

Order-LWE and the Hardness of Ring-LWE with Entropic Secrets

Madalina Bolboceanu ^{*} Zvika Brakerski [†] Renen Perlman [†] Devika Sharma [†]

Abstract

We propose a generalization of the celebrated Ring Learning with Errors (RLWE) problem (Lyubashevsky, Peikert and Regev, Eurocrypt 2010, Eurocrypt 2013), wherein the ambient ring is not the ring of integers of a number field, but rather an *order* (a full rank subring). We show that our Order-LWE problem enjoys worst-case hardness with respect to short-vector problems in invertible-ideal lattices *of the order*.

The definition allows us to provide a new analysis for the hardness of the abundantly used Polynomial-LWE (PLWE) problem (Stehlé et al., Asiacrypt 2009), different from the one recently proposed by Rosca, Stehlé and Wallet (Eurocrypt 2018). This suggests that Order-LWE may be used to analyze and possibly *design* useful relaxations of RLWE.

We show that Order-LWE can naturally be harnessed to prove security for RLWE instances where the “RLWE secret” (which often corresponds to the secret-key of a cryptosystem) is not sampled uniformly as required for RLWE hardness. We start by showing worst-case hardness even if the secret is sampled from a subring of the sample space. Then, we study the case where the secret is sampled from an *ideal* of the sample space or a coset thereof (equivalently, some of its CRT coordinates are fixed or leaked). In the latter, we show an interesting threshold phenomenon where the amount of RLWE *noise* determines whether the problem is tractable.

Lastly, we address the long standing question of whether high-entropy secret is sufficient for RLWE to be intractable. Our result on sampling from ideals shows that simply requiring high entropy is insufficient. We therefore propose a broad class of distributions where we conjecture that hardness should hold, and provide evidence via reduction to a concrete lattice problem.

1 Introduction

The Learning with Errors (LWE) problem, as introduced by Regev [Reg05], provides a convenient way to construct cryptographic primitives whose security is based on the hardness of *lattice problems*. The assumption that LWE is intractable was used as a basis for various cryptographic designs, including some cutting edge primitives such as fully homomorphic encryption (FHE) [BV11b], and attribute based encryption (ABE) for general policies [GVW13, BGG⁺14]. Two of the most appealing properties of the LWE problem are the existence of a reduction from *worst-case* lattice problems [Reg05, Pei09, BLP⁺13, PRSD17] (which is most relevant to this work), and its conjectured post-quantum security.

^{*}Bitdefender, Romania, mbolboceanu@bitdefender.com. A large portion of this study was conducted while visiting the Weizmann Institute of Science, Israel, supported by the European Union Horizon 2020 Research and Innovation Program via Project PROMETHEUS (Grant 780701).

[†]Weizmann Institute of Science, Israel, {zvika.brakerski,renen.perlman,devika.sharma}@weizmann.ac.il. Supported by the Israel Science Foundation (Grant No. 468/14), Binational Science Foundation (Grants No. 2016726, 2014276), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701).

On the other hand, one of the shortcomings of the LWE assumption is the relatively high computational complexity and large instance size (as a function of the security parameter) that it induces. This results, for example, in LWE-based encryption schemes having long keys and ciphertexts, and also high encryption complexity. It was known since the introduction of the NTRU cryptosystem [HPS98] and more rigorously in [LM06, PR06] that these aspects can be significantly improved by relying on lattices that stem from algebraic number theory.¹ In [LPR10, LPR13], Lyubashevsky, Peikert and Regev defined an algebraic number theoretic analog of the LWE problem, called Ring-LWE (RLWE). Similar to Regev’s original result, they showed that RLWE is as hard as solving worst-case *ideal lattice* problems.

Ring-LWE and its extensions quickly became a useful resource for the construction of various cryptographic primitives [BF11, BV11a, BGV12, GHS12, DDL13, AP13, HS14, BKLP15, ADPS16, BVWW16] (an extremely non-exhaustive list of examples). RLWE is appealing due to its improved efficiency, and its provable security guarantee based on the hardness of worst case (ideal) lattice problems. However, in concrete instantiations, parameters are not set based on provable hardness guarantees, but rather on the minimal parameters that prevent known and conceivable attacks, in order to achieve the best possible efficiency. In the case of RLWE, parameters are set way beyond the regime where we have provable guarantee in terms of choice of security parameter and, most relevant to this work, in terms of sampling secrets from different distributions than those for which provable security applies.² While a gap between the provable and concrete security properties of a cryptosystem is expected, one would at least like to make sure that changing the distribution does not make the problem qualitatively easy. In other words, we would like to show that the problem remains at least asymptotically hard with the new distributions.

Over the years, it has been shown that the LWE problem is quite robust to changes in the prescribed distribution, thus providing desired evidence for the safety of using the assumption in various settings. More precisely, it was shown that LWE hardness holds even if the secret (a vector that, very roughly, represents the coordinates of a hidden lattice point) is not sampled uniformly, as tradition, but is rather leaked [AGV09, DGK⁺10] or is chosen from a binary distribution of sufficient entropy [GKPV10, BLP⁺13] (with obvious loss coming from the secret having lower entropy). It is almost trivial to verify that if the LWE secret is chosen uniformly from a linear subspace of its prescribed space, then security degrades gracefully with the dimension of the space of secrets. (We note that there has also been much work on modifying the *noise distribution* of LWE, e.g. [BPR12, MP13]. However, the focus of this work is the distribution of secrets.)

Much less is known for RLWE. This is because its algebraic structure (which is the very reason for efficiency gains) prevents the use of techniques like randomness extraction that are instrumental to the aforementioned LWE robustness results.

This Work. Motivated by the task to investigate the behavior of the RLWE problem on non-uniform secret distributions, we present new tools to prove security in some cases and insecurity in others. The main tool that we introduce is a generalization of the RLWE problem that we call Order-LWE, and prove a worst-case hardness result for this problem.³ We show that the Order-LWE abstraction naturally implies a new proof for a previous result in the literature [RSW18] with

¹This inefficiency is common to cryptographic constructions based on “generic” lattices. Indeed, NTRU was introduced before the LWE assumption was formulated.

²Sometimes this is done not for efficiency but for functionality purposes, e.g. [BVWW16].

³As a reader with background in algebraic number theory would speculate, this is a setting where RLWE is instantiated relative to *orders* in a number field, rather than its ring of integers.

new and comparable parameters.

We justify that the formulation of Order-LWE is quite useful in exploring variants of RLWE where the distribution of secrets has some algebraic structure (a special case of secrets from a subfield was studied in [GHPS13]). We prove Order-LWE hardness (and thus worst-case hardness) when the secret is sampled from any subring of the prescribed space. We use this approach to address the fundamental question of whether any distribution of secrets with sufficiently high entropy implies RLWE hardness. We show that in some settings, RLWE with uniform secrets is intractable (under conservative worst-case ideal lattice assumptions), but a slight decrease in entropy leads to a complete break. This is the case when the distribution of secrets is supported over an *ideal* (or a coset thereof). On the other hand, we show that increasing the *noise* in the RLWE instance can compensate for the deficiency in secret entropy in this setting.

Finally, we address the more ambitious goal of proving security for secrets with no algebraic structure. Since, as we mentioned above, high entropy is insufficient as condition by itself for security, we identify a family of high-entropy distributions that capture (at least approximately) many of the relaxed variants of RLWE. We show that a particular (average case) hardness assumption implies hardness for this class of distributions.

We provide an overview of our results and techniques next.

1.1 Background

Recall that in the LWE problem, a secret vector \mathbf{s} is sampled from \mathbb{Z}_q^n for some modulus q ; an adversary gets oracle access to samples of the form $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle + e \pmod{q})$ where each $\mathbf{a} \in \mathbb{Z}_q^n$ is uniform and e is a small integer, say sampled from a discrete Gaussian with parameter $\ll q$. The adversary's goal is to distinguish this oracle from the one where $b \in \mathbb{Z}_q$ is random.

In the RLWE problem, the sample spaces are also vector spaces over \mathbb{Z}_q but with a ring structure. In this high level overview, for the sake of simplicity of notation and algebraic structure, we restrict to the case where the ring is the ring of integers in the power-of-two cyclotomic field $\mathbb{Q}[x]/(x^n + 1)$. An interested reader may see Section 2 for precise definitions in the general case. The cyclotomics is a particularly simple case: the so called *ring of integers* in this case is the ring of polynomials $R = \mathbb{Z}[x]/(x^n + 1)$. In this setting, the RLWE problem with modulus q is as follows: sample a random secret $s \in R_q = R/qR$, and provide the adversary with oracle access to samples of the form $(a, as + e) \in R_q^2$, where a is uniform and e is sampled from some “small” noise distribution (for our purposes, think of e as polynomial with Gaussian coefficients $\ll q$). The arithmetics is over R_q , and the goal is to distinguish these samples from uniform R_q^2 samples.⁴ For this overview, we will assume for simplicity that q is a prime, and focus on the setting (which is most commonly used in cryptography) where q splits completely as an ideal in R into a product of n distinct prime ideals. (In the case of the cyclotomics, this condition amounts to $q \equiv 1 \pmod{2n}$.) By the Chinese Remainder Theorem, the quotient $R_q := R/qR \simeq \mathbb{Z}[x]/(x^n + 1, q)$ is isomorphic to \mathbb{Z}_q^n and hence an element in R_q can be represented as a vector of n elements in \mathbb{Z}_q with pointwise addition and multiplication. This is called the CRT representation of elements in R_q and for $c \in R_q$, we denote its CRT coordinates by $c[1], \dots, c[n]$.

⁴An informed reader may notice that in the formal RLWE definition, s needs to be sampled from the dual of R_q , and e needs to be small in the so called “canonical embedding”. However, in the cyclotomic setting these distinction makes little difference and our choice makes the presentation simpler. Another simplifying choice for the exposition is to only consider discrete noise distributions.

1.2 Our Results

An overview of our results follows.

The Order-LWE Problem. We formulate a version of R -LWE where R is replaced by an order \mathcal{O} in K . An order in a number field is a subring of R which has full rank (i.e. can be described as a \mathbb{Z} -span of exactly n elements).⁵ We furthermore generalize the modulus of the RLWE equations, and instead of taking the equations modulo (the ideal generated by) the integer q , we allow to mod out by any ideal in the order. We call this problem Order-LWE and denote it as \mathcal{O} -LWE. More precisely, we define two variants of the problem \mathcal{O} -LWE and \mathcal{O}^\vee -LWE which have a duality relation between them (and which are equivalent to each other when $\mathcal{O} = R$). As explained above, R -LWE is a special case of \mathcal{O} -LWE (and, in a different notation, of \mathcal{O}^\vee -LWE).

Recalling that Ring-LWE was shown to be as hard to solve as worst-case lattice problems over ideal lattices from the ring R , we propose an analogous claim for Order-LWE. Using similar techniques as those used for proving Ring-LWE hardness, but with some necessary adaptations, we show that solving \mathcal{O} -LWE is at least as hard as solving short-vector problems on a class of lattices that is defined by the set of invertible ideals in the order \mathcal{O} . This result generalizes the known result on R -LWE (note that in R all ideals are invertible). For \mathcal{O}^\vee -LWE, worst-case hardness follows for lattices whose *dual* is an invertible \mathcal{O} -ideal (again, in R this holds for all ideals). We mention that these sets of lattices coincide in the case when the dual of the order \mathcal{O} is an invertible \mathcal{O} -ideal.

We show that using a larger order makes the \mathcal{O}^\vee -LWE problem harder, and in that sense R -LWE is harder than any other \mathcal{O}^\vee -LWE problem. This is the case even though formally the set of duals of (invertible) \mathcal{O} -ideals is disjoint from the set of R -ideals. We believe that this is due to the fact that any \mathcal{O} -ideal lattice can be (efficiently) mapped to an R ideal that contains it as a sublattice.

See Section 3 for a formal and general definition of \mathcal{O} -LWE, its dual \mathcal{O}^\vee -LWE and the respective worst-case hardness results.

A Corollary: New Hardness for Polynomial-LWE. Our definition of Order-LWE gives insight on the hardness of other computational problems underlying cryptographic constructions; specifically, the Polynomial-LWE problem (PLWE) [SSTX09, BV11a]. In PLWE, s and a are simply random polynomials with integer coefficients modulo a polynomial f and an integer q , and the noise e is a polynomial with small coefficients. It is evident that the PLWE problem provides the simplest interface for LWE over polynomial rings. In many useful cases, for example the power-of-two cyclotomic case, it is straightforward to relate PLWE and RLWE. However, for general polynomials f the connection is far from immediate, since the ring of integers of an arbitrary number field does not look like $\mathbb{Z}[x]/(f)$. Recently, Rosca, Stehlé and Wallet [RSW18] showed a reduction relating the hardness of PLWE in the general case from RLWE and thus from worst-case lattice problems.

We observe that we can straightforwardly address this problem using our Order-LWE machinery. The ambient space for the PLWE problem is the ring $\mathbb{Z}[x]/(f)$, for a polynomial $f \in \mathbb{Z}[x]$. This ring is a subring of full rank of the ring of integers of $K := \mathbb{Q}[x]/(f)$, and hence indeed an order. Therefore translation between PLWE and \mathcal{O} -LWE has two aspects: “reshaping” the noise distribution (identically to [RSW18]), and syntactic mapping of the secret to the dual domain. The [RSW18] reduction requires a few additional steps and in this sense our reduction is more direct.

⁵The full-rank condition arises naturally in applications as we discuss below.

Pinpointing the exact relation of our reduction to the one of [RSW18] is not straightforward. The class of lattices for which we show worst-case hardness is different (and in fact formally disjoint) from the class of lattices in [RSW18]. This is because the hardness result of Order-LWE deals with the worst case lattice problems on invertible \mathcal{O} -ideals. However, any \mathcal{O} -ideal can be translated into R -ideal which contains it as a sublattice. It therefore appears that R -lattices as in [RSW18] may provide stronger evidence of intractability. On the other hand, the approximation factor achieved by our reduction is never larger and in many cases should be much smaller than that achieved by [RSW18], depending on the specific number field.

This result suggests that perhaps it is instructive to think about orders where objects can be represented and operated on efficiently (more efficiently than over R), and in those orders \mathcal{O} -LWE could be a simple way to argue about the security of a cryptosystem with simpler interface than RLWE. We did not explore this avenue further. See Section 4 for the full details of our PLWE proof and comparison with [RSW18].

Ring-LWE with Secrets From a Subring/Order. We consider the hardness of the Ring-LWE problem, in the setting where the secret s is sampled from some subset with algebraic structure. As we described above, the proper distribution of secrets is uniform over the ring R_q (R modulo q). In this paper we consider a subring of this ring, but we note that this subring must still contain qR , since Ring-LWE equations are taken modulo q . This naturally imposes full-rank condition on the subring and thus orders naturally arise again. Indeed, we consider distributions that are uniform over $\mathcal{O}_q = \mathcal{O}/q\mathcal{O}$ for an order \mathcal{O} .

To motivate the setting of sampling the secret from a subring and illustrate its importance, we start with an analogy with (standard) LWE. In the LWE context, if the secret is sampled from a k -dimensional linear subspace of \mathbb{Z}_q^n , the problem easily translates to an LWE instance where n is replaced by k . In the ring setting, the rich algebraic structure makes the task of defining and analyzing such straightforward transformations much more involved. Previous works [BGV12, GHPS13, AP13] considered the notion of *ring-switching* which implies the hardness of RLWE when the secret is sampled from the ring of integers of a *subfield* of the field K . However, such transformations do not apply when K has no subfields of dimension k or no proper subfields at all. Our proposed setting allows to sample s from a subring of R_q that is isomorphic to \mathbb{Z}_q^k , for $1 \leq k \leq n$ and thus provides an algebraic analog of the linear subspace property.

One can view the subring property also in terms of the CRT coordinates (when q splits over R). A subring of R_q that is isomorphic to \mathbb{Z}_q^k is in one-to-one correspondence with an onto mapping $\alpha : [n] \rightarrow [k]$ as follows; sample k elements r_1, \dots, r_k from \mathbb{Z}_q uniformly, and set the CRT coordinates of an element $s \in R_q$ as $s[j] = r_{\alpha(j)}$, for $j \in [n]$. One can verify that this set forms a subring.

