, it is not difficult to verify that H (with the inner product induced on it by $\mathbb{C}^{\mathbb{Z}_m^*}$) is isomorphic to $\mathbb{R}^{[n]}$ as an inner product space. For $m = 2$ this is trivial, and for $m > 2$ this can be seen via the \mathbb{Z}_m^* -by- $[n]$ unitary basis matrix $\frac{1}{\sqrt{2}} \begin{pmatrix} I & \sqrt{-1}J \\ J & -\sqrt{-1}I \end{pmatrix}$ of H , where here the \mathbb{Z}_m^* -indexed rows are in increasing order according to their canonical representatives in $\{1, \dots, m-1\}$, the $[n]$ -indexed columns are in increasing order by index, I is the identity matrix, and J is the reversal matrix (obtained by reversing the rows of I).

We equip H with the ℓ_2 and ℓ_∞ norms induced on it from $\mathbb{C}^{\mathbb{Z}_m^*}$. Namely, for $\mathbf{x} \in H$ we have $\|\mathbf{x}\|_2 = \sum_i (|x_i|^2)^{1/2} = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$, and $\|\mathbf{x}\|_\infty = \max_i |x_i|$.

2.2 Gaussians and Subgaussian Random Variables

For $s > 0$, define the Gaussian function $\rho_s : H \rightarrow (0, 1]$ as $\rho_s(\mathbf{x}) = \exp(-\pi \|\mathbf{x}\|^2 / s^2)$. By normalizing this function we obtain the *continuous* Gaussian probability distribution D_s of parameter s , whose density is given by $s^{-n} \cdot \rho_s(\mathbf{x})$.

For much of our analysis it is convenient to use the standard notion of *subgaussian* random variables, relaxed slightly as in [28]. For any $\delta \geq 0$, we say that a random variable X (or its distribution) over \mathbb{R} is δ -*subgaussian* with parameter $s > 0$ if for all $t \in \mathbb{R}$, the (scaled) moment-generating function satisfies

$$\mathbb{E}[\exp(2\pi t X)] \leq \exp(\delta) \cdot \exp(\pi s^2 t^2).$$

Notice that the $\exp(\pi s^2 t^2)$ term on the right is exactly the (scaled) moment-generating function of the one-dimensional Gaussian distribution of parameter s over \mathbb{R} .

Decoding In many applications we need to perform the following algorithmic task, which is essentially that of bounded-distance decoding. Let Λ be a known fixed lattice, and let $\mathbf{x} \in H$ be an unknown short vector. The goal is to recover \mathbf{x} , given $\mathbf{x} \bmod \Lambda$. Although there are several possible algorithms for this task, here we focus on a slight extension of the so-called “round-off” algorithm, originally due to Babai [4]. This is due to its high efficiency and because for our lattices it performs optimally (or nearly so). The algorithm is very simple: let $\{\mathbf{v}_i\}$ be a fixed set of n short, linearly independent vectors in the dual lattice Λ^\vee . Denote the dual vectors of $\{\mathbf{v}_i\}$ by $\{\mathbf{b}_i\}$, and let $\Lambda' \supseteq \Lambda$ be the (super)lattice generated by $\{\mathbf{b}_i\}$. Given an input $\mathbf{t} = \mathbf{x} \bmod \Lambda$, we express $\mathbf{t} \bmod \Lambda'$ in the basis $\{\mathbf{b}_i\}$ as $\sum_i \bar{a}_i \mathbf{b}_i$ where $\bar{a}_i \in \mathbb{R}/\mathbb{Z}$ (so $\bar{a}_i = \langle \mathbf{x}, \mathbf{v}_i \rangle \bmod 1$), and output $\sum_i \llbracket \bar{a}_i \rrbracket \mathbf{b}_i \in H$.

Lemma 2.1. *Let $\Lambda \subset H$ be a lattice, let $\{\mathbf{v}_i\}$ be a set of n linearly independent vectors in its dual, and let $d_{\max} = \max_i \|\mathbf{v}_i\|$. For any \mathbf{x} of length less than $1/(2d_{\max})$, the above round-off algorithm succeeds in recovering \mathbf{x} from $\mathbf{x} \bmod \Lambda$. Moreover, for any $\delta > 0$, if \mathbf{x} is a random vector such that $\langle \mathbf{x}, \mathbf{v}_i \rangle$ is δ -subgaussian with parameter s for every i (in particular, if \mathbf{x} itself is δ -subgaussian with parameter s/d_{\max}), then the round-off algorithm succeeds with probability at least $1 - 2n \exp(\delta - \pi/(2s)^2)$, which is $1 - \text{negl}(n)$ when $\delta = O(1)$ and $s = 1/\omega(\sqrt{\log n})$.*

Discretization We now consider another algorithmic task related to the one in the previous subsection. This task shows up in applications, such as when converting a continuous Gaussian into a discrete Gaussian-like distribution. Given a lattice $\Lambda = \mathcal{L}(\mathbf{B})$ represented by a “good” basis $\mathbf{B} = \{\mathbf{b}_i\}$, a point $\mathbf{x} \in H$, and a point $\mathbf{c} \in H$ representing a lattice coset $\Lambda + \mathbf{c}$, the goal is to discretize \mathbf{x} to a point $\mathbf{y} \in \Lambda + \mathbf{c}$, written $\mathbf{y} \leftarrow \lfloor \mathbf{x} \rfloor_{\Lambda + \mathbf{c}}$, so that the length (or subgaussian parameter) of $\mathbf{y} - \mathbf{x}$ is not too large. To do this, we sample a relatively short offset vector \mathbf{f} from the coset $\Lambda + \mathbf{c}' = \Lambda + (\mathbf{c} - \mathbf{x})$, and output $\mathbf{y} = \mathbf{x} + \mathbf{f}$. We require that the method used to choose \mathbf{f} be efficient and depend only on the desired coset $\Lambda + \mathbf{c}'$, not on the particular representative used to specify it. In the full version of the paper, we describe several valid ways of sampling \mathbf{f} , offering tradeoffs between efficiency and output guarantees.