In Section 5, we show that the RLWE problem with secret sampled from an order \mathcal{O} is harder than the \mathcal{O}^\vee -LWE problem, which in turn is harder than the worst case problems on the duals of (invertible) \mathcal{O} -ideal lattices. In fact, given two orders $\mathcal{O}' \subseteq \mathcal{O}$, we show that \mathcal{O}^\vee -LWE is at least as hard as \mathcal{O}'^\vee -LWE (albeit with increase in noise which is comparable to the norm of a minimal generating set of \mathcal{O}^\vee over \mathcal{O}'^\vee). Since $\mathcal{O} \subseteq R$, this shows that R -LWE is harder than \mathcal{O}^\vee -LWE with appropriate noise increase. This result is in a similar flavor to the one of [GHPS13] which shows that R -LWE in a field K is harder than R' -LWE in a subfield $K' \subseteq K$. Since R' , the ring of integers of K' , is contained in R as a subring, our result implies hardness in this setting as well.

Ring-LWE with Secrets From Ideals, and High-Entropy Secrets. As we already mentioned, understanding the behavior of Ring-LWE in the setting where the secret is sampled from an arbitrary high-entropy distribution is a central subject of inquiry in the area. We show that if we sample s from a “dense” ideal (or equivalently zero-out a few of its CRT coordinates), then we may end up with a distribution that is high-entropy on one hand but makes Ring-LWE insecure on the other. More concretely, consider RLWE samples where the secret s is sampled from R_q such that its j -th CRT coordinate, $s[j]$, is uniformly chosen from \mathbb{Z}_q , for all $j \in T$, a randomly chosen subset of $[n]$, and $s[j] = 0$, for $j \notin T$. This is equivalent to choosing s uniformly from $\mathcal{P}_q := \mathcal{P}/qR$, where \mathcal{P} is the ideal of R that contains qR .⁶ If $|T| = k$ and we denote $\epsilon = \frac{k}{n}$, then the distribution of secret will have min-entropy $(1 - \epsilon)n \log q$. A good running example is ϵ which is a small constant (e.g. $\epsilon = 0.1$).

One can view this as a more structured analog of an LWE instance with a composite modulus $q = p_1 p_2$, where the secret s is a multiple of p_1 . It is straightforward to see that, in such a LWE instance, if the magnitude of the error e is sufficiently smaller than p_1 , this instance can be easily solved by dividing b by p_1 and rounding to the nearest integer, thereby yielding a noiseless set of equations modulo p_2 . However, if the noise magnitude is sufficiently larger than p_1 , then the instance is secure (intuitively, since one can essentially view it as scaling up of a mod p_2 LWE instance).

Things are more involved in the ring setting as we cannot just round to the nearest integer. Instead, we can interpret the factors of q as lattices. If the lattice corresponding to \mathcal{P} has a good decoding basis, it means that we can recover e when small enough. We also provide a ring analog of the complementary result by showing that if the noise is sufficiently large, then RLWE hardness holds. The latter is done by viewing the instance, again, as a scaled up RLWE instance modulo \mathcal{Q} , only now \mathcal{Q} is not an integer but an ideal. Our \mathcal{O} -LWE generalization of RLWE allows us to derive RLWE hardness in this setting.

Thus, we show that high entropy of secrets alone is insufficient to argue RLWE security and demonstrate an interplay between the entropy of the secret and the amplitude of noise. Interestingly, we exhibit a threshold phenomenon where the RLWE instance with secret sampled from an ideal is insecure if the error is modestly below the threshold of roughly $q^{-(1-\epsilon)}$, and secure if it is modestly above this value. The “modest” factors depend on the number field, but correspond to a fixed polynomial in the degree n in the cyclotomic case. The formal and general analysis of these results appears in Section 6.

Ring-LWE with Secrets From a k -Wise Independent Distribution. Given that a general result for high-entropy distributions cannot be achieved, we consider in the final section of the paper a subclass of high-entropy distributions. These distributions do not adhere to uniform sampling from an algebraic structure but instead have the following property; the marginal distribution over any subset of k CRT coordinates is jointly (statistically close to) uniform.⁷ In terms of entropy, such distributions must have min-entropy at least $k \log q$, and this entropy is also spread evenly across all k -tuples of CRT coordinates.

We speculate that the k -wise independence condition is sufficient for obtaining RLWE hardness. However, we are unable to show this via worst-case hardness. Instead, we define an average case

⁶As we hinted above, s is actually an element of the dual of R which is not a ring and doesn’t have ideals, however there is a natural translation between the dual and primal domain that captures the CRT/ideal structure. See Section 6 for the formal treatment.

⁷A computational variant is also possible, but needs to carefully define the indistinguishability experiment.

problem, which we call Decisional Bounded Distance Decoding on a Hidden Lattice (HLBDD) and show that the RLWE problem with secret sampled from a k -wise distribution is at least as hard as this problem. In HLBDD, the adversary needs to distinguish between a random oracle on R_q and an oracle of the following form. Upon initialization of the oracle, a set $T \subseteq [n]$ of cardinality k is sampled. For every oracle call, the oracle generates elements v, e as described next, and returns $v + e \pmod{q}$. For the element v , the CRT coordinate $v[j]$ is random if $j \in T$, and 0 otherwise. The element e is a small noise element, say Gaussian. This can be viewed as the decisional version of the bounded distance decoding (BDD) problem on the ideal lattice $\mathcal{I} := \prod_{j \in T} \mathfrak{p}_j$ (where $\{\mathfrak{p}_i\}_i$ are the prime factors of the ideal qR), since the element v is sampled from \mathcal{I} . We stress that this is the hidden version as T is sampled randomly at the invocation of the oracle, causing \mathcal{I} to be hidden. As in the standard BDD problem, we can consider HLBDD with worst-case noise and also with arbitrary noise distributions. Given the current understanding of the hardness of lattice problems, discrete Gaussian noise seems like a natural choice.

The HLBDD assumption is similar to one made in [HPS⁺14]. However, they only require $k = n/2$, whereas we attempt to take k to be very small, e.g. $k = n^{0.1}$. We assert that the hardness of the problem relies crucially on the set T being chosen at random in the beginning of the experiment rather than being fixed throughout. In other words, we cannot allow preprocessing that depends on T . This is because computing a good basis for the ideal lattice \mathcal{I} , defined by T , makes the HLBDD problem easy. It is also important to mention that T itself does not need to be known to the adversary; in this sense HLBDD resembles the approximate GCD problem [DGHV10]. Lastly, we note that it is sufficient for our purposes to limit the adversary to only make 2 oracle calls. Namely, the problem is to distinguish two samples $(v_1 + e_1, v_2 + e_2)$ from two uniform elements in R_q . Despite our efforts, we were unable to find additional corroboration to the hardness of this problem and we leave it as an interesting open problem to characterize its hardness.

Let us try to motivate and justify our assertion that the class of k -wise independent distributions is meaningful. Indeed, this class captures the spirit of some of the heuristic entropic distributions that were considered for RLWE. For example, consider the representation of the secret s as a formal polynomial modulo q (recall that $R_q \simeq \mathbb{Z}[x]/(f, q)$ is a ring of polynomials). If each coefficient of s is sampled from a Gaussian so that the total distribution has sufficient entropy (slightly above the necessary $k \log q$), then this distribution will be k -wise independent (as follows from a standard “smoothing” argument). This shows that sampling secrets with very low norm does not violate security under our new assumption. While it was previously known that sampling the secret from the noise distribution keeps security intact (also known as RLWE in Hermite Normal Form [ACPS09]), we are not aware of a proof of security when s is chosen from a narrower distribution than the error. This can be seen as a step in the direction of matching the robustness of LWE results [GKPV10, BLP⁺13], that show that LWE remains hard even with high entropy binary secrets. We note that low norm secrets are of importance in the FHE literature (e.g. [BGV12, HS14, HS15]). In fact, in the HElib implementation [HS14, HS15] the secret is chosen to be a random extremely sparse polynomial. Heuristically, it seems plausible that random sparse polynomials should translate into k -wise independent distributions but we do not have a proof for this speculation as yet.⁸

Another example of an interesting k -wise independent distribution is the “entropic RLWE” formulation that came up in the obfuscation literature [BVWW16]. That setting consists of a

⁸Intuitively, this follows from the fact that the translation between the coefficient and CRT representation is a linear transformation defined by a Vandermonde matrix. In order to prove k -wise independence we need to show that any subset of k rows of the Vandermonde matrix constitutes a deterministic extractor from a uniform distribution over a Hamming ball. Analogous theorems exist in other contexts, but we do not have a proof as of yet.

large number of public elements s_1, \dots, s_m , sampled from the noise distribution (which is Gaussian in the polynomial coefficient representation and thus can be shown to be k -wise independent in the CRT representation). The secret is generated by sampling a binary vector $\vec{z} = (z_1, \dots, z_m)$ and outputting $s = \prod s_i^{z_i}$. Using the leftover hash lemma, one can show that so long as \vec{z} has entropy sufficiently larger than $k \log q$, the resulting distribution will be k -wise independent as well. It is worth noting that in order to achieve the strongest notion of security for their obfuscator, [BVWW16] use \vec{z} with entropy $\ll \log q$ to which our technique does not directly apply.

See Section 7 for more details on this result.

1.3 Paper Organization

Section 2 contains preliminaries and definitions. The Order-LWE problem is formally defined in Section 3, where the worst case hardness reduction is proved as well. The new hardness result for PLWE appears in Section 4. We then present our results on sampling secrets from subrings in Section 5, on sampling secrets from ideals in Section 6, and finally on sampling secrets from k -wise independent distributions in Section 7.

2 Preliminaries

For a vector \mathbf{x} in \mathbb{C}^n and $p \in [1, +\infty)$, we mean by ℓ_p norm $\|\mathbf{x}\|_p = (\sum_i |x_i|^p)^{1/p}$ and by ℓ_∞ norm $\|\mathbf{x}\|_\infty = \max_i |x_i|$. We refer to ℓ_2 norm if p is omitted. Let \mathcal{D} be a distribution. When writing $x \leftarrow \mathcal{D}$ we mean sampling an element x according to the distribution \mathcal{D} . Similarly, for a finite set Ω , we denote by $x \xleftarrow{\$} \Omega$ sampling an element x from Ω uniformly at random. For $\varepsilon > 0$, we say that \mathcal{D} is a B -bounded distribution if $\Pr_{x \leftarrow \mathcal{D}}[\|x\| > B]$ is negligible. For two distributions \mathcal{D}_1 and \mathcal{D}_2 over the same measurable set Ω , we consider their *statistical distance* as $\Delta(\mathcal{D}_1, \mathcal{D}_2) = \frac{1}{2} \int_\Omega |\mathcal{D}_1(x) - \mathcal{D}_2(x)| dx$. When the support of \mathcal{D} is a finite set Ω , we define the *entropy* of \mathcal{D} to be $H(\mathcal{D}) := \sum_{\omega \in \Omega} \mathcal{D}(\omega) \cdot \log_2(1/\mathcal{D}(\omega))$. In a similar way, its *min-entropy* is defined by $H_\infty(\mathcal{D}) := \min_{\omega \in \Omega} \log_2(1/\mathcal{D}(\omega))$. It is easy to verify that $H(\mathcal{D}), H_\infty(\mathcal{D}) \leq \log_2 |\Omega|$ with equality if and only if \mathcal{D} is the uniform distribution over Ω . When discussing computational problems, we consider by default the standard (nonuniform) polynomial time adversarial model.

2.1 The Space H

When working with number fields from a geometric perspective, we usually work with the following space $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ for some numbers $s_1 + 2s_2 = n$, defined as

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

Note that H , equipped with the inner product induced 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 coordinate, and 0 elsewhere; then for $j \in [s_1]$, we take $\mathbf{h}_j = \mathbf{e}_j \in \mathbb{C}^n$, and for $s_1 < j \leq s_1 + s_2$ we take $\mathbf{h}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_j + \mathbf{e}_{j+s_2})$ and $\mathbf{h}_{j+s_2} = \frac{1}{\sqrt{-2}}(\mathbf{e}_j - \mathbf{e}_{j+s_2})$.

We will also equip H with the ℓ_p norm induced on it from \mathbb{C}^n .

2.2 Lattices

We define a *lattice* as a discrete additive subgroup of H . Equivalently, a lattice is the \mathbb{Z} -span of some set of k linearly independent *basis* vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\} \subseteq H$:

$$\mathcal{L} = \left\{ \sum z_i \mathbf{b}_i \mid z_i \in \mathbb{Z} \right\}.$$

We refer to k as the *rank* of the lattice, and to n as its *dimension*. If $k = n$, we say that the lattice is *full-rank*. The *determinant* of a lattice \mathcal{L} , denoted by $\det(\mathcal{L})$, is defined as $\sqrt{\det(\langle \mathbf{b}_i, \mathbf{b}_j \rangle)_{i,j}}$. It is a standard fact that this does not depend on the choice of the basis.

The *minimum distance* $\lambda_1^p(\mathcal{L})$ of a lattice \mathcal{L} with respect to the ℓ_p norm, $p \in [1, \infty]$, is the length of a shortest nonzero lattice vector. More generally, we define the *i th successive minimum* as

$$\lambda_i^p(\mathcal{L}) := \inf\{r > 0 \mid \dim(\text{span}(\mathcal{L} \cap \overline{\mathbf{B}}^p(0, r))) \geq i\},$$

where $\overline{\mathbf{B}}^p(0, r)$ is the closed ball of radius r around 0 in norm ℓ_p . When the ℓ_2 norm is used, we denote the i th successive minimum by λ_i .

The *basis minimum* of a lattice \mathcal{L} , denoted by $g^p(\mathcal{L})$, is the smallest radius r such that $\overline{\mathbf{B}}^p(0, r)$ contains a basis of the lattice. The *dual* of a lattice $\mathcal{L} \subset H$ is defined as $\mathcal{L}^* = \{\mathbf{x} \in H \mid \langle \mathcal{L}, \mathbf{x} \rangle \subseteq \mathbb{Z}\}$, which is also a lattice of same rank as \mathcal{L} . We recall the following inequality relating $\lambda_1(\mathcal{L})$ and $\lambda_n(\mathcal{L}^*)$.

Lemma 2.1. *For any lattice \mathcal{L} of dimension n , $\lambda_1(\mathcal{L}) \cdot \lambda_n(\mathcal{L}^*) \geq 1$.*

2.2.1 Gaussians

For $r > 0$, define the Gaussian function $\rho_r : H \rightarrow (0, 1]$ as $\rho_r(\mathbf{x}) := \exp(-\pi \|\mathbf{x}\|^2 / r^2)$. By normalizing this function, we obtain the *continuous Gaussian probability distribution* of width r , denoted by D_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. Define $G = \{\mathbf{r} \in (\mathbb{R}^+)^n \mid r_{s_1+s_2+i} = r_{s_1+i}, \forall i \in [s_2]\}$; note this has symmetry mirroring that of H . For consistency with prior works, we sometimes use $r \in \mathbb{R}^+$ as shorthand for the all- r s vector $r\mathbf{1} \in G$. For $\mathbf{r} \in G$, a sample from $D_{\mathbf{r}}$ is given by $\sum x_i \mathbf{h}_i$, where each x_i is chosen independently from the (one-dimensional) Gaussian distribution D_{r_i} over \mathbb{R} . We equip partial ordering on G defined by $\mathbf{r}' \geq \mathbf{r}$ if $r'_i \geq r_i$ for all i . By total Gaussian measure on a set X we mean $\rho_{\mathbf{r}}(X) := \sum_{\mathbf{x} \in X} \rho_{\mathbf{r}}(\mathbf{x})$.

Micciancio and Regev [MR07] introduced a lattice quantity called the *smoothing parameter*, and related to it various lattice quantities.

Definition 2.2 (Smoothing Condition). *For a lattice $\mathcal{L} \subset H$, positive real $\varepsilon > 0$ and $\mathbf{r} \in G$, we write $\mathbf{r} \geq \eta_\varepsilon(\mathcal{L})$ if $\rho_{1/\mathbf{r}}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon$, where $1/\mathbf{r} = (1/r_1, \dots, 1/r_n)$.*

A first application of the smoothing parameter suggests that the total Gaussian measure on any translate of the lattice is almost the same, as long as the Gaussian parameter \mathbf{r} exceeds the smoothing parameter.

Lemma 2.3 (Special case of [Pei10, Lemma 2.4]). *For any lattice \mathcal{L} of dimension n , $\varepsilon \in (0, 1)$, $\mathbf{c} \in H$ and $\mathbf{r} \geq \eta_\varepsilon(\mathcal{L})$ we have $\rho_{\mathbf{r}}(\mathcal{L} + \mathbf{c}) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \rho_{\mathbf{r}}(\mathcal{L})$.*

For a lattice \mathcal{L} in H and $\mathbf{r} \in G$ we define the *discrete Gaussian distribution* over \mathcal{L} as being the distribution having as support \mathcal{L} and as mass function $D_{\mathcal{L},\mathbf{r}}(\mathbf{x}) := \rho_{\mathbf{r}}(\mathbf{x})/\rho_{\mathbf{r}}(\mathcal{L})$, for any $\mathbf{x} \in \mathcal{L}$.