2.3 Algebraic Number Theory Background

Cyclotomic Number Fields and Polynomials. For a positive integer m , the m th cyclotomic number field is a field extension $K = \mathbb{Q}(\zeta_m)$ obtained by adjoining an element ζ_m of order m (i.e., a primitive m th root of unity) to the rationals. (Note that we view ζ_m as an abstract element, and not, for example, as any particular value in \mathbb{C} .) The minimal polynomial of ζ_m is the m th cyclotomic polynomial

$$\Phi_m(X) = \prod_{i \in \mathbb{Z}_m^*} (X - \omega_m^i) \in \mathbb{Z}[X], \quad (2)$$

where $\omega_m \in \mathbb{C}$ is any primitive m th root of unity in \mathbb{C} , e.g., $\omega_m = \exp(2\pi\sqrt{-1}/m)$. Therefore, there is a natural isomorphism between K and $\mathbb{Q}[X]/(\Phi_m(X))$, given by $\zeta_m \mapsto X$. Since $\Phi_m(X)$ has degree $n = |\mathbb{Z}_m^*| = \varphi(m)$, we can view K as a vector space of degree n over \mathbb{Q} , which has $\{1, \zeta_m, \dots, \zeta_m^{n-1}\}$ as a basis. This is called the *power basis* of K .

For the m th cyclotomic number field $K = \mathbb{Q}(\zeta_m)$ of degree $n = \varphi(m)$, the *ring of integers* is $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$, and hence has the power basis $\{\zeta_m^j\}_{j \in [n]}$ as a \mathbb{Z} -basis.

Non-Prime-Power Cyclotomics. Let m have prime-power factorization $m = \prod_{\ell} m_{\ell}$, i.e., the m_{ℓ} are powers of distinct primes. Then $K = \mathbb{Q}(\zeta_m)$ may be seen as the *tensor product* $\otimes_{\ell} K_{\ell}$ of the fields $K_{\ell} = \mathbb{Q}(\zeta_{m_{\ell}})$, in the following way. First, view each K_{ℓ} as a subfield of K , via the ring embedding $\zeta_{m_{\ell}} \mapsto \zeta_m^{m/m_{\ell}}$. Then viewing K and K_{ℓ} as vector spaces over \mathbb{Q} , the tensor product $\otimes_{\ell} K_{\ell}$ is isomorphic to K , under the map $(\otimes_{\ell} a_{\ell}) \mapsto \prod_{\ell} a_{\ell}$.⁷ In particular, if B_{ℓ} are \mathbb{Q} -bases of K_{ℓ} respectively (e.g., the power bases), then their tensor product $\otimes_{\ell} B_{\ell} = \{\prod_{\ell} b_{\ell} \in K : b_{\ell} \in B_{\ell}\}$ is a \mathbb{Q} -basis of K . Moreover, endowing $\otimes_{\ell} K_{\ell}$ with the multiplication operation induced by the mixed-product property $(\otimes_{\ell} a_{\ell}) \cdot (\otimes_{\ell} b_{\ell}) = \otimes_{\ell} (a_{\ell} \cdot b_{\ell})$ also makes the above mapping from $\otimes_{\ell} K_{\ell}$ to K a field isomorphism, as desired.

Equivalently, in terms of polynomial rings we may view $K \cong \mathbb{Q}[X]/(\Phi_m(X))$ instead as

$$K \cong \mathbb{Q}[X_1, X_2, \dots]/(\Phi_{m_1}(X_1), \Phi_{m_2}(X_2), \dots), \quad (3)$$

where there is one indeterminant X_{ℓ} and modulus $\Phi_{m_{\ell}}(X_{\ell})$ per prime divisor of m , and where $X_{\ell} \mapsto X^{m/m_{\ell}}$ defines an isomorphism with $\mathbb{Q}[X]/(\Phi_m(X))$. Notice that by tensoring the power bases $\{X_{\ell}^j\}_{j \in [\varphi(m_{\ell})]}$ of each K_{ℓ} , we get the basis $\{X_1^{j_1} X_2^{j_2} \dots\}_{j_{\ell} \in [\varphi(m_{\ell})]}$. Mapping this basis to $\mathbb{Q}[X]/(\Phi_m(X))$ yields the basis $\{X^{\sum_{\ell} (m/m_{\ell})j_{\ell}}\}_{j_{\ell} \in [\varphi(m_{\ell})]}$, which is *not* necessarily the power basis $\{X^j\}_{j \in [\varphi(m)]}$, since the powers of X appearing in each basis can be different modulo m . (For example, take $m = 3 \cdot 5$.)

Embeddings and Geometry. Here we describe the *embeddings* of a cyclotomic number field, which induce a ‘canonical’ geometry on it.