The following lemma justifies the name “smoothing parameter” adapted to discrete Gaussians. This is a slightly generalized result of [GPV08, Corollary 2.8] in the case of elliptical Gaussians, whose proof is the same but uses Lemma 2.3 instead of [GPV08, Lemma 2.7].

Lemma 2.4. *For any two full rank lattices $\mathcal{L}' \subseteq \mathcal{L}$ in H , $\varepsilon \in (0, 1/2)$ and $\mathbf{r} \geq \eta_{\varepsilon}(\mathcal{L}')$ in G , the statistical distance between $D_{\mathcal{L},\mathbf{r}}$ mod \mathcal{L}' and the uniform distribution over \mathcal{L}/\mathcal{L}' is at most 2ε .*

Theorem 2.5 (Special case of [Pei10, Theorem 3.1]). *Let $\mathcal{L}_1 \subset H$ be a lattice, and let $\mathbf{r}_1, \mathbf{r}_2 \in G$. Define $\mathbf{r} \in G$ by $\mathbf{r}_i^2 = (\mathbf{r}_1)_i^2 + (\mathbf{r}_2)_i^2$, and $\mathbf{r}_3 \in G$ by $1/(\mathbf{r}_3)_i := 1/(\mathbf{r}_1)_i + 1/(\mathbf{r}_2)_i$. Assume that $\mathbf{r}_1 \geq \eta_{\varepsilon}(\mathcal{L}_1)$ for some positive $\varepsilon \leq 1/2$, and let $\mathbf{c}_1, \mathbf{c}_2 \in H$ be arbitrary. Consider the following probabilistic experiment:*

Choose $\mathbf{x}_2 \leftarrow D_{\mathbf{r}_2}$, then choose $\mathbf{x}_1 \leftarrow \mathbf{x}_2 + D_{\mathcal{L}_1 + \mathbf{c}_1 - \mathbf{x}_2, \mathbf{r}_1}$.

Then the marginal distribution of \mathbf{x}_1 is within statistical distance of 8ε of $D_{\mathcal{L}_1 + \mathbf{c}_1, \mathbf{r}}$ and the marginal distribution of \mathbf{x}_2 is as before and the conditional distribution of \mathbf{x}_2 given $\mathbf{x}_1 = \bar{\mathbf{x}}_1 \in \mathcal{L}_1 + \mathbf{c}_1$ is statistically close to $D_{\mathbf{r}}$.

Theorem 2.6 (Adapted from [Lan14, Lemma 1.42]). *Let \mathcal{L} be a lattice in H of dimension n , $\mathbf{r} \in G$ and $s > 0$. Define by \mathbf{t} having $t_i = \sqrt{r_i^2 + s^2}$ for any i . Assume that $\min_i r_i s / t_i \geq \eta_{\varepsilon}(\mathcal{L})$, for some $\varepsilon \leq 1/2$. Consider the distribution Y over H obtained by sampling from $D_{\mathcal{L},\mathbf{r}}$ and then adding an element drawn from D_s . Then Y is statistically close to $D_{\mathbf{t}}$ within 4ε .*

The following is a standard fact from [Reg05, Claim 2.13].

Lemma 2.7. *For any lattice $\mathcal{L} \subset H$ and $\varepsilon \in (0, 1)$, we have $\eta_{\varepsilon}(\mathcal{L}) \geq \sqrt{\log(1/\varepsilon)}/\lambda_1(\mathcal{L}^*)$.*

Here we state some facts used in the “noise swallowing” applications from in Section 7.

Lemma 2.8 (Special case of [AGHS13, Lemma 3]). *For a lattice \mathcal{L} in H of dimension n and $\mathbf{r} \in G$ such that $\min_i r_i \geq \eta_{\varepsilon}(\mathcal{L})$, where $0 < \varepsilon < 1$, then $\Pr_{x \leftarrow D_{\mathcal{L},\mathbf{r}}}(\|x\| \geq \max_i r_i \cdot \sqrt{n}) \leq \frac{1+\varepsilon}{1-\varepsilon} \cdot 2^{-n}$.*

Lemma 2.9 (Special case of [LD18, Proposition 2.1]). *For c an element from \mathbb{R} and D_s an 1-dimensional continuous Gaussian over \mathbb{R} , $\Delta(D_s, D_s + c) \leq C \cdot |c|/s$, where C is equal to $\sqrt{2\pi}/2$.*

This lemma shows that as long as $|c|/s$ is negligible, D_s is statistically close to $D_s + c$.

Lemma 2.10. *For c an element from H and $D_{\mathbf{r}'}$ an n -dimensional continuous Gaussian over H , $\Delta(D_{\mathbf{r}'}, D_{\mathbf{r}'} + c) \leq C\sqrt{n} \cdot \|c\|/\min r'_i$.*

Proof. Since c is in H , consider that $c = c_1 h_1 + c_2 h_2 + \dots + c_n h_n$, for some $c_i \in \mathbb{R}$. Then $\Delta(D_{\mathbf{r}'}, D_{\mathbf{r}'} + c) \leq \sum_i \Delta(D_{r'_i}, D_{r'_i} + c_i) \leq C \cdot \sum_i |c_i|/r'_i \leq C \cdot \sqrt{n} \|c\|/\min r'_i$, where for the second inequality we applied Lemma 2.9 and for the last inequality we used Cauchy Schwarz inequality and the fact that h_1, h_2, \dots, h_n form an orthonormal basis. \square

This lemma shows that as long as $\|c\| \cdot \sqrt{n}/\min r'_i$ is negligible, $D_{\mathbf{r}'}$ is statistically close to $D_{\mathbf{r}'} + c$.

Lemma 2.11. *Let c be an element from R and $\mathbf{r} \in G$, where $\mathbf{r} \geq \sigma$, for some $\sigma \geq \eta_\varepsilon(R)$. Then $D_{R,\mathbf{r}}$ and $D_{R,\mathbf{r}} + c$ are statistically close as long as $\|c\| \cdot \sqrt{n}/\min r_{\sigma,i}$ is negligible, where $\mathbf{r}_\sigma \in G$ is defined by $r_{\sigma,i}^2 := r_i^2 - \sigma^2$.*

Proof. Let z be drawn from $D_{\mathbf{r}_\sigma}$ and e drawn from $D_{R-z,\sigma}$. According to Theorem 2.5, $z + e$ is statistically close to $D_{R,\mathbf{r}}$. Also, $z + e + c$ is again statistically close to $D_{R,\mathbf{r}} + c$. Notice that by triangle inequality we have $\Delta(D_{R,\mathbf{r}}, D_{R,\mathbf{r}} + c) \leq \Delta(D_{R,\mathbf{r}}, z + e) + \Delta(D_{R,\mathbf{r}} + c, z + e + c) + \Delta(z + e, z + e + c)$. The previous arguments show that the first two terms are negligible. Since $c \in R$, then the cosets $R - z$ and $R - (z + c)$ coincide so one can add the same $e \leftarrow D_{R-z,\sigma}$ when discretizing z and $z + c$. Now using the fact that $\Delta(z + e, z + e + c) \leq \Delta(z, z + c)$, by Lemma 2.10 it follows that as long as $\|c\| \cdot \sqrt{n}/\min r_{\sigma,i}$ is negligible, the conclusion holds. \square

Lemma 2.12. *Let c be drawn from $D_{R,\mathbf{t}}$, for some $\mathbf{t} \in G$, and $\mathbf{r} \in G$, where $\mathbf{r} \geq \sigma$, for some $\sigma \geq \eta_\varepsilon(R)$. Then $D_{R,\mathbf{r}}$ and $D_{R,\mathbf{r}} + c$ are statistically close as long as $\|\mathbf{t}\|_\infty \cdot n/\min r_{\sigma,i}$ is negligible.*

Proof. By Lemma 2.8, $\|c\| \leq \sqrt{n} \cdot \|\mathbf{t}\|_\infty$ with non-negligible probability. Lemma 2.11 helps in concluding the proof. \square

This lemma shows that if we add a discrete Gaussian sample to a discrete Gaussian distribution with wide parameter, this gets swallowed. The following definition describes this phenomenon.

Definition 2.13. *Consider \mathbf{r} and $\mathbf{t} \in G$ two Gaussian parameters in G . We say that \mathbf{r} swallows \mathbf{t} if $n \cdot \|\mathbf{t}\|_\infty / \min r_{\sigma,i}$ is negligible in n .*

2.2.2 Lattice Problems

Let \mathcal{L} be a lattice in H represented by a basis \mathbf{B} and let $\mathbf{e} + \mathcal{L}$ be a lattice coset represented by its unique representative $\bar{\mathbf{e}} = (\mathbf{e} + \mathcal{L}) \cap \mathcal{P}(\mathbf{B})$ in the fundamental parallelepiped $\mathcal{P}(\mathbf{B}) := \mathbf{B} \cdot [-1/2, 1/2]^n$ of \mathbf{B} . We state the standard lattice problems.

Definition 2.14 (Gap Shortest Vector Problem). *For an approximation factor $\gamma = \gamma(n) \geq 1$ and a family of lattices \mathfrak{L} , the \mathfrak{L} -GapSVP $_\gamma$ problem is: given a lattice $\mathcal{L} \in \mathfrak{L}$ and length $d > 0$, output YES if $\lambda_1(\mathcal{L}) \leq d$ and NO if $\lambda_1(\mathcal{L}) \geq \gamma d$.*

Definition 2.15 (Shortest Independent Vectors Problem). *For an approximation factor $\gamma = \gamma(n) \geq 1$ and a family of lattices \mathfrak{L} , the \mathfrak{L} -SIVP $_\gamma$ problem is: given a lattice $\mathcal{L} \in \mathfrak{L}$, output n linearly independent lattice vectors of norm at most $\gamma \cdot \lambda_n(\mathcal{L})$.*

Definition 2.16 (Discrete Gaussian Sampling). *For a family of lattices \mathfrak{L} and a function γ that maps lattices from \mathfrak{L} to $G := \{\mathbf{r} \in (\mathbb{R}^+)^n : r_{s_1+s_2+i} = r_{s_1+i}, \text{ for } 1 \leq i \leq s_2\}$, the \mathfrak{L} -DGS $_\gamma$ problem is: given a lattice $\mathcal{L} \in \mathfrak{L}$ and a parameter $\mathbf{r} \geq \gamma(\mathcal{L})$, output an independent sample from a distribution that is within negligible statistical distance of $D_{\mathcal{L},\mathbf{r}}$.*

Definition 2.17 (Bounded Distance Decoding). *For a family of lattices \mathfrak{L} and a function δ that maps lattices from \mathfrak{L} to positive reals, the \mathfrak{L} -BDD $_\delta$ problem is: given a lattice $\mathcal{L} \in \mathfrak{L}$, a distance bound $d \leq \delta(\mathcal{L})$, and a coset $\mathbf{e} + \mathcal{L}$ where $\|\mathbf{e}\| \leq d$, output \mathbf{e} .*

Lemma 2.18 (Babai's round-off algorithm [Bab86], [LPR13, Claim 2.10]). *For every family of lattices \mathfrak{L} , there is an efficient algorithm that given as input a lattice $\mathcal{L} \in \mathfrak{L}$, a set of linearly independent vectors $\{v_1, v_2, \dots, v_n\}$ in \mathcal{L}^* and a coset $\mathbf{e} + \mathcal{L}$ such that $|\langle \mathbf{e}, \bar{v}_i \rangle| \leq \frac{1}{2}$, solves \mathfrak{L} -BDD $_\delta$ for $\delta(\mathcal{L}) = \frac{1}{2\lambda_n(\mathcal{L}^*)}$.*

Definition 2.19 (Gaussian Decoding Problem [PRSD17]). *For a lattice $\mathcal{L} \subset H$ and a Gaussian parameter $g > 0$, the $\text{GDP}_{\mathcal{L},g}$ problem is: given a coset $\mathbf{e} + \mathcal{L}$ where $\mathbf{e} \in H$ was drawn from D_g , find \mathbf{e} .*

2.3 Algebraic Number Theory

2.3.1 Ideals and Orders in Number Fields

Below we present the definitions and facts used in this work, all of which are standard in the area of algebraic number theory. As such, we omit the proofs, which can be found in various textbooks. Unlike previous works, we focus our presentation on *orders* in number fields, and refer to [Cona] for an introduction of invertible ideals in an order, and to [Ste08] for a more thorough background and proofs for the statements that follow.

A *number field* is a field extension $K = \mathbb{Q}(\zeta)$ obtained by adjoining an element ζ to the rationals \mathbb{Q} , where ζ satisfies the relation $f(\zeta) = 0$ for some irreducible polynomial $f(x) \in \mathbb{Q}[x]$, called *minimal polynomial* of ζ , which is monic without the loss of generality. The *degree* n of the number field is the degree of f .

Let K be some number field of degree n . An *order* \mathcal{O} in K is a subring of \mathcal{O}_K that contains a \mathbb{Q} -basis of K , i.e. $\mathcal{O} = \bigoplus_{i=1}^n \mathbb{Z}g_i$ for some $\{g_1, \dots, g_n\} \subset \mathcal{O}$ such that $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q} = K$. The set of orders in K has a unique maximal element (under inclusion), which is called the *maximal-order*, and is denoted by \mathcal{O}_K . An element in a number field $x \in K$ is said to be *integral* if it is the root of some monic polynomial with (rational) integer coefficients. The set of all integral elements in K has a ring structure, called the *ring of integers* and it turns out to be \mathcal{O}_K .

Let \mathcal{O} be some order in K . An ideal $\mathcal{I} \subseteq \mathcal{O}$ is an additive subgroup that is closed under multiplication by \mathcal{O} , i.e. $x \cdot a \in \mathcal{I}$ for every $x \in \mathcal{O}$ and $a \in \mathcal{I}$. Ideals in \mathcal{O} are sometimes called *integral* (as opposed to fractional ideal defined below). Every ideal in \mathcal{O} could be generated by n elements over \mathbb{Z} .

The *sum* of two ideals $\mathcal{I}, \mathcal{J} \subseteq \mathcal{O}$ is defined by $\mathcal{I} + \mathcal{J} := \{x + y \mid x \in \mathcal{I}, y \in \mathcal{J}\}$, and their *product* is defined by $\mathcal{I} \cdot \mathcal{J} := \{\sum x_i y_i \mid x_i \in \mathcal{I}, y_i \in \mathcal{J}\}$. Their *intersection* is simply their set theoretic intersection, and their *quotient* is defined by $(\mathcal{I} : \mathcal{J}) := \{x \in K \mid x\mathcal{J} \subseteq \mathcal{I}\}$. All of the former sets, excluding the quotient, are ideals in \mathcal{O} .

An integral ideal $\mathfrak{p} \subset \mathcal{O}$ is *prime* if for every pair of elements $x, y \in \mathcal{O}$, whenever $xy \in \mathfrak{p}$, then either $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$. Every integral ideal \mathcal{I} of \mathcal{O} contains a product of prime ideals $\mathcal{I} \supseteq \prod \mathfrak{p}_i$. Integral ideals \mathcal{I}, \mathcal{J} of \mathcal{O} are *coprime*, if $\mathcal{I} + \mathcal{J} = \mathcal{O}$ and therefore we also have, $\mathcal{I}\mathcal{J} = \mathcal{I} \cap \mathcal{J}$ and $(\mathcal{I} \cap \mathcal{J})\mathcal{L} = \mathcal{I}\mathcal{L} \cap \mathcal{J}\mathcal{L}$, for any ideal \mathcal{L} . For an integral ideal $\mathcal{I} \subseteq \mathcal{O}$, the set of *associated primes* of \mathcal{I} is the set of all prime ideals of \mathcal{O} that contain \mathcal{I} .

The *norm* of an ideal $\mathcal{I} \subset \mathcal{O}$ is its index as a subgroup of, i.e. $N(\mathcal{I}) := [\mathcal{O} : \mathcal{I}] = |\mathcal{O}/\mathcal{I}|$. For the special case where $\mathcal{O} = \mathcal{O}_K$ is the maximal order, the norm is a multiplicative function, i.e. $N(\mathcal{I}\mathcal{J}) = N(\mathcal{I}) \cdot N(\mathcal{J})$ for any integral \mathcal{I}, \mathcal{J} .

A *fractional ideal* $\mathcal{I} \subset K$ of \mathcal{O} is a set such that $d\mathcal{I} \subset \mathcal{O}$ for some $d \in \mathcal{O}$. We define its norm to be $N(\mathcal{I}) := N(d\mathcal{I})/|N(d)|$. Note that for any fractional ideals \mathcal{I}, \mathcal{J} , their sum, product, quotient and intersection are again fractional ideals.