The m th cyclotomic number field $K = \mathbb{Q}(\zeta_m)$ of degree $n = \varphi(m)$ has exactly n ring homomorphisms (embeddings) $\sigma_i: K \rightarrow \mathbb{C}$ that fix every element of \mathbb{Q} . Concretely, for each $i \in \mathbb{Z}_m^*$ there is an embedding σ_i defined by $\sigma_i(\zeta_m) = \omega_m^i$, where $\omega_m \in \mathbb{C}$ is some fixed primitive m th root of unity. Clearly, the embeddings come in pairs of complex conjugates, i.e., $\sigma_i = \overline{\sigma_{m-i}}$. The *canonical embedding* $\sigma: K \rightarrow \mathbb{C}^{\mathbb{Z}_m^*}$ is defined as

$$\sigma(a) = (\sigma_i(a))_{i \in \mathbb{Z}_m^*}.$$

⁷ The tensor product of two vector spaces K, L over a common base field can be defined as the set of all finite sums of pure tensors $a \otimes b$ for $a \in K, b \in L$, where \otimes is bilinear. The tensor product of multiple vector spaces is defined similarly.

When K is viewed as the tensor product of subfields K_ℓ , $\sigma = \bigotimes_\ell \sigma^{(\ell)}$ is the tensor product of the canonical embeddings $\sigma^{(\ell)}$ of K_ℓ . In this case, the index set of σ is $\prod_\ell \mathbb{Z}_{m_\ell}^*$, which corresponds to \mathbb{Z}_m^* via the Chinese remainder theorem.

The *trace* $\text{Tr} = \text{Tr}_{K/\mathbb{Q}}: K \rightarrow \mathbb{Q}$ can be defined as the sum of the embeddings: $\text{Tr}(a) = \sum_i \sigma_i(a)$. Clearly, $\text{Tr}(a + b) = \text{Tr}(a) + \text{Tr}(b)$ and $\text{Tr}(c \cdot a) = c \cdot \text{Tr}(a)$ for all $a, b \in K$ and $c \in \mathbb{Q}$. Moreover,

$$\text{Tr}(a \cdot b) = \sum_i \sigma_i(a)\sigma_i(b) = \langle \sigma(a), \overline{\sigma(b)} \rangle.$$

Duality. For any fractional ideal \mathcal{I} in K , its *dual* is defined as

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

It is easy to verify that \mathcal{I}^\vee is a fractional ideal, and that $(\mathcal{I}^\vee)^\vee = \mathcal{I}$.

For any \mathbb{Q} -basis $B = \{b_j\}$ of K , we denote its dual basis by $B^\vee = \{b_j^\vee\}$, which is characterized by $\text{Tr}(b_i \cdot b_j^\vee) = 1$ if $i = j$, and 0 otherwise. It is immediate that $(B^\vee)^\vee = B$, and if B is a \mathbb{Z} -basis of some fractional ideal \mathcal{I} , then B^\vee is a \mathbb{Z} -basis of its dual ideal \mathcal{I}^\vee . An important fact is that if $a = \sum_j a_j \cdot b_j$ (where $a_j \in \mathbb{R}$) is the unique representation of some $a \in K_\mathbb{R}$ in basis B , then $a_j = \text{Tr}(a \cdot b_j^\vee)$ by linearity of Tr .

Except in the trivial number field $K = \mathbb{Q}$, the ring of integers R is not self-dual, nor are an ideal and its inverse dual to each other. However, an ideal and its inverse *are* related by multiplication with the dual ideal R^\vee of the ring: for any fractional ideal \mathcal{I} , its dual is $\mathcal{I}^\vee = \mathcal{I}^{-1} \cdot R^\vee$. (Notice that for $\mathcal{I} = R$ this holds trivially, since $R^{-1} = R$.) A standard fact is that $R^\vee = \langle t^{-1} \rangle$ is a principal ideal generated by t^{-1} for some (non-unique) $t \in R$. When $R \cong \bigotimes_\ell R_\ell$ is viewed as the tensor product of rings of integers $R_\ell \subset K_\ell$ (where $K \cong \bigotimes_\ell K_\ell$), its dual ideal has an analogous tensorial form, as $R^\vee = \bigotimes_\ell R_\ell^\vee$.

2.4 Ring-LWE

We now provide the formal definition of the ring-LWE problem and recall the worst-case hardness result shown in [26]. We remark that our definition here differs very slightly from the one used in [26]: we scale the b component by a factor of q , so that it is an element of $K_\mathbb{R}/qR^\vee$ and not $K_\mathbb{R}/R^\vee$ as in [26]. This is done for convenience when later discretizing the b component, and the two definitions are easily seen to be equivalent.

Definition 2.2 (Ring-LWE Distribution). For a “secret” $s \in R_q^\vee$ (or just R^\vee) and a distribution ψ over $K_\mathbb{R}$, a sample from the ring-LWE distribution $A_{s,\psi}$ over $R_q \times (K_\mathbb{R}/qR^\vee)$ is generated by choosing $a \leftarrow R_q$ uniformly at random, choosing $e \leftarrow \psi$, and outputting $(a, b = a \cdot s + e \bmod qR^\vee)$.

Definition 2.3 (Ring-LWE, Average-Case Decision). The average-case decision version of the ring-LWE problem, denoted $R\text{-DLWE}_{q,\psi}$, is to distinguish with non-negligible advantage between independent samples from $A_{s,\psi}$, where $s \leftarrow R_q^\vee$ is uniformly random, and the same number of uniformly random and independent samples from $R_q \times (K_\mathbb{R}/qR^\vee)$.

Theorem 2.4. *Let K be the m th cyclotomic number field having dimension $n = \varphi(m)$ and $R = \mathcal{O}_K$ be its ring of integers. Let $\alpha = \alpha(n) > 0$, and let $q = q(n) \geq 2$, $q = 1 \pmod m$ be a $\text{poly}(n)$ -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 the problem of solving R -DLWE $_{q,\psi}$ given only ℓ samples, where ψ is the Gaussian distribution $D_{\xi q}$ for $\xi = \alpha \cdot (n\ell / \log(n\ell))^{1/4}$.*

In cryptographic applications it is often useful to work with a version of ring-LWE whose error distribution is discrete. In the full version of the paper, we show that for a wide family of discrete error distributions, it is easy to deduce the hardness of the discrete version from that of the continuous one. Another important variant of ring-LWE, known as the “normal form,” is the one in which the secret, instead of being uniformly distributed, is chosen from the error distribution (discretized to R^\vee). Showing that this variant of ring-LWE is as hard as the original one follows from the techniques of [3].

3 The Powerful, CRT, and Decoding Bases

In this section we study certain \mathbb{Z} -bases of certain (fractional) ideals \mathcal{I} in $K = \mathbb{Q}(\zeta_m)$, which are also \mathbb{Z}_q -bases of the quotients $\mathcal{I}_q = \mathcal{I}/q\mathcal{I}$ for any positive integer q . Fixing such a basis \vec{b} and viewing it as a (column) vector over K , we can represent any $a \in \mathcal{I}$ uniquely as $a = \langle \vec{b}, \mathbf{a} \rangle = \vec{b}^T \cdot \mathbf{a}$ for some coefficient vector \mathbf{a} over \mathbb{Z} . Similarly, any $\bar{a} \in \mathcal{I}_q$ is represented uniquely as $\bar{a} = \langle \vec{b}, \bar{\mathbf{a}} \rangle$ for some $\bar{\mathbf{a}}$ over \mathbb{Z}_q . Our algorithms that work with field elements simply store and operate on these coefficient vectors, while also keeping track of the corresponding basis, which will be among the few we consider below. Notice that by linearity, if we have some $a \in \mathcal{I}$ represented by coefficient vector \mathbf{a} in basis \vec{b} , then \mathbf{a} is also the representation of $ra \in r\mathcal{I}$ in the basis $r\vec{b}$, so we can switch between the two values at essentially no cost.

3.1 The Powerful Basis

Here we define a certain useful \mathbb{Q} -basis of K , and \mathbb{Z} -basis of R . We call it the “powerful” basis, due to its decomposition in terms of the power bases of K_ℓ , and the fast algorithms associated with it. (We are aware of only one occurrence in the literature of this basis; it coincides with what Bosma [6] calls the “canonical” basis of R .)

Definition 3.1. *The powerful basis \vec{p} of $K = \mathbb{Q}(\zeta_m)$ and $R = \mathbb{Z}[\zeta_m]$ is defined as follows:*

- For a prime power m , define \vec{p} to be the power basis $(\zeta_m^j)_{j \in [\varphi(m)]}$, treated as a vector over $R \subset K$.
- For m having prime-power factorization $m = \prod_\ell m_\ell$, define $\vec{p} = \bigotimes_\ell \vec{p}_\ell$, the tensor product of the power(ful) bases \vec{p}_ℓ of each $K_\ell = \mathbb{Q}(\zeta_{m_\ell})$.

For any power $\mathcal{I} = (R^\vee)^k$ of $R^\vee = \langle t^{-1} \rangle$, the powerful basis of \mathcal{I} is $t^{-k} \cdot \vec{p}$.

By definition of the tensor product, \vec{p} is a vector with index set $\prod_\ell [\varphi(m_\ell)]$. So to specify an entry of \vec{p} we need one index $j_\ell \in [\varphi(m_\ell)]$ per prime divisor of m , and the specified entry is $p_{(j_\ell)} = \prod_\ell \zeta_{m_\ell}^{j_\ell}$. Note that because $\zeta_{m_\ell} = \zeta_m^{m/m_\ell} \in K$, it is possible to “flatten” the index set to a size- $\varphi(m)$ subset of $[m]$, where index (j_ℓ) maps to $j = \sum_\ell (m/m_\ell) \cdot j_\ell \bmod m$, and $p_j = \zeta_m^j$. We note that unless m is a prime power, the flattened index set is *not* $[\varphi(m)]$, so the powerful basis differs from the power basis, although it still consists of powers of ζ_m . For instance, for $m = 15$ and $\zeta = \zeta_{15}$, the powerful basis consists of $\zeta^0, \zeta^3, \zeta^5, \zeta^6, \zeta^8, \zeta^9, \zeta^{11}$, and ζ^{14} . For our purposes, it is preferable to maintain the structured index set.

In the full version we prove the following lemma describing the good geometric properties of the powerful basis.

Lemma 3.2. *The length of each element p_j of \vec{p} in ℓ_2 norm is $\|p_j\| = \sqrt{\varphi(m)} = \sqrt{n}$, and in ℓ_∞ norm is $\|p_j\|_\infty = 1$. The largest singular value of $\sigma(\vec{p}^T)$ is $s_1(\vec{p}) = \sqrt{\hat{m}}$, and the smallest singular value is $s_n(\vec{p}) = \sqrt{m/\text{rad}(m)}$, where $\hat{m} = m/2$ if m is even, and $\hat{m} = m$ otherwise.*

We point out while the *power* basis elements also all have ℓ_2 and ℓ_∞ norms \sqrt{n} and 1 (respectively), the power basis can be poorly conditioned. E.g., for $m = 1155 = 3 \cdot 5 \cdot 7 \cdot 11$ its ratio of largest to smallest singular value is $\approx 21.4\sqrt{m}$, whereas for the powerful basis it is exactly \sqrt{m} .