A fractional ideal \mathcal{I} is *invertible* if there exists a fractional ideal \mathcal{J} such that $\mathcal{I} \cdot \mathcal{J} = \mathcal{O}$. If there exists such \mathcal{J} , then it is unique and equal to $(\mathcal{O} : \mathcal{I})$, and is denoted by \mathcal{I}^{-1} . The set of invertible ideals of \mathcal{O} forms a multiplicative group, with \mathcal{O} being the unit element. It is denoted by

$\mathfrak{I}(\mathcal{O})$. In the special case where $\mathcal{O} = \mathcal{O}_K$ is the maximal order, *every* fractional ideal is invertible. Moreover, every fractional ideal \mathcal{I} of \mathcal{O}_K has *unique factorization* into prime ideals $\mathcal{I} = \prod \mathfrak{p}_i^{e_i}$ for some prime ideals \mathfrak{p}_i and integers $e_i \in \mathbb{Z}$. However, this does not hold for non-maximal orders. There are fractional ideals which are *not invertible*, and there are invertible ideals that are *not a product of prime ideals*. In fact, every invertible ideal in \mathcal{O} is not invertible in any other order \mathcal{O}' .

2.3.2 Embeddings and Geometry

A number field $K = \mathbb{Q}(\zeta)$ of degree n has exactly n field embeddings (injective homomorphisms) $\sigma_i : K \rightarrow \mathbb{C}$. Concretely, these embeddings map ζ to each of the complex root of its minimal polynomial f . An embedding whose images lies in \mathbb{R} (corresponding to a real root of f) is called a *real embedding*; otherwise it is called a *complex embedding*. Because complex roots of f come in conjugate pairs, so too do the complex embeddings. The number of real embeddings is denoted s_1 and the number of pairs of complex embeddings is denoted s_2 , so we have $n = s_1 + 2s_2$. By convention, we let $\{\sigma_j\}_{j \in [s_1]}$ be the real embeddings, and we order the complex embeddings so that $\sigma_{s_1+s_2+j} = \overline{\sigma_{s_1+j}}$ for $j \in [s_2]$. The *canonical embedding* $\sigma : K \rightarrow \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ is then defined as

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

By identifying elements of K with their canonical embeddings on H , we can speak of the norms on K . For any $x \in K$ and any $p \in [1, \infty]$, the ℓ_p norm of x is simply $\|x\|_p = \|\sigma(x)\|_p$. Notice that due to the component-wise multiplication of the embedded elements, we have $\|xy\|_p \leq \|x\|_\infty \cdot \|y\|_p \leq \|x\|_p \cdot \|y\|_p$, for any $x, y \in K$ and $p \in [1, \infty]$.

Using the canonical embedding also allows us to think of the Gaussian distribution $D_{\mathbf{r}}$ over H , or its discrete analogue over lattice in H , as a distribution over K . Strictly speaking, the distribution $D_{\mathbf{r}}$ is not over K , but rather over the field tensor product $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R}$, which is isomorphic to H .

2.3.3 Trace and Norm

The *trace* $Tr = Tr_{K/\mathbb{Q}} : K \rightarrow \mathbb{Q}$, and the *norm* $N = N_{K/\mathbb{Q}} : K \rightarrow \mathbb{Q}$ of an element $x \in K$ are the sum and product, respectively, of the embeddings:

$$Tr(x) := \sum_{i=1}^n \sigma_i(x) \quad N(x) := \prod_{i=1}^n \sigma_i(x) .$$

It is easy to check that the trace is additive, while the norm is multiplicative. Moreover, the (absolute) norm of an element coincides with the norm of the ideal generated by it, in any order \mathcal{O} . That is $|N(x)| = N(x\mathcal{O})$. Also, for all $x, y \in K$,

$$Tr(x \cdot y) = \sum_{i=1}^n \sigma_i(x) \cdot \sigma_i(y) = \langle \sigma(x), \overline{\sigma(y)} \rangle .$$

2.3.4 Orders and Their Ideals as Lattices

Recall that a fractional ideal \mathcal{I} of any order \mathcal{O} has a \mathbb{Z} -basis $U = \{u_1, \dots, u_n\}$. Therefore, under the canonical embedding σ , the ideal yields a full-rank lattice $\sigma(\mathcal{I})$ having basis $\{\sigma(u_1), \dots, \sigma(u_n)\} \subset H$.

In particular, orders themselves are lattices of dimension n in the \mathbb{R} vector space H . We often identify \mathcal{I} with $\sigma(\mathcal{I})$. We call such ideals *ideal lattices*.

The (absolute) *discriminant* $\Delta(\mathcal{O})$ of an order \mathcal{O} is defined to be the square of the determinant of $\sigma(\mathcal{O})$. Equivalently $\Delta(\mathcal{O}) = |\det(\text{Tr}(b_i \cdot b_j))|$, where b_1, \dots, b_n is any basis of \mathcal{O} . Consequently, the determinant of any ideal lattice \mathcal{I} of \mathcal{O} is $N(\mathcal{I})\sqrt{\Delta(\mathcal{O})}$. We denote by Δ_K the *discriminant of K* which is the discriminant of the ring of integers $\Delta(\mathcal{O}_K)$. Moreover, the discriminant of an order \mathcal{O} is related to the discriminant of K by $\Delta(\mathcal{O}) = [\mathcal{O}_K : \mathcal{O}]^2 \Delta_K$.

For the special case of cyclotomic number fields we have the following bound.

Lemma 2.20 ([Was83, Proposition 2.7]). *Let K be a cyclotomic number field of degree n , then $\Delta_K \leq n^n$.*

The following lemma gives upper and lower bounds on the minimum distance of fractional ideals over orders. These bounds are generalizations of the bounds stated in [PR07, Lemma 6.1, Lemma 6.2] and proved using the same framework.

Lemma 2.21. *Let K be some number field of degree n , let \mathcal{O} be an order, and \mathcal{I} a fractional ideal over \mathcal{O} . Then, in any ℓ_p norm, for $p \in [1, \infty]$:*

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

As a corollary we get the following bound of the smoothing parameter, which is a generalization of [PR07, Lemma 6.5] in the case of fractional ideals over orders and its proof is the same.

Lemma 2.22. *Let K be some number field of degree n , let \mathcal{O} be an order and \mathcal{I} a fractional ideal over \mathcal{O} . Then $\eta_{\varepsilon}(\mathcal{I}) \leq N(\mathcal{I})^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n}$, where $\varepsilon = 2^{-n}$.*

We also give an upper bound for the basis minimum $g^p(\mathcal{I})$, which is obtained from [PR07, Lemma 6.3] by generalizing to ideals over orders.

Lemma 2.23. *Let K be a number field of degree n . Then, in any ℓ_p norm, for $p \in [1, \infty]$, there exists a constant C such that for every order \mathcal{O} in K and every \mathcal{O} -fractional ideal \mathcal{I} we have*

$$g^p(\mathcal{I}) \leq C \cdot n^{1/p} \cdot N^{1/n}(\mathcal{I}) \cdot \Delta_{\mathcal{O}}^{1/n} \cdot \sqrt{\log n}.$$

In particular, in the case of ℓ_{∞} norm we have:

$$g^{\infty}(\mathcal{I}) \leq C \cdot N^{1/n}(\mathcal{I}) \cdot \Delta_{\mathcal{O}}^{1/n} \cdot \sqrt{\log n}.$$

All the computational problems defined for general lattices are immediately specialized to ideal lattices. We use standard asymptotic convention where complexity is measured respective to some implicit security parameter which is polynomially related to n . That is, we always consider an infinite ensemble of number fields with asymptotically increasing degree n , and measure all other values as a function of n .

2.3.5 Duality

Let K be a number field, and $\mathcal{O} \subset K$ be some order. For any fractional ideal \mathcal{I} of \mathcal{O} , its *dual* is defined as

$$\mathcal{I}^\vee = \{x \in K \mid \text{Tr}(x\mathcal{I}) \subset \mathbb{Z}\}.$$

It follows that \mathcal{I}^\vee is a fractional ideal and that $\sigma(\mathcal{I}^\vee) = \overline{\sigma(\mathcal{I})^*}$. We recall that the dual of the ring of integers \mathcal{O}_K^\vee is called *co-different* ideal and its norm is $N(\mathcal{O}_K^\vee) = \Delta_K^{-1}$.

We mention some useful properties regarding the duals of ideals.

Lemma 2.24 ([Conb, Section 3][Cona, Section 4]). *Let K be a number field and $\mathcal{O} \subset K$ an order. For any \mathcal{I}, \mathcal{J} fractional ideals of \mathcal{O} the following hold:*

1. $(\mathcal{I}^\vee)^\vee = \mathcal{I}$.
2. $\mathcal{I} \subset \mathcal{J} \iff \mathcal{J}^\vee \subset \mathcal{I}^\vee$.
3. $(\mathcal{I} + \mathcal{J})^\vee = \mathcal{I}^\vee \cap \mathcal{J}^\vee$.
4. $(\mathcal{I} \cap \mathcal{J})^\vee = \mathcal{I}^\vee + \mathcal{J}^\vee$.
5. $\mathcal{I} \cdot \mathcal{I}^\vee = \mathcal{O}^\vee$ if $\{x \in K \mid x\mathcal{I} \subset \mathcal{I}\} = \mathcal{O}$. If $\mathcal{O} = \mathcal{O}_K$, the equality holds for any ideal \mathcal{I} .
6. Further assuming that \mathcal{I} is invertible, $(\mathcal{I}\mathcal{J})^\vee = \mathcal{I}^{-1}\mathcal{J}^\vee$.

From the last item, the following is an immediate corollary:

Corollary 2.25. *Let \mathcal{I} be a fractional ideal over an order \mathcal{O} and let $0 \neq \alpha \in K$ be a nonzero field element. Then*

$$(\alpha\mathcal{I})^\vee = \frac{1}{\alpha}\mathcal{I}^\vee.$$

The dual of an order is not necessarily an invertible ideal. However there are some special cases where it is, described in the lemma below:

Lemma 2.26 ([Conb, Section 3]). *If $\mathcal{O} = \mathbb{Z}[x]/(f)$ for some monic irreducible $f \in \mathbb{Q}[x]$, then $\mathcal{O}^\vee = \frac{1}{f'(\alpha)}\mathcal{O}$ where $\alpha \in \mathbb{C}$ is some root of f . Therefore $(\mathcal{O}^\vee)^{-1} = f'(\alpha)\mathcal{O}$ and in particular \mathcal{O}^\vee is an invertible \mathcal{O} -ideal.*

We now give a simple lemma involving the smoothing parameter of product of ideals.

Lemma 2.27. *Let \mathcal{O} be an order in a number field K . Let \mathcal{I}, \mathcal{J} be fractional \mathcal{O} -ideals. Assume that \mathcal{I} is invertible. Then, for every $\varepsilon > 0$,*

$$\frac{\eta_\varepsilon(\mathcal{J})}{\lambda_1^\infty(\mathcal{I}^{-1})} \leq \eta_\varepsilon(\mathcal{I} \cdot \mathcal{J}) \leq \lambda_1^\infty(\mathcal{I}) \cdot \eta_\varepsilon(\mathcal{J})$$

where λ_1^∞ denotes the length of the shortest vector in the lattice with respect to the ℓ_∞ -norm.

Proof. By the definition of the smoothing parameter, and Lemma 2.24,

$$\begin{aligned} \eta_\varepsilon(\mathcal{I} \cdot \mathcal{J}) &= \arg \min_{s>0} \{ \rho_{1/s}((\mathcal{I} \cdot \mathcal{J})^\vee \setminus \{0\}) \leq \varepsilon \} \\ &= \arg \min_{s>0} \{ \rho_{1/s}(\mathcal{I}^{-1} \cdot \mathcal{J}^\vee \setminus \{0\}) \leq \varepsilon \} \end{aligned}$$

Let $v \in \mathcal{I}$ be such that $\|v\| = \lambda_1^\infty(\mathcal{I})$. Since $v\mathcal{O} \subseteq \mathcal{I}$, then $v^{-1}\mathcal{O} \supseteq \mathcal{I}^{-1}$, and so $v^{-1}\mathcal{J}^\vee \supseteq \mathcal{I}^{-1} \cdot \mathcal{J}^\vee$. Hence, we get that,

$$\arg \min_{s>0} \{ \rho_{1/s}((\mathcal{I}^{-1} \cdot \mathcal{J}^\vee) \setminus \{0\}) \leq \varepsilon \} \leq \arg \min_{s>0} \{ \rho_{1/s}((v^{-1} \cdot \mathcal{J}^\vee) \setminus \{0\}) \leq \varepsilon \} .$$

For every $x \in \mathcal{J}^\vee$, we have that

$$\|v^{-1}x\| \geq \min_{i \in [n]} |\sigma_i(v^{-1})| \|x\| = \|x\| / \|v\|_\infty = \|x\| / \lambda_1^\infty(\mathcal{I}) .$$

Thus,

$$\arg \min_{s>0} \{ \rho_{1/s}((v^{-1} \cdot \mathcal{J}^\vee) \setminus \{0\}) \leq \varepsilon \} \leq \arg \min_{s>0} \left\{ \rho_{\lambda_1^\infty(\mathcal{I})/s}(\mathcal{J}^\vee \setminus \{0\}) \leq \varepsilon \right\} ,$$

and the upper bound follows.

Using this bound we obtained, we have $\eta_\varepsilon(\mathcal{I}^{-1} \cdot \mathcal{IJ}) \leq \lambda_1^\infty(\mathcal{I}^{-1}) \cdot \eta_\varepsilon(\mathcal{IJ})$. By reversing the inequality we get the lower bound. \square

2.3.6 The Conductor Ideal

In this section we recall the definition of the *conductor ideal* as well as a few basic facts. We note that there is a much richer theory and we only mention the properties used in this work. See [Cona] for more details.

Definition 2.28. *The conductor of an order \mathcal{O} is defined to be the ideal*

$$\mathcal{C}_\mathcal{O} = (\mathcal{O} : R) = (R^\vee : \mathcal{O}^\vee) .$$

It is the maximal set that is both an \mathcal{O}_K and an \mathcal{O} -ideal.

We mention here properties of the \mathcal{O} ideals coprime to the conductor ideal. We begin with stating the following theorem which shows that R ideals which are coprime to the conductor ideal are in 1-to-1 correspondence to \mathcal{O} ideals which are coprime to the conductor ideal and then we present some corollaries of this result.

Theorem 2.29 ([Cona, Theorem 3.8]). *The nonzero ideals over \mathcal{O} coprime to $\mathcal{C}_\mathcal{O}$ and the nonzero ideals over R coprime to $\mathcal{C}_\mathcal{O}$ are in multiplicative bijection by mapping $\mathcal{I} \mapsto \mathcal{I}R$ and $\mathcal{J} \mapsto \mathcal{J} \cap \mathcal{O}$.*

Corollary 2.30 ([Cona, Corollary 3.10]). *If qR is coprime to the conductor, then $qR \cap \mathcal{O} = q\mathcal{O}$.*

Corollary 2.31 ([Cona, Corollary 3.11]). *The ideals over \mathcal{O} coprime to $\mathcal{C}_\mathcal{O}$ have unique factorization into prime ideals over \mathcal{O} .*

We conclude the discussion on the conductor ideal with the following lemma.

Lemma 2.32. $\mathcal{C}_\mathcal{O}^{-1}R^\vee = R\mathcal{O}^\vee$, where $\mathcal{C}_\mathcal{O}^{-1}$ is the inverse of the conductor ideal as R ideal.

Proof. By the definition of the dual of an ideal we have: $(R\mathcal{O}^\vee)^\vee = \{x \in K \mid \text{Tr}(xR\mathcal{O}^\vee) \subseteq \mathbb{Z}\} = \{x \in K \mid xR \subseteq (\mathcal{O}^\vee)^\vee\} = \{x \in K \mid xR \subseteq \mathcal{O}\} = \mathcal{C}_\mathcal{O}$. Hence $\mathcal{C}_\mathcal{O}^\vee = R\mathcal{O}^\vee$, so $\mathcal{C}_\mathcal{O}^{-1}R^\vee = R\mathcal{O}^\vee$. \square

Remark 2.33. *By the proof of Lemma 2.32, one can get that $R^\vee = \mathcal{C}_\mathcal{O}R\mathcal{O}^\vee = \mathcal{C}_\mathcal{O}\mathcal{O}^\vee$.*

2.3.7 Cancellation of Ideals

The next lemma is a generalization of [LPR10, Lemma 2.15]. It is crucially used to make a BDD instance and a DGS sample into an Order-LWE instance in the hardness result in Section 3. Generally speaking, the lemma allows us to cancel invertible factors in the quotient $\mathcal{IL}/\mathcal{IJL}$ to yield an isomorphism onto \mathcal{L}/\mathcal{JL} by multiplying by an appropriate “tweak” factor. The proof of the lemma uses a generalization of the Chinese Remainder Theorem adapted for ideals over orders.