3.2 The CRT Basis and Fast Operations

In ring-LWE and its applications, we work in R_q and R_q^\vee , and sometimes in \mathcal{I}_q for $\mathcal{I} = (R^\vee)^k$, where $q = 1 \bmod m$ is a prime integer. Here we define Chinese remainder (CRT) bases for these quotients, and describe how they yield fast addition and multiplication.

Recalling that $R \cong \bigotimes_\ell R_\ell$ where $m = \prod_\ell m_\ell$ is the prime-power factorization of m and R_ℓ is the m_ℓ th cyclotomic ring, it is easy to verify that the quotient ring $R_q \cong \bigotimes_\ell (R_\ell/qR_\ell)$. Therefore we may focus on the case of prime-power m . A standard fact is that the ideal $\langle q \rangle \subset R$ factors into the product of n distinct prime ideals \mathfrak{q}_i , for $i \in \mathbb{Z}_m^*$.

Definition 3.3. *For a positive integer m , the Chinese remainder (or CRT) \mathbb{Z}_q -basis \vec{c} of R_q is as follows:*

- For a prime power m , $\vec{c} = (c_i)_{i \in \mathbb{Z}_m^*}$ is characterized by $c_i = 1 \bmod \mathfrak{q}_i$ and $c_i = 0 \bmod \mathfrak{q}_j$ for $i \neq j$. (Its existence is guaranteed by the Chinese Remainder Theorem.)
- For m having prime-power factorization $m = \prod_\ell m_\ell$, define $\vec{c} = \bigotimes_i \vec{c}_\ell$, the tensor product of the CRT bases \vec{c}_ℓ of each R_ℓ/qR_ℓ .

For any power $\mathcal{I} = (R^\vee)^k$ of $R^\vee = \langle t^{-1} \rangle$, the CRT \mathbb{Z}_q -basis of \mathcal{I}_q is $t^{-k} \cdot \vec{c}$.

Similarly to the powerful basis, \vec{c} is a vector over R_q having the Cartesian product $\prod_\ell \mathbb{Z}_{m_\ell}^*$ as its index set, which may be flattened to the set \mathbb{Z}_m^* using the bijective correspondence $(j_\ell) \leftrightarrow j = \sum_\ell (m/m_\ell) \cdot j_\ell \in \mathbb{Z}_m^*$. But it is usually more convenient to retain the structured index set.

In the full version of the paper we give a novel, fast ‘‘CRT transformation’’ algorithm for converting between the powerful and CRT bases of R_q , or more generally \mathcal{I}_q for $\mathcal{I} = (R^\vee)_q^k$. The algorithm is analogous to a combination of the Cooley-Tukey and Good-Thomas (mixed radix) FFT algorithms, but specialized to evaluate at only the *primitive* m th roots of unity in a ring. The algorithm is simpler and more efficient than converting between the *power* and CRT bases, which involves reducing modulo the cyclotomic polynomial $\Phi_m(X)$.

Working in the CRT basis yields very fast arithmetic operations. Suppose that m is a prime power. Since $c_i^2 = c_i \in R_q$ and $c_i \cdot c_{i'} = 0 \in R_q$ for distinct $i, i' \in \mathbb{Z}_m^*$, the CRT basis has the property that if $a, b \in R_q$ have coefficient vectors \mathbf{a}, \mathbf{b} (respectively) over \mathbb{Z}_q in the CRT basis—i.e., $a = \langle \vec{c}, \mathbf{a} \rangle$ and $b = \langle \vec{c}, \mathbf{b} \rangle$ —then the coefficient vector of $a \cdot b \in R_q$ is the componentwise product $\mathbf{a} \odot \mathbf{b}$ over \mathbb{Z}_q . (Addition is componentwise as well, simply by linearity.) Moreover, this extends immediately to powers of R^\vee : if \mathbf{a}, \mathbf{b} are the respective coefficient vectors of $a \in (R^\vee)_q^{k_1}$, $b \in (R^\vee)_q^{k_2}$ in the respective CRT bases $t^{-k_1} \cdot \vec{c}$ and $t^{-k_2} \cdot \vec{c}$, then $\mathbf{a} \odot \mathbf{b}$ is the coefficient vector of $a \cdot b \in (R^\vee)_q^k$ in the CRT basis $t^{-k} \cdot \vec{c}$, where $k = k_1 + k_2$.

3.3 The Decoding Basis of R^\vee

When working with ring-LWE we need to perform a variety of operations over $R^\vee = \langle t^{-1} \rangle$ or R_q^\vee . For certain operations it is best to use the following important \mathbb{Z} -basis of R^\vee (and \mathbb{Z}_q -basis of R_q^\vee).

Definition 3.4. *The decoding basis of R^\vee is $\vec{d} = \vec{p}^\vee$, the dual of the powerful basis \vec{p} of R .⁸*

The decoding basis therefore has the same index set as \vec{p} . When m is a prime power, \vec{d} is simply the dual of the power basis $\vec{p} = (\zeta_m^j)_{j \in [\varphi(m)]}$ of R . In general, because \vec{p} is the tensor product of the power bases for prime-power cyclotomics R_ℓ , and $(\vec{a} \otimes \vec{b})^\vee = (\vec{a}^\vee \otimes \vec{b}^\vee)$, it follows that \vec{d} is the tensor product of the decoding bases for each R_ℓ^\vee .

In the full version of the paper, we prove several important and useful properties of the decoding basis, summarized as follows:

- There are very fast linear transformations (requiring fewer than nd scalar additions, where d is the number of prime divisors of m) for converting between the decoding basis \vec{d} and the powerful basis $t^{-1}\vec{p}$ of R^\vee .

⁸ Note that unlike the powerful and CRT bases, we do not define a decoding basis for any other power of R^\vee ; see Section 3.4 for discussion.

- Short elements (as always, in the sense of the canonical embedding) of K have optimally small coefficients with respect to \vec{d} , making it a best choice for decoding R^\vee . Moreover, \vec{d} also yields (nearly) optimal decoding in higher powers of R^\vee .
- Continuous Gaussians (especially spherical ones) as represented in the decoding basis can be sampled very simply and efficiently.

The first fact, combined with the fast CRT transformation, means that we can efficiently convert among the decoding, power, and CRT bases of R^\vee (or R_q^\vee) as needed. The latter two facts mean that the decoding basis is an excellent choice for generating and decoding error terms (e.g., in encryption and decryption, respectively). By contrast, the power(ful) basis and other natural bases of R or R^\vee do not typically enjoy the above properties (except when m is a power of 2), and while they can in principle be used for all the same tasks, it would come at a potentially large loss in tightness and/or computational efficiency.

3.4 Decoding R^\vee and Its Powers

Recall from Section 2.2 the “round-off” decoding procedure, which uses short linearly independent vectors in a dual lattice Λ^\vee to recover a sufficiently short \mathbf{x} given $\mathbf{x} \bmod \Lambda$. To decode from K/R^\vee to K , we apply the procedure using the decoding basis \vec{d} of R^\vee ; i.e., the linearly independent dual elements (in $(R^\vee)^\vee = R$) are those of the powerful basis \vec{p} . Recall from Lemma 2.1 that the tolerable decoding distance (or subgaussian parameter) depends inversely on the maximum length of the dual elements, and that by Lemma 3.2, every p_j in the powerful basis has length $\|p_j\| = \sqrt{n}$. From this we get corresponding bounds on the decoding operation, as summarized below in Lemma 3.6. We remark that the decoding basis is an optimal choice here.

In some applications (e.g., homomorphic encryption), we need to solve the more general problem of decoding K/\mathcal{I} to K , where $\mathcal{I} = (R^\vee)^k = \langle t^{-k} \rangle$ for some (usually small) $k \geq 1$. The naïve way to do this would be to apply the round-off procedure with the \mathbb{Z} -basis $t^{1-k}\vec{d}$ of \mathcal{I} . This, however, turns out to be highly suboptimal for many values of m , because the elements of the dual basis $t^{k-1}\vec{p}$ might be much longer than the shortest nonzero elements of $\mathcal{I}^\vee = \langle t^{k-1} \rangle$.

Instead, in the round-off algorithm we use the *scaled decoding basis* $\hat{m}^{1-k}\vec{d}$, which generates the superideal $\mathcal{J} = \hat{m}^{1-k}R^\vee = t^{-k}g^{1-k} \supseteq \mathcal{I}$, and whose dual elements are $\hat{m}^{k-1}\vec{p} \subseteq \mathcal{I}^\vee$. (Recall that $\hat{m} = m/2$ if m is even, and $\hat{m} = m$ otherwise. It is easy to show that $\hat{m} = t \cdot g$ for some $g \in R$; see the full version.) The lengths of the dual elements are therefore $\hat{m}^{k-1}\sqrt{n}$, from which one gets the bounds summarized in Lemma 3.6 below.

We summarize the above discussion in the following definition and lemma. As it will be more convenient for applications, here we consider a “scaled up and discretized” version of the decoding procedure, where we decode from \mathcal{I}_q to \mathcal{I} for some $q \geq 1$. So the unknown short element is guaranteed to be in \mathcal{I} , and the output is also expected to be in \mathcal{I} . The only difference this makes in the above procedure (apart from the obvious scaling by q) is that for $k \geq 2$, since the

scaled decoding basis may generate a strict superideal of \mathcal{I} , when the round-off procedure fails to decode correctly it might produce an element that is not in \mathcal{I} . In such a case we just consider the output to be undefined.

Definition 3.5 (Decoding \mathcal{I}_q to \mathcal{I}). For $\bar{a} \in \mathcal{I}_q$ where $\mathcal{I} = (R^\vee)^k$ for some $k \geq 1$, let $\bar{a} = \langle \hat{m}^{1-k} \vec{d}, \bar{\mathbf{a}} \rangle \bmod q\mathcal{J}$ for some $\bar{\mathbf{a}}$ over \mathbb{Z}_q , where $\mathcal{J} = \hat{m}^{1-k} R^\vee$. Define $\llbracket \bar{a} \rrbracket$ to be $\langle \hat{m}^{1-k} \vec{d}, \llbracket \bar{\mathbf{a}} \rrbracket \rangle$ if it is in \mathcal{I} , otherwise $\llbracket \bar{a} \rrbracket$ is undefined (where $\llbracket \bar{\mathbf{a}} \rrbracket$ is a vector over \mathbb{Z} , as defined in the beginning of Section 2).