Theorem 2.34 (Chinese Remainder Theorem). *Let \mathcal{I} be a fractional \mathcal{O} -ideal, and let $\mathfrak{p}_1, \dots, \mathfrak{p}_k$ be k distinct prime ideals in \mathcal{O} . The canonical \mathcal{O} -module homomorphism*

$$\mathcal{I} / \left(\prod \mathfrak{p}_i \right) \mathcal{I} \rightarrow \bigoplus \mathcal{I} / \mathfrak{p}_i \mathcal{I}$$

is an isomorphism.

Lemma 2.35. *Let \mathcal{I}, \mathcal{J} be integral ideals in an order \mathcal{O} and let \mathcal{L} be a fractional \mathcal{O} -ideal. Assume that \mathcal{I} is invertible. Given the associated primes of \mathcal{J} , $\mathfrak{p}_1, \mathfrak{p}_2, \dots, \mathfrak{p}_k$, and an element $t \in \mathcal{I} \setminus \bigcup_{i=1}^k \mathfrak{p}_i \mathcal{I}$ the map*

$$\begin{aligned} \theta_t : \mathcal{L} / \mathcal{JL} &\rightarrow \mathcal{IL} / \mathcal{IJL} \\ x &\mapsto t \cdot x \end{aligned}$$

is well-defined, and induces an isomorphism of \mathcal{O} -modules. Moreover, θ_t is efficiently inverted given $\mathcal{I}, \mathcal{J}, \mathcal{L}$ and t . Finally, such t can be computed given \mathcal{I} and $\mathfrak{p}_1, \mathfrak{p}_2, \dots, \mathfrak{p}_k$.

Proof. Assume that t further satisfies that $t\mathcal{I}^{-1} + \mathcal{J} = \mathcal{O}$. Therefore $t\mathcal{L} + \mathcal{IJL} = \mathcal{IL}$, and the multiplication by t map induces an \mathcal{O} -module isomorphism

$$\frac{\mathcal{L}}{\mathcal{JL}} \xrightarrow{\sim} \frac{t\mathcal{L}}{t\mathcal{JL}} = \frac{t\mathcal{L}}{t\mathcal{L} \cap \mathcal{IJL}} \simeq \frac{t\mathcal{L} + \mathcal{IJL}}{\mathcal{IJL}} = \frac{\mathcal{IL}}{\mathcal{IJL}}$$

where for the first equality we used the fact that $t\mathcal{I}^{-1}$ and \mathcal{J} are coprime and hence $t\mathcal{JL} = (t\mathcal{I}^{-1} \cap \mathcal{J})\mathcal{IL} = t\mathcal{L} \cap \mathcal{IJL}$. The isomorphism above is the inclusion map used in the second law of isomorphism and for the last equality we used $t\mathcal{L} + \mathcal{IJL} = \mathcal{IL}$.

We now show that the condition $t \in \mathcal{I} \setminus \bigcup_{i=1}^k \mathfrak{p}_i \mathcal{I}$ implies that $t\mathcal{I}^{-1} + \mathcal{J} = \mathcal{O}$. If it is not the case, then $t\mathcal{I}^{-1} + \mathcal{J}$ must be included in a maximal ideal \mathfrak{m} of \mathcal{O} . Hence $t\mathcal{I}^{-1}$ and \mathcal{J} are included in \mathfrak{m} . The last inclusion suggests that \mathfrak{m} is an associated prime of \mathcal{J} , so it is a prime ideal \mathfrak{p}_i . Therefore, the first inclusion suggests that t belongs to $\mathfrak{p}_i \mathcal{I}$, which comes in contradiction with the choice of t .

Finally we show how to compute such t . Since \mathcal{I} is invertible and $\mathfrak{p}_i \subsetneq \mathcal{O}$, the quotient $\mathcal{I}/\mathfrak{p}_i \mathcal{I} \neq 0$. For each $i \in [k]$, choose a non-zero $t_i \in \mathcal{I}/\mathfrak{p}_i \mathcal{I}$. Let $t \in \mathcal{I}$ be the pre-image of $(t_1, t_2, \dots, t_k) \in \bigoplus \mathcal{I}/\mathfrak{p}_i \mathcal{I}$, under the Chinese remainder theorem. It is clear that $t \notin \bigcup_{i=1}^k \mathfrak{p}_i \mathcal{I}$. □

2.4 A Counting Lemma

Here we consider a rational prime q and $\mathfrak{q}_1, \dots, \mathfrak{q}_k$ its associated primes when considered as an \mathcal{O} ideal. The following lemma assures us of existence of small elements in an ideal \mathcal{Q}^{-1} outside the ideals $\mathcal{Q}^{-1}\mathfrak{q}_i$, for any i .

Lemma 2.36. *Let K be a degree n number field and \mathcal{O} an order, let q be a rational prime. Let \mathcal{Q} be an invertible \mathcal{O} -ideal and $\mathfrak{q}_1, \dots, \mathfrak{q}_k$ be the associated primes of $q\mathcal{O}$. Then there exists $u \in \mathcal{Q}^{-1} \setminus \bigcup_i \mathcal{Q}^{-1}\mathfrak{q}_i$ of norm $\|u\|_\infty \leq O\left(\frac{n\sqrt{\log n}\Delta_{\mathcal{O}}^{1/n}}{N(\mathcal{Q})^{1/n}}\right)$.*

Proof. We prove the lemma using a counting argument. We show that the number of elements in \mathcal{Q}^{-1} with bounded ℓ_∞ -norm is larger than the total number of elements in the union $\bigcup_i \mathcal{Q}^{-1}\mathfrak{q}_i$. This gives an existential proof, which is sufficient for our use.

We denote by \mathcal{C}_r the hypercube (of dimension n) with edge length $2r$, i.e. $\mathcal{C}_r = (-r, r)^n$. Note that a point $x \in \mathcal{C}_r$ if and only if $\|x\|_\infty \leq r$. We turn to bound the number of lattice points with bounded ℓ_∞ -norm, or equivalently the number of lattice points that intersect the hypercube.

Claim 2.36.1. *For every $r > 0$ and \mathcal{L} a fractional ideal over \mathcal{O} we have that*

$$\frac{(2(r - ng^\infty(\mathcal{L})))^n}{\det(\mathcal{L})} \leq \left| \mathcal{L} \cap \mathcal{C}_r \right| \leq \frac{(2(r + ng^\infty(\mathcal{L})))^n}{\det(\mathcal{L})}.$$

Proof. Let $\{v_1, \dots, v_n\} \subset \mathcal{L}$ be a basis in the lattice \mathcal{L} , such that $\|v_i\|_\infty \leq g^\infty(\mathcal{L})$. We denote the fundamental parallelepiped, which we call it shortly a box, defined by $\{v_1, \dots, v_n\}$ by $\mathcal{B}_{\vec{v}}$. Note that its volume is $\det(\mathcal{L})$. It is easy to see that the translates of the box by lattice points are disjoint. Moreover, if x is in $\mathcal{B}_{\vec{v}}$, then $x = \alpha_1 v_1 + \dots + \alpha_n v_n$, for some scalars from $\alpha_i \in [0, 1)$, so $\|x\|_\infty \leq \sum \|\alpha_i v_i\|_\infty \leq \sum \|v_i\|_\infty \leq n \cdot g^\infty(\mathcal{L})$. Therefore, the parallelepiped is contained in $\mathcal{C}_{ng^\infty(\mathcal{L})}$.

Starting from the upper bound, we take the box $\mathcal{B}_{\vec{v}}$ around each lattice point. We get that

$$\bigcup_{x \in \mathcal{L} \cap \mathcal{C}_r} x + \mathcal{B}_{\vec{v}} \subseteq \mathcal{C}_{r+ng^\infty(\mathcal{L})}.$$

Indeed, for every $x \in \mathcal{L} \cap \mathcal{C}_r$ and $y \in x + \mathcal{B}_{\vec{v}}$, we have that $\|y\|_\infty \leq \|x\|_\infty + ng^\infty(\mathcal{L}) \leq r + ng^\infty(\mathcal{L})$, therefore $y \in \mathcal{C}_{r+ng^\infty(\mathcal{L})}$. Hence,

$$\left| \mathcal{L} \cap \mathcal{C}_r \right| \cdot \det(\mathcal{L}) \leq (2(r + ng^\infty(\mathcal{L})))^n,$$

and the upper bound follows.

For the lower bound, we use the same covering, and note that

$$\mathcal{C}_{r-ng^\infty(\mathcal{L})} \subseteq \bigcup_{x \in \mathcal{L} \cap \mathcal{C}_r} x + \mathcal{B}_{\vec{v}}.$$

Let $y \in \mathcal{C}_{r-ng^\infty(\mathcal{L})}$ and let $x \in \mathcal{L}$ be the lattice point such that $y \in x + \mathcal{B}_{\vec{v}}$. Since $\|x\|_\infty \leq \|y\|_\infty + ng^\infty(\mathcal{L}) \leq r$, then $x \in \mathcal{C}_r$. We get that

$$(2(r - ng^\infty(\mathcal{L})))^n \leq \left| \mathcal{L} \cap \mathcal{C}_r \right| \cdot \det(\mathcal{L}),$$

which concludes the proof. ■

We now use those bounds to count the number of lattice points. Indeed, we get that,

$$\begin{aligned} \left| \mathcal{Q}^{-1} \cap \mathcal{C}_r \right| &\geq \frac{(2(r - ng^\infty(\mathcal{Q}^{-1})))^n}{\det(\mathcal{Q}^{-1})} \\ &\geq \frac{(2(r - Cn\sqrt{\log n}\Delta_{\mathcal{O}}^{1/n}N(\mathcal{Q}^{-1})^{1/n}))^n}{\det(\mathcal{Q}^{-1})}, \end{aligned}$$

where the last inequality is by Lemma 2.23. On the other hand, we have

$$\begin{aligned}
\left| \bigcup_{i=1}^k \mathcal{Q}^{-1} \mathbf{q}_i \cap \mathcal{C}_r \right| &\leq \sum_{i=1}^k \left| \mathcal{Q}^{-1} \mathbf{q}_i \cap \mathcal{C}_r \right| \\
&\leq \sum_{i=1}^k \frac{(2(r + ng^\infty(\mathcal{Q}^{-1} \mathbf{q}_i)))^n}{\det(\mathcal{Q}^{-1} \mathbf{q}_i)} \\
&= \sum_{i=1}^k \frac{(2(r + ng^\infty(\mathcal{Q}^{-1} \mathbf{q}_i)))^n}{N(\mathbf{q}_i) \det(\mathcal{Q}^{-1})} \\
&\leq \sum_{i=1}^k \frac{(2(r + Cn\sqrt{\log n} \Delta_{\mathcal{Q}}^{1/n} N(\mathbf{q}_i)^{1/n} N(\mathcal{Q}^{-1})^{1/n}))^n}{N(\mathbf{q}_i) \det(\mathcal{Q}^{-1})} \\
&= \sum_{i=1}^k \frac{N(\mathbf{q}_i) (2(r/N(\mathbf{q}_i)^{1/n} + Cn\sqrt{\log n} \Delta_{\mathcal{Q}}^{1/n} N(\mathcal{Q}^{-1})^{1/n}))^n}{N(\mathbf{q}_i) \det(\mathcal{Q}^{-1})} \\
&\leq \sum_{i=1}^k \frac{(2(r/q^{1/n} + Cn\sqrt{\log n} \Delta_{\mathcal{Q}}^{1/n} N(\mathcal{Q}^{-1})^{1/n}))^n}{\det(\mathcal{Q}^{-1})} \\
&= k \cdot \frac{(2(r/q^{1/n} + Cn\sqrt{\log n} \Delta_{\mathcal{Q}}^{1/n} N(\mathcal{Q}^{-1})^{1/n}))^n}{\det(\mathcal{Q}^{-1})}
\end{aligned}$$

where the first inequality is the union bound, the second inequality holds because of Claim 2.36.1, the third inequality is due to Lemma 2.23. The fourth inequality follows since $N(\mathbf{q}_i) \geq q$ for all i (since $\mathbf{q}_i | q\mathcal{O}$).

Comparing the two inequalities we get that for $r \geq \frac{n\sqrt{\log n} \Delta_{\mathcal{Q}}^{1/n}}{N(\mathcal{Q})^{1/n}} \cdot \frac{C(1+k^{1/n})}{1-(k/q)^{1/n}} = \Omega\left(\frac{n\sqrt{\log n} \Delta_{\mathcal{Q}}^{1/n}}{N(\mathcal{Q})^{1/n}}\right)$, there exists $u \in (\mathcal{Q}^{-1} \setminus \bigcup_i \mathcal{Q}^{-1} \mathbf{q}_i) \cap \mathcal{C}_r$. \square

2.5 The Ring-LWE Problem

Let $q \geq 2$ be a (rational) integer. Let $\mathbb{T} = K_{\mathbb{R}}/R^{\vee}$ denote a torus in the Minkowski space. For any fractional ideal \mathcal{I} of R , let $\mathcal{I}_q := \mathcal{I}/q\mathcal{I}$.

Definition 2.37 (Ring-LWE Distribution). *For $s \in R_q^{\vee}$, referred to as “the secret”, and an error distribution ψ over $K_{\mathbb{R}}$, a sample from the R -LWE distribution $A_{s,\psi}$ over $R_q \times \mathbb{T}$ is generated by sampling $a \xleftarrow{\$} R_q$, $e \leftarrow \psi$, and outputting $(a, b = a \cdot s/q + e \bmod R^{\vee})$.*

Definition 2.38 (Ring-LWE, Average-Case Decision Problem). *Let φ be a distribution over R_q^{\vee} , and let Υ be a distribution over a family of error distributions, each over $K_{\mathbb{R}}$. The average-case Ring-LWE decision problem, denoted $R\text{-LWE}_{q,\varphi,\Upsilon}$, is to distinguish between independent samples from $A_{s,\psi}$ for a random choice of a “secret” $s \leftarrow \varphi$, and an error distribution $\psi \leftarrow \Upsilon$, and the same number of uniformly random and independent samples from $R_q \times \mathbb{T}$.*

We recall the error distribution defined in [PRSD17, Definition 6.1]. We stick to the notation used in the reference.

Definition 2.39. Fix an arbitrary $f(n) = \omega(\sqrt{\log n})$. For a real $\alpha > 0$, a distribution sampled from Υ_α is an elliptical Gaussian $D_{\mathbf{r}}$, where $\mathbf{r} \in G$ is sampled as follows: for each $1 \leq i \leq s_1$, sample $x_i \leftarrow D_1$ and set $r_i^2 = \alpha^2(x_i^2 + f^2(n))/2$. For each $s_1 + 1 \leq i \leq s_1 + s_2$, sample $x_i, y_i \leftarrow D_{1/\sqrt{2}}$ and set $r_i^2 = r_{i+s_2}^2 = \alpha^2(x_i^2 + y_i^2 + f^2(n))/2$.

Theorem 2.40 ([PRSD17, Theorem 6.2]). Let K be an arbitrary field of degree n and $R = \mathcal{O}_K$ its ring of integers. Let $\alpha = \alpha(n) \in (0, 1)$, and let $q = q(n) \geq 2$ be a (rational) integer such that $\alpha q \geq 2\omega(1)$. There is a polynomial-time quantum reduction from $\mathfrak{J}(R)$ -DGS $_\gamma$ to R -LWE $_{q, U(R_q^\vee), \Upsilon_\alpha}$, where

$$\gamma = \max \left\{ \eta(\mathcal{L}) \cdot \sqrt{2}/\alpha \cdot \omega(1), \sqrt{2n}/\lambda_1(\mathcal{L}^\vee) \right\}.$$

3 Order-LWE: Definition, Variants and Worst-Case Hardness

The ring of integers R of a number field K plays a central role in the definition and use of the Ring-LWE problem. However, the ring of integers is a special member of a family of rings in a number field, known as *orders*. We present a generalization of Ring-LWE which we call Order-LWE, and show that similar to Ring-LWE it also enjoys worst-case hardness, but with respect to a different set of lattices. Generalizing the problem to the setting of orders also exposes a difference between two variants of Ring-LWE that are indeed identical when considering the ring of integers, but are distinct for general orders. Some background on algebraic number theory and particularly on orders can be found in Section 2.3.