Lemma 3.6. For any $k \geq 1$ let $\mathcal{I} = (R^\vee)^k$ and let $q \geq 1$ be arbitrary. Then for any $a \in \mathcal{I}$ of length less than $q/(2\hat{m}^{k-1}\sqrt{n})$, we have $\llbracket a \bmod q\mathcal{I} \rrbracket = a$. Moreover, if a is δ -subgaussian with parameter s , then for any $b \in (R^\vee)^\ell$ where $\ell \geq 0$, we have $\llbracket a \cdot b \bmod q(R^\vee)^{k+\ell} \rrbracket = a \cdot b$ except with probability at most

$$2n \exp(\delta - \pi q^2 / (2s \cdot \hat{m}^{k+\ell-1} \|b\|_2)^2).$$

4 Regularity

In this section we state a certain ‘‘regularity theorem’’ (whose proof appears in the full version) that is useful in cryptographic applications of ring-LWE, such as when adapting the ‘dual’ cryptosystem and IBEs of Gentry et al. [19] and others. Independently, a closely related statement (specialized to power-of-2 cyclotomics) was recently shown in [33] with a different proof.

The theorem says the following. Assume we are working with the m th cyclotomic of degree $n = \varphi(m)$, and let $q \geq 1$ be a prime integer. Let $a_1, \dots, a_{\ell-1}$ be chosen uniformly and independently from R_q . Then, with high probability over the choice of the a_i ’s, the distribution of $b_0 + \sum_{i=1}^{\ell-1} b_i a_i$ is within statistical distance $2^{-\Omega(n)}$ of uniform, where the b_i are chosen from a discrete Gaussian distribution on R of width essentially $nq^{1/\ell}$ (in the canonical embedding). Equivalently, the lemma says that if a_0 is any fixed invertible element of R_q and $a_1, \dots, a_{\ell-1}$ are uniformly and independently chosen from R_q , then $\sum_{i=0}^{\ell-1} b_i a_i$ is within $2^{-\Omega(n)}$ of uniform, where the b_i are chosen as before. The equivalence follows by simply dividing by a_0 . (The lemma we prove is actually more general, and applies to the joint distribution of $k \geq 1$ sums as above; see Theorem 4.1 and Corollary 4.2 for the exact statement.)

This regularity statement is already interesting and non-trivial when ℓ is as small as 2, and is close to being tight: for instance, in case m is a power of 2, a width of at least $\sqrt{n}q^{1/\ell}$ is required just for entropy reasons. To see this, recall that R is a rotation of $\sqrt{n}\mathbb{Z}^n$, so roughly speaking, a discrete Gaussian of width t covers $(t/\sqrt{n})^n$ points.

One might wonder about the significance of the b_0 term, and why we do not analyze the regularity of $\sum_{i=0}^{\ell-1} b_i a_i$ when all the a_i are chosen uniformly from R . In fact, a regularity lemma for exactly such sums was shown by Micciancio [27]. (His work is specialized to the ring $R = \mathbb{Z}[x]/\langle x^n - 1 \rangle$, but can be extended to other rings, as observed in [34].) Unfortunately, such sums have a much worse

regularity property, and in particular require super-constant ℓ to get negligible distance to uniformity. To see why this is the case, assume q is a prime satisfying $q = 1 \pmod m$, so that $\langle q \rangle$ splits completely into n ideals of norm q each. Letting \mathfrak{q} denote one of these prime factors, notice that with probability $q^{-\ell}$, all the a_i are in \mathfrak{q} . In this case, $\sum_{i=1}^m b_i a_i$ is in \mathfrak{q} with certainty, and its distribution is therefore very far from uniform. By adding the b_0 term we avoid this “common divisor” problem and get much better regularity, providing exponentially small distance to uniformity already for ℓ as small as 2.

The following is the regularity theorem. Here, for a matrix $A \in R_q^{[k] \times [\ell]}$ we define $\Lambda^\perp(A) = \{\vec{z} \in R^{[\ell]} : A\vec{z} = 0 \pmod{qR}\}$, which we identify with a lattice in H^ℓ . Its dual lattice (which is again a lattice in H^ℓ) is denoted by $\Lambda^\perp(A)^\vee$.

Theorem 4.1. *Let R be the ring of integers in the m th cyclotomic number field K of degree n , and $q \geq 2$ an integer. For positive integers $k \leq \ell \leq \text{poly}(n)$, let $A = [I_{[k]} \mid \bar{A}] \in (R_q)^{[k] \times [\ell]}$, where $I_{[k]} \in (R_q)^{[k] \times [k]}$ is the identity matrix and $\bar{A} \in (R_q)^{[k] \times [\ell-k]}$ is uniformly random. Then for all $r > 2n$,*

$$\mathbb{E}_{\bar{A}}[\rho_{1/r}(\Lambda^\perp(A)^\vee)] \leq 1 + 2(r/n)^{-n\ell} q^{kn+2} + 2^{-\Omega(n)}.$$

In particular, if $r > 2n \cdot q^{k/\ell+2/(n\ell)}$ then $\mathbb{E}_{\bar{A}}[\rho_{1/r}(\Lambda^\perp(A)^\vee)] \leq 1 + 2^{-\Omega(n)}$, and so by Markov’s inequality, $\eta_{2^{-\Omega(n)}}(\Lambda^\perp(A)) \leq r$ except with probability at most $2^{-\Omega(n)}$.

Using [31, Claim 3.8], we obtain the following corollary, which is often more useful in applications.