In the original R -LWE definition [LPR10], the secret s was sampled from the dual of the ring of integers R^\vee (modulo q), and the coefficients a were sampled from R (modulo q). We similarly define \mathcal{O} -LWE as a sequence of noisy linear univariate equations where the secret is sampled from \mathcal{O}^\vee and the coefficients are sampled from \mathcal{O} . As pointed out in [LPR10], a dual version where s is sampled from R and a from R^\vee can also be defined, and is equivalent to the original one. Indeed some followup works used the alternative definition (e.g. [GHPS13]). In the context of orders, we show that this distinction can make a difference. We denote the dual version by \mathcal{O}^\vee -LWE. While we are able to show worst-case hardness reductions for both \mathcal{O} -LWE and \mathcal{O}^\vee -LWE, the classes of lattices for which worst-case hardness holds is different for the two variants; one is the dual of the other. Our definition also generalizes R -LWE in another dimension, by allowing to take equations modulo arbitrary ideals, and not necessarily modulo (an ideal generated by) a rational integer q . In this section we define the variants of Order-LWE and present the worst-case hardness results.

To set up the problems, let K be a number field, and let \mathcal{O} be an order in it. Let \mathcal{Q} be an integral \mathcal{O} -ideal, and let $u \in (\mathcal{O} : \mathcal{Q}) := \{x \in K : x\mathcal{Q} \subseteq \mathcal{O}\}$. For fractional \mathcal{O} -ideals \mathcal{J} and \mathcal{L} , define $\mathcal{J}_\mathcal{L} := \mathcal{J}/\mathcal{L}$, and let $\mathbb{T}_{\mathcal{O}^\vee} := K_\mathbb{R}/\mathcal{O}^\vee$.

Definition 3.1 (\mathcal{O} -LWE Distribution). For $s \in \mathcal{O}_\mathcal{Q}^\vee$ and an error distribution ψ over $K_\mathbb{R}$, a sample from the \mathcal{O} -LWE distribution $\mathcal{O}_{s, \psi, u}$ over $\mathcal{O}_\mathcal{Q} \times \mathbb{T}_{\mathcal{O}^\vee}$ is generated by sampling $a \xleftarrow{\$} \mathcal{O}_\mathcal{Q}$, $e \leftarrow \psi$ and outputting $(a, b = u \cdot (a \cdot s) + e \bmod \mathcal{O}^\vee)$.

Definition 3.2 (\mathcal{O} -LWE, Average-Case Decision Problem). Let φ be a distribution over $\mathcal{O}_\mathcal{Q}^\vee$ and let Υ be a distribution over a family of error distributions, each over $K_\mathbb{R}$. The average-case \mathcal{O} -LWE decision problem, denoted \mathcal{O} -LWE $_{(\mathcal{Q}, u), \varphi, \Upsilon}$, is to distinguish between independent samples from $\mathcal{O}_{s, \psi, u}$, for a random choice of a “secret” $s \leftarrow \varphi$, and an error distribution $\psi \leftarrow \Upsilon$, and the same number of uniformly random and independent samples from $\mathcal{O}_\mathcal{Q} \times \mathbb{T}_{\mathcal{O}^\vee}$.

When the secret is sampled from the uniform distribution over \mathcal{O}_Q^\vee , we sometimes omit it from the subscript. Observe that when $\mathcal{O} = \mathcal{O}_K$, $Q = q\mathcal{O}_K$ and $u = 1/q$, the \mathcal{O} -LWE problem coincides with the Ring-LWE problem.

In our definition of an \mathcal{O} -LWE distribution, the secret $s \in \mathcal{O}_Q^\vee$ and $a \in \mathcal{O}_Q$. One can also consider a dual variant of \mathcal{O} -LWE where $a \in \mathcal{O}_Q^\vee$ and $s \in \mathcal{O}_Q$. In general, these two variants are not equivalent, unlike in the case of Ring-LWE (see Remark 3.5), but for special orders \mathcal{O} they are, namely for orders \mathcal{O} such that their duals \mathcal{O}^\vee are invertible as \mathcal{O} -ideals. For example, if f is the minimal polynomial of the number field K , then the ring $\mathcal{O} = \mathbb{Z}[x]/(f)$ is an order in K , whose dual is invertible according to Lemma 2.26.

Definition 3.3 (\mathcal{O}^\vee -LWE Distribution). *For $s \in \mathcal{O}_Q$ and an error distribution ψ over $K_\mathbb{R}$, a sample from the \mathcal{O}^\vee -LWE distribution $\mathcal{O}_{s,\psi,u}^\vee$ over $\mathcal{O}_Q^\vee \times \mathbb{T}_{\mathcal{O}^\vee}$ is generated by sampling $a \xleftarrow{\$} \mathcal{O}_Q^\vee$, $e \leftarrow \psi$, and outputting $(a, b = u \cdot a \cdot s + e \bmod \mathcal{O}^\vee)$.*

Definition 3.4 (\mathcal{O}^\vee -LWE, Average-Case Decision Problem). *Let φ be a distribution over \mathcal{O}_Q , and let Υ be a distribution over a family of error distributions, each over $K_\mathbb{R}$. The average-case \mathcal{O}^\vee -LWE decision problem, denoted $\mathcal{O}^\vee\text{-LWE}_{(Q,u),\varphi,\Upsilon}$, is to distinguish between independent samples from $\mathcal{O}_{s,\psi,u}^\vee$, for a random choice of a “secret” $s \leftarrow \varphi$, and an error distribution $\psi \leftarrow \Upsilon$, and the same number of uniformly random and independent samples from $\mathcal{O}_Q^\vee \times \mathbb{T}_{\mathcal{O}^\vee}$.*

As before, when the secret is sampled from the uniform distribution over \mathcal{O}_Q , we sometimes omit it from the subscript. Similar to the case of \mathcal{O} -LWE, when $\mathcal{O} = \mathcal{O}_K$, $Q = q\mathcal{O}_K$ and $u = 1/q$, the \mathcal{O}^\vee -LWE problem coincides with the variant of the Ring-LWE problem where a is sampled from R^\vee/qR^\vee and s is sampled from R/qR .

Remark 3.5. *The \mathcal{O} -LWE problem and the \mathcal{O}^\vee -LWE problem are equivalent as long as \mathcal{O}^\vee is an invertible \mathcal{O} -ideal. By Lemma 2.35, the invertibility of \mathcal{O}^\vee yields an isomorphism from \mathcal{O}_Q^\vee to \mathcal{O}_Q induced by multiplication by $t \in (\mathcal{O}^\vee)^{-1}$. Therefore, the samples of the form $(a, b = u \cdot a \cdot s + e \bmod \mathcal{O}^\vee)$ are transformed to $(a' = a \cdot t, b' = b = u \cdot a' \cdot s' + e \bmod \mathcal{O}^\vee)$, where $a' = a \cdot t \in \mathcal{O}_Q$ and $s' = s \cdot t^{-1} \in \mathcal{O}_Q^\vee$. In the particular case of \mathcal{O} being the ring of integers, we obtain the equivalence between Ring-LWE and the variant of Ring-LWE previously described.*

The \mathcal{O}^\vee -LWE definition is inspired by [GHPS13], where the authors show that for the variant of Ring-LWE with a from the dual and s from the ring, problem becomes harder as the number field grows. In Section 5, we prove an analogue of this result for the set of orders under inclusion, i.e., the bigger the order is, the harder the \mathcal{O}^\vee -LWE problem is. Since the ring of integers is the maximal order in the field, the Ring-LWE problem is harder than any \mathcal{O}^\vee -LWE problem.

3.1 Worst-Case Hardness for \mathcal{O} -LWE and \mathcal{O}^\vee -LWE

We now state the hardness results of the \mathcal{O} -LWE and \mathcal{O}^\vee -LWE problems and derive the hardness of the Ring-LWE problem (see Theorem 2.40) as a special case. We begin by generalizing Definition 2.39 of the error distribution Υ_α to be elliptical according to u .

Definition 3.6. *Fix an arbitrary $f(n) = \omega(\sqrt{\log n})$. For $\alpha > 0$ and $u \in K$, a distribution sampled from $\Upsilon_{u,\alpha}$ is an elliptical Gaussian $D_\mathbf{r}$, where $\mathbf{r} \in G$ is sampled as follows: for $i = 1, \dots, s_1$, sample $x_i \leftarrow D_1$ and set $r_i^2 = \alpha^2(x_i^2 + (f(n) \cdot |\sigma_i(u)| / \|u\|_\infty)^2)/2$. For $i = s_1 + 1, \dots, s_1 + s_2$, sample $x_i, y_i \leftarrow D_{1/\sqrt{2}}$ and set $r_i^2 = r_{i+s_2}^2 = \alpha^2(x_i^2 + y_i^2 + (f(n) \cdot |\sigma_i(u)| / \|u\|_\infty)^2)/2$.*

Note that when $u \in K$ satisfies $\sigma_1(u) = \dots = \sigma_n(u)$ (and therefore is rational), the distribution $\Upsilon_{u,\alpha}$ degenerates to Υ_α . Otherwise, $\Upsilon_{u,\alpha}$ is strictly narrower than Υ_α .

Recall that $\mathfrak{I}(\mathcal{O})$ is the set of invertible fractional ideals over the order \mathcal{O} . Our hardness results for \mathcal{O} -LWE and \mathcal{O}^\vee -LWE are as follows.

Theorem 3.7. *Let K be an arbitrary number field of degree n and $\mathcal{O} \subset K$ an order. Let \mathcal{Q} be an integral \mathcal{O} -ideal, $u \in (\mathcal{O} : \mathcal{Q})$ and let $\alpha \in (0, 1)$ be such that $\alpha / \|u\|_\infty \geq 2 \cdot \omega(1)$. There is a polynomial-time quantum reduction from $\mathfrak{I}(\mathcal{O})$ -DGS $_\gamma$ to \mathcal{O} -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$, where*

$$\gamma = \max \left\{ \eta(\mathcal{Q}\mathcal{L}) \cdot \sqrt{2} \|u\|_\infty / \alpha \cdot \omega(1), \sqrt{2n}/\lambda_1(\mathcal{L}^\vee) \right\}. \quad (1)$$

Theorem 3.8. *Let K be an arbitrary number field of degree n and $\mathcal{O} \subset K$ an order. Let \mathcal{Q} be an integral \mathcal{O} -ideal, $u \in (\mathcal{O} : \mathcal{Q})$ and let $\alpha \in (0, 1)$ be such that $\alpha / \|u\|_\infty \geq 2 \cdot \omega(1)$. There is a polynomial-time quantum reduction from $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ -DGS $_\gamma$ to \mathcal{O}^\vee -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$, where*

$$\gamma = \max \left\{ \eta(\mathcal{Q}\mathcal{L}) \cdot \sqrt{2} \|u\|_\infty / \alpha \cdot \omega(1), \sqrt{2n}/\lambda_1(\mathcal{L}^\vee) \right\}. \quad (2)$$

We note that the class $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ is exactly the class of all lattices whose dual is in $\mathfrak{I}(\mathcal{O})$. Thus we see that the effect of changing the domains of a and s to the dual of their previous domains is that the class of lattices for which the hardness result applies is the dual of the previous class. The classes are the same if \mathcal{O}^\vee itself is an invertible ideal in \mathcal{O} . An equivalence between the problems can be shown in this case directly, similar to the setting in Ring-LWE.

Remark 3.9. *Consider the special case where $\mathcal{O} = R$, the ideal $\mathcal{Q} = qR$ and $u = 1/q$. Then the \mathcal{O}^\vee -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ problem is equivalent to the R -LWE $_{q,\Upsilon_\alpha}$ problem as mentioned in Remark 3.5. Moreover, the sets $\mathfrak{I}(R) \cdot R^\vee$ and $\mathfrak{I}(R)$ are equal as all fractional R -ideals are invertible, and finally $\eta(\mathcal{Q}\mathcal{L}) \|u\|_\infty = \eta(\mathcal{L})$ shows that the parameters γ from Theorem 2.40, Theorem 3.8 and Theorem 3.7 coincide. In fact, the expression for γ in Theorem 3.7 is achieved in the latter two results when $\mathcal{Q} = q\mathcal{O}$ and $u = 1/q$, for any order \mathcal{O} .*

Remark 3.10. *Another important special case is the R -LWE distribution with an ideal modulus \mathcal{Q} in place of the integer modulus qR . Formally, we let \mathcal{O} be R , choose $u \in (R : \mathcal{Q}) = \mathcal{Q}^{-1}$ such that $\|u\|_\infty = \lambda_1^\infty(\mathcal{Q}^{-1})$ and $\alpha < \sqrt{\log n/n}$. Then the theorem above implies a reduction from $\mathfrak{I}(R)$ -DGS $_\gamma$ to R -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ with γ greater than at most $\Delta_K^{1/n}$ times the γ obtained when $\mathcal{Q} = qR$ and $u = 1/q$, as in Theorem 2.40. Indeed, the lower bound in Lemma 2.27 implies that the first term in Eq. (1) is at least $\eta(\mathcal{L}) \cdot \sqrt{2}/\alpha \cdot \omega(1)$, which is greater than the second term as long as $\alpha < \sqrt{\log n/n}$, since $\eta(\mathcal{L}) > \omega(\sqrt{\log n})/\lambda_1(\mathcal{L}^\vee)$ (Lemma 2.7). Now the first term in (1) is at most $\Delta_K^{1/n} \cdot \eta(\mathcal{L}) \cdot \sqrt{2}/\alpha \cdot \omega(1)$, according to Lemma 2.21 and the upper bound from Lemma 2.7. Making a similar comparison with modulus \mathcal{Q} , an invertible \mathcal{O} -ideal, we get reductions from $\mathfrak{I}(\mathcal{O})$ -DGS $_\gamma$ to \mathcal{O} -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ and from $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ -DGS $_\gamma$ to \mathcal{O}^\vee -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ with γ greater than at most $\Delta_{\mathcal{O}}^{1/n}$ times the γ obtained when $\mathcal{Q} = q\mathcal{O}$ and $u = 1/q$ in Theorem 3.7 and 3.8.*

3.2 Proving Worst-Case Hardness

We present a proof for Theorem 3.7. The proof for Theorem 3.8 is completely analogous and follows by replacing the point of reference from lattices in $\mathfrak{I}(\mathcal{O})$ to their dual (see Section 3.2.3). Our proofs are related to ones in previous works and in particular to the one in [PRSD17], see Section 3.2.4 for details.

Proof Overview. At the heart of the reduction, we use an iterative step that transforms discrete Gaussian samples into slightly narrower ones. Initially, we generate samples from a Gaussian distribution with large enough parameter such that these samples can be generated efficiently. Then we repeatedly apply the iterative step to generate narrower and narrower samples until we obtain the desired parameter. The iterative step is formally stated in Lemma 3.15.

The proof of the promised iteration has two components. First, we show how to transform an \mathcal{O} -LWE solver into a GDP solver given polynomially many discrete Gaussian sample using Lemma 3.13. This lemma assumes that we are given an efficient algorithm that transforms BDD like instances into \mathcal{O} -LWE samples. We describe such an algorithm in Lemma 3.16. The second step uses Lemma 3.14 to generate narrower discrete Gaussian samples via a GDP solver.

3.2.1 Building Blocks from Previous Works

Recall the definitions of the set G from Section 2.2.1 and of the Gaussian Decoding Problem (Definition 2.19). We will use the following components from previous works. First, we define a specific set of Gaussian parameters that are used in the proof.

Definition 3.11 ([PRSD17, Definition 6.4]). *For real numbers $\zeta > 0$, $r > 0$ and an integer $T \geq 1$, we define the following set $W_{r,\zeta,T} = \{\mathbf{r}_{i,j} : 1 \leq i \leq s_1 + s_2, 0 \leq j \leq T\} \subset G$, where each $\mathbf{r}_{i,j}$ is defined as*

$$(\mathbf{r}_{i,j})_k = \begin{cases} r \cdot (1 + \zeta)^j, & \text{if } k = i \\ r, & \text{else} \end{cases}$$

if $i \leq s_1$, and otherwise as

$$(\mathbf{r}_{i,j})_k = \begin{cases} r \cdot (1 + \zeta)^j, & \text{if } k \in \{i, i + s_2\} \\ r, & \text{else} \end{cases}.$$

Notice that all vectors \mathbf{r} from $W_{r,\zeta,T}$ have all coordinates greater than or equal to r . In our hardness proof we will consider $\zeta = 1/\text{poly}(n)$ and $T = \text{poly}(n)$ and for the simplicity of notation, their corresponding subscripts will be dropped.

The next lemma from [Reg05] shows that one can sample from a discrete Gaussian distribution over any lattice, given that the parameter is large enough.

Lemma 3.12 ([Reg05, Lemma 3.2]). *There exists a polynomial-time algorithm that given any n -dimensional lattice \mathcal{L} and parameter $r > 2^{2n}\lambda_n(\mathcal{L})$ outputs a sample from $D_{\mathcal{L},\mathbf{r}}$ for any $\mathbf{r} \in W_r$.*

The following lemma shows how to convert an \mathcal{O} -LWE solver along with sufficiently many Gaussian samples into a GDP solver.

Lemma 3.13 ([PRSD17, Adaptation of Lemma 6.6]⁹). *Consider an algorithm with the following interface. Its input is an integral ideal \mathcal{Q} , an element $u \in (\mathcal{O} : \mathcal{Q})$, a fractional ideal $\mathcal{L} \in \mathfrak{I}(\mathcal{O})$, a coset $e + \mathcal{L}^\vee$ and a bound $d \geq \|e\|_\infty$, a real number $r \geq \sqrt{2} \cdot \eta(\mathcal{Q}\mathcal{L})$, and polynomially many samples from the discrete Gaussian distribution $D_{\mathcal{L},\mathbf{r}}$ for some $\mathbf{r} \geq r$. The algorithm outputs samples that*

⁹The lemma in [PRSD17] shows that if some algorithm exists, then a Ring-LWE reduction exists, and then it explicitly presents the algorithm as a sub-claim of the lemma. We observe that deriving the reduction from the algorithm follows for Order-LWE almost verbatim by just replacing “Ring-LWE” with “Order-LWE” (and adapting the algorithm interface to refer to ideals in orders). This lemma is a statement of this observation.

are statistically close to the \mathcal{O} -LWE distribution $\mathcal{O}_{s,u,D_{\mathbf{r}'}}$, where s is distributed uniformly over $\mathcal{O}_{\mathcal{Q}}^{\vee}$ and the coordinates of \mathbf{r}' are given by $(r'_i)^2 := (r_i |\sigma_i(e)\sigma_i(u)|)^2 + (rd \|u\|_{\infty})^2$.

There exists a probabilistic polynomial-time algorithm that given an algorithm satisfying the above interface; an oracle that solves \mathcal{O} -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ on input an integral \mathcal{O} -ideal \mathcal{Q} , an element $u \in (\mathcal{O} : \mathcal{Q})$, and a real number $\alpha \in (0, 1)$; a real number $r \geq \sqrt{2} \cdot \eta(\mathcal{Q}\mathcal{L})$, a fractional ideal $\mathcal{L} \in \mathfrak{I}(\mathcal{O})$ and polynomially many samples from the discrete Gaussian distribution $D_{\mathcal{L},\mathbf{r}}$ for some $\mathbf{r} \geq r$; it solves GDP $_{\mathcal{L}^{\vee},g}$ for any $g = o(1) \cdot \alpha / (\|u\|_{\infty} \cdot r)$.

Finally, the next lemma states that one can sample from a Gaussian distribution over a lattice \mathcal{L} given a GDP solver over the dual lattice.

Lemma 3.14 ([PRSD17, Lemma 6.7]). *There is an efficient quantum algorithm that, given any n -dimensional lattice \mathcal{L} , a real number $g < \lambda_1(\mathcal{L}^{\vee}) / (2\sqrt{2n})$, a vector $\mathbf{r} \geq 1$, and an oracle that solves GDP $_{\mathcal{L}^{\vee},g}$ (with all but negligible probability), outputs an independent sample from $D_{\mathcal{L},\mathbf{r}/(2g)}$.*

3.2.2 Proof of Theorem 3.7

The proof of the theorem relies on the following lemma.

Lemma 3.15. *There exists an efficient quantum algorithm that given an oracle that solves \mathcal{O} -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ on input an integral \mathcal{O} -ideal \mathcal{Q} , an element $u \in (\mathcal{O} : \mathcal{Q})$, and a real number $\alpha \in (0, 1)$; along with a fractional ideal $\mathcal{L} \in \mathfrak{I}(\mathcal{O})$, a real number $r \geq \sqrt{2} \cdot \eta(\mathcal{Q}\mathcal{L})$ such that $r' := r \cdot \|u\|_{\infty} / \alpha \cdot \omega(1) > \sqrt{2n} / \lambda_1(\mathcal{L}^{\vee})$ and polynomially many samples from the discrete Gaussian distribution $D_{\mathcal{L},\mathbf{r}}$ for each $\mathbf{r} \in W_r$, and a vector $\mathbf{r}' \in G$ where $\mathbf{r}' \geq r'$, outputs an independent sample from $D_{\mathcal{L},\mathbf{r}'}$.*

The proof of Theorem 3.7 given Lemma 3.15 follows, we then prove the lemma.

Proof of Theorem 3.7. Let $\mathcal{L} \in \mathfrak{I}(\mathcal{O})$. We use an \mathcal{O} -LWE $_{(\mathcal{Q},u),\Upsilon_{u,\alpha}}$ solver to solve DGS_{γ} on \mathcal{L} , where γ is as in Theorem 3.7. To this end, we begin with samples from a wide enough Gaussian for each $\mathbf{r} \in W_r$, where $r \geq 2^{2n} \lambda_n(\mathcal{L})$. This bound ensures that sampling from $D_{\mathcal{L},\mathbf{r}}$ is an efficient process by Lemma 3.12. We iteratively run the algorithm described in Lemma 3.15 to generate samples from $D_{\mathcal{L},\mathbf{r}'}$, for each $\mathbf{r}' \in W_{r'}$, where $r' = r \cdot \omega(1) \cdot \|u\|_{\infty} / \alpha \leq r/2$. (recall that $\alpha / \|u\|_{\infty} \geq 2 \cdot \omega(1)$). By the conditions on r and r' , we can apply the iterative step as long as $r' \geq \gamma = \max \{ \eta(\mathcal{Q}\mathcal{L}) \cdot \sqrt{2} \|u\|_{\infty} / \alpha \cdot \omega(1), \sqrt{2n} / \lambda_1(\mathcal{L}^{\vee}) \}$. As the width of the Gaussian decreases exponentially in each iteration, after polynomially many applications of Lemma 3.15 we generate a sample from $D_{\mathcal{L},\gamma}$, as desired. \square

We now prove Lemma 3.15. The proof relies on the existence of an algorithm that transforms BDD instances into \mathcal{O} -LWE samples, as described in the following lemma which is proven below.

Lemma 3.16. *There is a probabilistic polynomial time algorithm that takes as input: an integral ideal \mathcal{Q} , $u \in (\mathcal{O} : \mathcal{Q})$, a fractional ideal $\mathcal{L} \in \mathfrak{I}(\mathcal{O})$, a coset $e + \mathcal{L}^{\vee}$ and a bound $d \geq \|e\|_{\infty}$, a real number $r \geq \sqrt{2} \cdot \eta(\mathcal{Q}\mathcal{L})$, and polynomially many samples from the discrete Gaussian distribution $D_{\mathcal{L},\mathbf{r}}$ for some $\mathbf{r} \geq r$, and outputs samples that are statistically close to the \mathcal{O} -LWE distribution $\mathcal{O}_{s,u,D_{\mathbf{r}'}}$, where the coordinates of \mathbf{r}' are given by $(r'_i)^2 := (r_i |\sigma_i(e)\sigma_i(u)|)^2 + (rd \|u\|_{\infty})^2$.*

We now prove Lemma 3.15 using Lemma 3.16.

Proof of Lemma 3.15. Let $\mathcal{L} \in \mathcal{J}(\mathcal{O})$. Consider the Gaussian samples from $D_{\mathcal{L}, \mathbf{r}}$, for \mathbf{r} in W_r , where $r \geq \sqrt{2}\eta(Q\mathcal{L})$, as mentioned in the statement of Lemma 3.15. These samples along with the given \mathcal{O} -LWE solver, yield a $\text{GDP}_{\mathcal{L}^\vee, g}$ -solver, for $g = 1/(2r') = o(1) \cdot \alpha/(\|u\|_\infty \cdot r)$. See Lemma 3.13. As mentioned earlier, the proof of Lemma 3.13 requires an efficient algorithm to transform a BDD instance into an \mathcal{O} -LWE sample. We describe this algorithm in Lemma 3.16 below. Since $r' > \sqrt{2n}/\lambda_1(\mathcal{L}^\vee)$, we have $g < \lambda_1(\mathcal{L}^\vee)/(2\sqrt{2n})$ and $\mathbf{r}' = \mathbf{1}/(2g)$, for $\mathbf{1} \in W_1$. Therefore, by Lemma 3.14, this $\text{GDP}_{\mathcal{L}^\vee, g}$ -solver yields an independent sample from $D_{\mathcal{L}, \mathbf{r}'}$, via a quantum algorithm. \square

Finally, we prove Lemma 3.16.

Proof of Lemma 3.16. By a scaling argument, we may assume that $\mathcal{L} \subseteq \mathcal{O}$ is an integral ideal. Let $y = x + e$ be a point in H with $x \in \mathcal{L}^\vee$ and $\|e\|_\infty \leq d$. We input y as a representative of the coset $\mathcal{L}^\vee + e$ in the algorithm. The key point that makes the algorithm work is the existence of an element $t \in \mathcal{L}$, guaranteed by Lemma 2.35, that induces the following isomorphisms

$$\theta_t : \mathcal{O}/\mathcal{Q}\mathcal{O} \xrightarrow{\cdot t} \mathcal{L}\mathcal{O}/\mathcal{Q}\mathcal{L}\mathcal{O} = \mathcal{L}/\mathcal{Q}\mathcal{L} \quad (3)$$

$$\theta_t : \mathcal{L}^\vee/\mathcal{Q}\mathcal{L}^\vee \xrightarrow{\cdot t} \mathcal{L}\mathcal{L}^\vee/\mathcal{Q}\mathcal{L}\mathcal{L}^\vee = \mathcal{O}^\vee/\mathcal{Q}\mathcal{O}^\vee \quad (4)$$

Recall that we assumed full knowledge of the set of associated primes of \mathcal{Q} when we fixed the modulus (since ideal factorization is solvable in quantum polynomial time). This implies that θ_t is an efficiently computable map.

The algorithm described in the statement works as follows: on input $z \leftarrow D_{\mathcal{L}, \mathbf{r}}$ and an error term $e' \leftarrow D_{rd\|u\|_\infty}$, it outputs the following tuple

$$(a = \theta_t^{-1}(z \bmod \mathcal{Q}\mathcal{L}), \quad b = u \cdot z \cdot y + e' \bmod \mathcal{O}^\vee) \in \mathcal{O}_{\mathcal{Q}} \times \mathbb{T}_{\mathcal{O}^\vee}$$

We show that (a, b) is indeed an \mathcal{O} -LWE sample, as claimed. The distribution $D_{\mathcal{L}, \mathbf{r}} \bmod \mathcal{Q}\mathcal{L}$ is statistically close to the uniform distribution over $\mathcal{L}_{\mathcal{Q}}$ by Lemma 2.4, as $\mathbf{r} \geq \eta(\mathcal{Q}\mathcal{L})$. Since θ_t^{-1} induces a bijection between $\mathcal{L}/\mathcal{Q}\mathcal{L}$ and $\mathcal{O}/\mathcal{Q}\mathcal{O}$, the variable a is distributed statistically close to $U(\mathcal{O}_{\mathcal{Q}})$.

In order to analyze the marginal distribution of b conditioned on some value of a , observe that

$$b = u \cdot z \cdot y + e' = u \cdot z \cdot x + u \cdot z \cdot e + e' \bmod \mathcal{O}^\vee.$$

Rewriting the first term

$$u \cdot z \cdot x = u \cdot (t^{-1}z) \cdot (tx) \bmod \mathcal{O}^\vee$$

we get that $s := t \cdot x \bmod \mathcal{Q}\mathcal{O}^\vee$ is a well-defined element in $\mathcal{O}_{\mathcal{Q}}^\vee$. Also, $a = z \cdot t^{-1} \bmod \mathcal{Q}\mathcal{O}$. Since \mathcal{O}^\vee is a fractional \mathcal{O} -ideal and $u \in (\mathcal{O} : \mathcal{Q})$, the first term of b , namely $u \cdot z \cdot x = u \cdot a \cdot s \bmod \mathcal{O}^\vee$ is well-defined. Finally, we analyze the error term $u \cdot z \cdot e + e'$. Assume that $e \neq 0$, otherwise the conclusion holds trivially. Write $u \cdot z \cdot e + e'$ as $(z + \frac{e'}{ue}) \cdot ue$, and observe that conditioned on a , the distribution of z is $D_{\mathcal{Q}\mathcal{L}+c, \mathbf{r}}$ for $c = \theta_t(a)$. As $e' \leftarrow D_{rd\|u\|_\infty}$ and $\|e\|_\infty \leq d$, the term $\frac{e'}{ue} \leftarrow D_{\mathbf{t}}$, where $\mathbf{t} = \left(\frac{rd\|u\|_\infty}{|\sigma_i(u)||\sigma_i(e)|} \right)_i > r$. Let $\frac{e'}{ue} = f + g$, where $f \leftarrow D_r$ and $g \leftarrow D_{\mathbf{t}'}$ with \mathbf{t}' defined as $\mathbf{t}'^2 = \mathbf{t}^2 - r^2$, are sampled from independent continuous Gaussian distributions. Thus we have,

$$\left(z + \frac{e'}{ue} \right) \cdot ue = (z + f + g) \cdot ue.$$

Since $r \geq \sqrt{2} \cdot \eta(\mathcal{QL})$ and $\mathbf{r} \geq r$, Theorem 2.6 implies that the distribution of $z + f$ is statistically close to $D_{\sqrt{\mathbf{r}_i^2 + r^2}}$. Therefore, the distribution of $z + f + g$ is statistically close to $D_{\sqrt{\mathbf{r}_i^2 + \frac{(rd\|u\|_\infty)^2}{|\sigma_i(e)\sigma_i(u)|^2}}}$, thus implying that $u \cdot z \cdot e + e'$ is sampled from a distribution statistically close to $D_{\mathbf{r}'}$. \square

3.2.3 Proving Theorem 3.8

The only difference required to prove Theorem 3.8 is a change in Lemma 3.16. This consists in applying Lemma 2.35 differently in the BDD-to- \mathcal{O}^\vee -LWE reduction (Lemma 3.16); namely replace the maps in Eq. (3), (4) with the following:

$$\theta_t : \mathcal{O}^\vee / \mathcal{Q}\mathcal{O}^\vee \xrightarrow{t} \mathcal{L}'\mathcal{O}^\vee / \mathcal{Q}\mathcal{L}'\mathcal{O}^\vee = \mathcal{L} / \mathcal{Q}\mathcal{L} \quad (5)$$

$$\theta_t : \mathcal{L}^\vee / \mathcal{Q}\mathcal{L}^\vee \xrightarrow{t} \mathcal{L}'\mathcal{L}^\vee / \mathcal{Q}\mathcal{L}'\mathcal{L}^\vee = \mathcal{O} / \mathcal{Q}\mathcal{O}, \quad (6)$$

where $\mathcal{L} = \mathcal{L}' \cdot \mathcal{O}^\vee \in \mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ for \mathcal{L}' an invertible fractional \mathcal{O} -ideal and $t \in \mathcal{L}'$.

3.2.4 Relation of Our Proof to Previous Works

Our proof of Theorems 3.7, 3.8 follows the blueprint analogous to the hardness proofs of both LWE [Reg05], and R -LWE [LPR10, PRSD17]. Specifically, the proof using the iterative step given in Lemma 3.15, that is further proved via a two step argument are common ingredients in all these previous works (and in fact rely on some of lemmas from there without any change).

The main difference between our proof and the previous works is Lemma 3.16 which generalizes Lemma 4.8 from [LPR10] and Lemma 6.8 from [PRSD17]. These latter lemmas (from previous works) show how to use Gaussian samples to transform BDD instances into Ring-LWE samples. Our work (Lemma 3.16) generalizes this algorithm to generate \mathcal{O} -LWE samples for any order \mathcal{O} rather than just for the ring of integers. The proof of the lemma is almost identical to the ones from previous works. The novelty comes from using the cancellation lemma adapted to the case of fractional ideals over orders (Lemma 2.35).

4 New Worst-Case Hardness for Polynomial-LWE

The *Polynomial Learning with Errors* problem, or PLWE in short, introduced by Stehlé et al. [SSTX09]¹⁰ is closely related to both the Ring-LWE and Order-LWE problems. PLWE has an advantage of having very simple interface which is useful for manipulations and thus also for applications and implementations. In a recent work, Rosca, Stehlé and Wallet [RSW18] showed a reduction from worst-case ideal-lattice problems to PLWE. In this section, we show that the hardness of \mathcal{O} -LWE that we proved in Section 3 implies a different worst-case hardness result for PLWE, essentially by relating it to a different class of lattices than those considered in [LPR10, PRSD17]. In what follows we start with an informal description of the PLWE problem, the current hardness result of PLWE, our result and a comparison. This is followed by a more detailed and formal treatment.

¹⁰As “ideal-LWE”. The name PLWE was used in [BV11a].