Corollary 4.2. *Let R , n , q , k , and ℓ be as in Theorem 4.1. Assume that $A = [I_{[k]} \mid \bar{A}] \in (R_q)^{[k] \times [\ell]}$ is chosen as in Theorem 4.1. Then, with probability $1 - 2^{-\Omega(n)}$ over the choice of \bar{A} , the distribution of $A\vec{x} \in R_q^{[k]}$ where each coordinate of $\vec{x} \in R_q^{[\ell]}$ is chosen from a discrete Gaussian distribution of radius $r > 2n \cdot q^{k/\ell+2/(n\ell)}$ over R , satisfies that the probability of each of the q^{nk} possible outcomes is in the interval $(1 \pm 2^{-\Omega(n)})q^{-nk}$ (and in particular is within statistical distance $2^{-\Omega(n)}$ of the uniform distribution over $R_q^{[k]}$).*

5 Example Cryptosystem

Here we give an example application of our toolkit which works in arbitrary cyclotomic rings. In particular, we give a public-key cryptosystem whose public key and ciphertext each consists of only two ring elements. In the full version, we also give a simple adaptation of the “dual-style” LWE-based public-key cryptosystem of [19], which uses our regularity theorem of Section 4, and which can serve as a foundation for (hierarchical) identity-based encryption. Additionally, in the full version we provide another (much more involved) example of a symmetric-key “somewhat homomorphic” cryptosystem and all the associated “modulus reduction” and “key switching” algorithms.

Let q be a positive integer that is coprime with every odd prime dividing m , and let p be a positive integer coprime with q . The message space is R_p . Let ψ be a continuous LWE error distribution over $K_{\mathbb{R}}$, and let $\lfloor \cdot \rfloor$ denote a valid discretization to (cosets of) R^{\vee} or pR^{\vee} . The cryptosystem is defined as follows.

- **Gen**: choose a uniformly random $a \leftarrow R_q$. Choose $s \leftarrow \lfloor \psi \rfloor_{R^{\vee}}$ and $e \leftarrow \lfloor p \cdot \psi \rfloor_{pR^{\vee}}$. Output $(a, b = \hat{m}(a \cdot s + e) \bmod qR) \in R_q \times R_q$ as the public key, and s as the secret key.
- **Enc** $_{(a,b)}(\mu \in R_p)$: choose $z \leftarrow \lfloor \psi \rfloor_{R^{\vee}}$, $e' \leftarrow \lfloor p \cdot \psi \rfloor_{pR^{\vee}}$, and $e'' \leftarrow \lfloor p \cdot \psi \rfloor_{t^{-1}\mu + pR^{\vee}}$. Let $u = \hat{m}(a \cdot z + e') \bmod qR$ and $v = b \cdot z + e'' \in R_q^{\vee}$. Output $(u, v) \in R_q \times R_q^{\vee}$.
- **Dec** $_s(u, v)$: compute $v - u \cdot s = \hat{m}(e \cdot z - e' \cdot s) + e'' \bmod qR^{\vee}$, and decode it to $d = \llbracket v - u \cdot s \rrbracket \in R^{\vee}$ (see Definition 3.5). Output $\mu = t \cdot d \bmod pR$.

Lemma 5.1. *The above cryptosystem is IND-CPA secure assuming the hardness of R -DLWE $_{q,\psi}$.*

Lemma 5.2. *Suppose that $\lfloor \psi \rfloor_{c+R^{\vee}}$ is δ -subgaussian with parameter $r \geq 1$ and $\delta = O(1)$, for any coset $c + R^{\vee}$. Then assuming $q > \hat{m}pr^2 \cdot \omega(\sqrt{n} \log n)$, the decryption procedure is correct with probability negligibly close to one (over all the random choices of **Gen** and **Enc**).*

References

- [1] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, pages 553–572, 2010.
- [2] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. *Theory of Computing Systems*, 48(3):535–553, April 2011. Preliminary version in STACS 2009.
- [3] 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.
- [4] László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.
- [5] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, pages 719–737, 2012.
- [6] Wieb Bosma. Canonical bases for cyclotomic fields. *Appl. Algebra Eng. Commun. Comput.*, 1:125–134, 1990.
- [7] Xavier Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Public Key Cryptography*, pages 499–517, 2010.
- [8] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *ICTS*, pages 309–325, 2012.
- [9] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *CRYPTO*, pages 505–524, 2011.
- [10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *J. Cryptology*, 25(4):601–639, 2012. Preliminary version in Eurocrypt 2010.

- [11] Léo Ducas and Alain Durmus. Ring-LWE in polynomial rings. In *Public Key Cryptography*, pages 34–51, 2012.
- [12] Paul Erdős. On the coefficients of the cyclotomic polynomial. *Bulletin of the American Mathematical Society*, 52(2):179–184, 1946.
- [13] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. crypto.stanford.edu/craig.
- [14] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [15] Craig Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *CRYPTO*, pages 116–137, 2010.
- [16] Craig Gentry, Shai Halevi, Chris Peikert, and Nigel P. Smart. Ring switching in BGV-style homomorphic encryption. In *SCN*, pages 19–37, 2012. Full version at <http://eprint.iacr.org/2012/240>.
- [17] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully homomorphic encryption with polylog overhead. In *EUROCRYPT*, pages 465–482, 2012.
- [18] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the AES circuit. In *CRYPTO*, pages 850–867, 2012.
- [19] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [20] Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *CHES*, pages 530–547, 2012.
- [21] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
- [22] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
- [23] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155, 2006.
- [24] Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, pages 37–54, 2008.
- [25] Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72, 2008.
- [26] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, pages 1–23, 2010.
- [27] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.
- [28] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.
- [29] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342, 2009.
- [30] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487, 2007.
- [31] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.
- [32] N.P. Smart and F. Vercauteren. Fully homomorphic SIMD operations. Cryptology ePrint Archive, Report 2011/133, 2011. <http://eprint.iacr.org/>.
- [33] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, pages 27–47, 2011.
- [34] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635, 2009.