4.1 Overview

Consider a number field K defined by an irreducible polynomial f , so that $K = \mathbb{Q}[x]/(f)$. Recall that the Ring-LWE distribution involves elements a and s of the ring of integers $R := \mathcal{O}_K$ and its dual R^\vee , respectively. The \mathcal{O} -LWE distribution is defined similarly, but with a and s coming from an arbitrary order in K and its dual, respectively. In the PLWE setting, both a and s are elements of the ring $\mathcal{O} := \mathbb{Z}[x]/(f)$, i.e. polynomials with integer coefficients in the number field. There are number fields for which $R \neq \mathcal{O}$, however it is always true that \mathcal{O} is an order of K . We highlight that in PLWE, unlike in Ring-LWE and Order-LWE, both a and s are elements of the order itself.¹¹

The aforementioned [RSW18] presented a reduction from Ring-LWE to PLWE (see Theorem 4.2 for the formal statement). Their reduction is based on the so called “Cancellation Lemma” (Lemma 2.35) which, informally, allows to “reshape” orders and ideals at the cost of increasing the size of the error. As mentioned above, in PLWE both a and s are elements of \mathcal{O} , whereas in Ring-LWE a and s are elements of R and R^\vee respectively. The reduction of [RSW18] applies the Cancellation Lemma to reshape both R and R^\vee into \mathcal{O} . We mention that using the Cancellation Lemma to reshape ideals of the ring of integers R is a known technique (see [Pei16, Section 2.3.2]). The novel contribution of [RSW18] is both in analyzing the increase of the error and in reshaping ideals of one order into another.

We suggest an alternative reduction from \mathcal{O} -LWE to PLWE in Theorem 4.4. Our reduction is also based on the reshaping procedure, but with a single application of the Cancellation Lemma. More specifically, we only need to reshape \mathcal{O}^\vee into \mathcal{O} . We show below that our reduction increases the error by a smaller factor than in the reduction of [RSW18] from Ring-LWE. See Proposition 4.9 for the formal statement.

4.2 Hardness of PLWE

The formal definitions and hardness results follow, along with a more detailed and formal comparison of the results. We let K be a number field of degree n defined by a polynomial f . We denote $\mathcal{O} := \mathbb{Z}[x]/(f)$, and $R := \mathcal{O}_K$. The PLWE distribution and problem are defined as follows.

Definition 4.1 (PLWE Distribution and Problem [SSTX09]). *For a rational integer $q \geq 2$, a ring element $s \in \mathcal{O}_q$, and an error distribution ψ over $K_{\mathbb{R}}/\mathcal{O}$, the PLWE distribution over $\mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}$, denoted by $P_{s,\psi}$, is sampled by independently choosing a uniformly random $a \xleftarrow{\$} \mathcal{O}_q$ and an error term $e \leftarrow \psi$, and outputting $(a, b = (a \cdot s)/q + e \bmod \mathcal{O})$.*

For a distribution Υ over a family of error distributions, each over $K_{\mathbb{R}}/\mathcal{O}$, the PLWE decision problem, denoted $\text{PLWE}_{q,\Upsilon}$, is to distinguish between independent samples from $P_{s,\psi}$ for a random choice of $s \xleftarrow{\$} \mathcal{O}_q$, and an error distribution $\psi \leftarrow \Upsilon$, and the same number of uniformly random and independent samples from $\mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}$.

For a distribution φ and an element $t \in K$ we denote by $t \cdot \varphi$ the distribution obtained by sampling an element $x \leftarrow \varphi$ and outputting $t \cdot x$. Similarly, for a family distribution Υ , we denote by $t \cdot \Upsilon$ the family obtained by multiplying each distribution by t .

¹¹Another difference between Ring/Order-LWE and PLWE is that in the latter, the error distribution is specified using the so called *coefficients embedding*, and not the *canonical embedding*. For the sake of simplicity, we focus on a variant of PLWE which uses the canonical embedding, (called PLWE^σ in [RSW18]) but we call it likewise, and we avoid the distinction between the embeddings. Both hardness results in this section can be further extended to the hardness of PLWE.

We now turn to present and compare the two worst-case to average-case reductions. Recall the definition of the conductor ideal, denoted by $\mathcal{C}_{\mathcal{O}}$, from Definition 2.28. [RSW18] showed the following reduction from Ring-LWE to PLWE.

Theorem 4.2 ([Pei16, Section 2.3.2][RSW18, Theorem 4.2]). *Let $q \geq 2$ be some rational integer such that $qR + \mathcal{C}_{\mathcal{O}} = R$, and let Υ be a distribution over a family of error distributions, each over $K_{\mathbb{R}}/\mathcal{O}$. There exists a probabilistic polynomial time reduction from $R\text{-LWE}_{q,\Upsilon}$ to $\text{PLWE}_{q,t_1 t_2^2 \cdot \Upsilon}$, where $t_1 \in (R : R^{\vee}) \setminus \bigcup_i \mathfrak{p}_i(R : R^{\vee})$ and $t_2 \in \mathcal{C}_{\mathcal{O}} \setminus \bigcup_i \mathfrak{p}_i \mathcal{C}_{\mathcal{O}}$, where \mathfrak{p}_i 's are the prime ideals of qR .*

Combining the reduction above with the hardness of Ring-LWE stated in Theorem 2.40 we get the following:

Corollary 4.3 (Worst-Case Hardness of PLWE from Ring-LWE). *With the same notations as above, let $\alpha \in (0, 1)$ such that $\alpha q \geq 2 \|t_1 t_2^2\|_{\infty} \omega(1)$. There is a reduction from $\mathfrak{I}(R)\text{-DGS}_{\gamma}$ to $\text{PLWE}_{q,\Upsilon_{\alpha}}$ for any*

$$\gamma = \max \left\{ \eta(\mathcal{L}) \cdot \sqrt{2}/\alpha \cdot \|t_1 t_2^2\|_{\infty} \cdot \omega(1), \sqrt{2n}/\lambda_1(\mathcal{L}^{\vee}) \right\}.$$

We now state the hardness result based on the hardness of \mathcal{O} -LWE. First, using a reduction similar to the one from [Pei16, Section 2.3.2], we obtain an analogous reduction from \mathcal{O} -LWE to PLWE.

Theorem 4.4. *Let $q \geq 2$ be some rational integer, and let Υ be a distribution over a family of error distributions, each over $K_{\mathbb{R}}/\mathcal{O}$. There exists a probabilistic polynomial time reduction from $\mathcal{O}\text{-LWE}_{q,\Upsilon}$ to $\text{PLWE}_{q,t \cdot \Upsilon}$, where $t \in (\mathcal{O} : \mathcal{O}^{\vee}) \setminus \bigcup_i \tilde{\mathfrak{p}}_i(\mathcal{O} : \mathcal{O}^{\vee})$, where $\tilde{\mathfrak{p}}_i$'s are the associated primes of $q\mathcal{O}$.*

Proof. Recall that $\mathcal{O} = \mathbb{Z}[x]/f$, and so by Lemma 2.26 \mathcal{O}^{\vee} is an invertible \mathcal{O} -ideal. Using Lemma 2.35, there exists an element $t \in (\mathcal{O} : \mathcal{O}^{\vee}) \setminus \bigcup_i \tilde{\mathfrak{p}}_i(\mathcal{O} : \mathcal{O}^{\vee})$ such that the mapping $x \mapsto t \cdot x$ is an isomorphism $\mathcal{O}^{\vee}/q\mathcal{O}^{\vee} \rightarrow \mathcal{O}/q\mathcal{O}$.

The reduction follows from the following efficient transformation that takes as input a pair $(a, b) \in \mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}^{\vee}$ and outputs another $(a', b') \in \mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}$. We show that this transformation maps the uniform distribution over $\mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}^{\vee}$ to the uniform distribution over $\mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}$, and $\mathcal{O}_{s,\varphi}$ to $P_{t \cdot s, t \cdot \varphi}$ for any element $s \in \mathcal{O}^{\vee}/q\mathcal{O}^{\vee}$ and distribution φ . Note that in the latter, if s was sampled uniformly over $\mathcal{O}^{\vee}/q\mathcal{O}^{\vee}$, then by our choice of t , $t \cdot s$ is distributed uniformly over $\mathcal{O}/q\mathcal{O}$.

The transformation is the following. Given a pair $(a, b) \in \mathcal{O}_q \times K_{\mathbb{R}}/\mathcal{O}^{\vee}$, the transformation outputs

$$a' = a, \quad b' = t \cdot b \pmod{\mathcal{O}}.$$

Note that if (a, b) is distributed uniformly over $\mathcal{O} \times K_{\mathbb{R}}/\mathcal{O}^{\vee}$ then (a', b') is distributed uniformly over $\mathcal{O} \times K_{\mathbb{R}}/\mathcal{O}$, by the choice of t . On the other hand, if $(a, b = 1/q \cdot a \cdot s + e \pmod{\mathcal{O}^{\vee}})$ are sampled from $\mathcal{O}_{s,\varphi}$, then the distribution of $(a', b' = 1/q \cdot a \cdot (t \cdot s) + t \cdot e \pmod{\mathcal{O}})$ is $P_{t \cdot s, t \cdot \varphi}$. \square

Now, using the hardness of \mathcal{O} -LWE from Theorem 3.7 we obtain:

Corollary 4.5 (Worst-Case Hardness of PLWE from Order-LWE). *Let $q \geq 2$ be some rational integer, and let $\alpha \in (0, 1)$ be such that $\alpha q \geq 2 \|t\|_{\infty} \omega(1)$. Then there is a reduction from $\mathfrak{I}(\mathcal{O})\text{-DGS}_{\gamma}$ to $\text{PLWE}_{q,\Upsilon_{\alpha}}$ for any*

$$\gamma = \max \left\{ \eta(\mathcal{L}) \cdot \sqrt{2}/\alpha \cdot \|t\|_{\infty} \cdot \omega(1), \sqrt{2n}/\lambda_1(\mathcal{L}^{\vee}) \right\}.$$

4.3 On the Existence of Small Multipliers

In the following, we consider α a root of the defining polynomial f of the field K , so $\mathbb{Z}[x]/(f) = \mathbb{Z}[\alpha]$. We also denote by $\tilde{\mathfrak{p}}_1, \dots, \tilde{\mathfrak{p}}_k$ the associated primes of $q\mathcal{O}$, where q is a rational prime.

We assert that a short element t as in Theorem 4.4 exists and can be found using the combinatorial argument from Lemma 2.36.

Corollary 4.6. *There exists an element t in $(\mathcal{O} : \mathcal{O}^\vee) \setminus \bigcup_i (\mathcal{O} : \mathcal{O}^\vee) \tilde{\mathfrak{p}}_i$ whose norm is bounded by*

$$\|t\| \leq O(n \cdot \sqrt{n \log n} \cdot N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n}).$$

Proof. In the setup of Lemma 2.36, let $\mathcal{O} = \mathbb{Z}[X]/(f) = \mathbb{Z}[\alpha]$, $\mathcal{Q} = (\mathcal{O} : \mathcal{O}^\vee)^{-1}$ and $\tilde{\mathfrak{p}}_i$ be the associated primes of $q\mathcal{O}$. This yields a short t in $(\mathcal{O} : \mathcal{O}^\vee) \setminus \bigcup_i (\mathcal{O} : \mathcal{O}^\vee) \tilde{\mathfrak{p}}_i$, whose norm $\|t\|_\infty \leq O\left(\frac{n\sqrt{\log n} \Delta_{\mathcal{O}}^{1/n}}{N(\mathcal{Q})^{1/n}}\right)$.

Note that $N(\mathcal{Q}) = N((\mathcal{O} : \mathcal{O}^\vee)^{-1}) = N(\mathcal{O}^\vee)$. Therefore, by Lemma 2.26, we have $N(\mathcal{Q}) = N(1/f'(\alpha)\mathcal{O}) = 1/N(f'(\alpha))$, and the length of t is bounded above by

$$\|t\| \leq \sqrt{n} \cdot \|t\|_\infty \leq O(n \cdot \sqrt{n \log n} \cdot N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n}).$$

□

A short element t as in Theorem 4.4 can also be obtained by sampling via a Gaussian distribution over $(\mathcal{O} : \mathcal{O}^\vee)$ with an appropriately small parameter, exactly as in [RSW18, Theorem 3.1]. We present a generalized version of this result here, for the sake of completion. We use it to derive another bound for our multiplier t . We would like to point out that the assumption that q is coprime to the conductor is a crucial one here, whereas it was not needed in Corollary 4.6 above.

Proposition 4.7 ([RSW18, Adaptation of Theorem 3.1]). *Let \mathcal{I} and \mathcal{J} be two integral ideals over an order \mathcal{O} such that \mathcal{I} is an invertible ideal and \mathcal{J} is coprime to the conductor. Let $\mathfrak{p}_1, \dots, \mathfrak{p}_k$ be the primes from the unique prime decomposition of \mathcal{J} and ordered in a non-decreasing order according to their norm. Let $\delta \in \left[\frac{4n+\log \Delta_{\mathcal{O}}}{\log N(\mathcal{J})}, 1\right]$ and*

$$s = \begin{cases} (N(\mathcal{J})^{1/2} N(\mathcal{I}) \Delta_{\mathcal{O}})^{1/n}, & \text{if } N(\mathfrak{p}_k) \geq N(\mathcal{J})^{1/2+\delta} \\ (N(\mathcal{J})^{1/2+2\delta} N(\mathcal{I}) \Delta_{\mathcal{O}})^{1/n}, & \text{else} \end{cases}$$

Then

$$\Pr_{t \leftarrow D_{\mathcal{I},s}} [t\mathcal{I}^{-1} + \mathcal{J} = \mathcal{O}] \geq 1 - \frac{k}{N(\mathfrak{p}_1)} - 2^{-n+4}.$$

The proof of the proposition is identical to that of [RSW18, Theorem 3.1], once Lemma 2.5, and Lemma 2.7 from [RSW18] are generalized to \mathcal{O} -ideal lattices in Lemma 2.21 and Lemma 2.22, respectively.

Corollary 4.8. *Assuming q is coprime to the conductor, there exists an element t in $(\mathcal{O} : \mathcal{O}^\vee) \setminus \bigcup_i (\mathcal{O} : \mathcal{O}^\vee) \tilde{\mathfrak{p}}_i$ whose norm is bounded with high probability by*

$$\|t\| \leq \sqrt{q} \cdot \sqrt{n} \cdot q^{2\delta} \cdot N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n},$$

where $\delta \in \left[\frac{4n+\log \Delta_{\mathcal{O}}}{n \log q}, 1\right]$.

Proof. Let $\mathcal{O} = \mathbb{Z}[x]/(f) = \mathbb{Z}[\alpha]$, $\mathcal{I} = (\mathcal{O} : \mathcal{O}^\vee) = (\mathcal{O}^\vee)^{-1}$ and $\mathcal{J} = q\mathcal{O}$, where q is coprime to the conductor. As above, by Lemma 2.26, $N(\mathcal{I}) = N((\mathcal{O} : \mathcal{O}^\vee)) = N((\mathcal{O}^\vee)^{-1}) = N(f'(\alpha)\mathcal{O}) = N(f'(\alpha))$. Also, the associated primes of $q\mathcal{O}$ are exactly the unique primes of $q\mathcal{O}$ since q is coprime to the conductor. We sample t from a discrete Gaussian over $(\mathcal{O} : \mathcal{O}^\vee)$ of parameter

$$s = \begin{cases} \sqrt{q} \cdot N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n}, & \text{if } N(\tilde{\mathfrak{p}}_k) \geq q^{n(1/2+\delta)} \\ \sqrt{q} \cdot q^{2\delta} \cdot N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n}, & \text{else,} \end{cases}$$

where $\delta \in \left[\frac{4n+\log \Delta_{\mathcal{O}}}{n \log q}, 1 \right]$. By Proposition 4.7, t satisfies $t(\mathcal{O} : \mathcal{O}^\vee)^{-1} + q\mathcal{O} = \mathcal{O}$, with high probability. It is straight forward to check that $t \in (\mathcal{O} : \mathcal{O}^\vee) \setminus \bigcup_i (\mathcal{O} : \mathcal{O}^\vee) \tilde{\mathfrak{p}}_i$. Moreover, for $\epsilon = 2^{-2n}$, Lemma 2.22 implies that $\eta_\epsilon(\mathcal{I}) \leq N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n} \leq s$. Therefore, by Lemma 2.8, with high probability,

$$\|t\| \leq \sqrt{q} \cdot \sqrt{n} \cdot q^{2\delta} \cdot N(f'(\alpha))^{1/n} \cdot \Delta_{\mathcal{O}}^{1/n}.$$

□

The bounds from Corollaries 4.6 and 4.8 do have some common factors, namely $\Delta_{\mathcal{O}}^{1/n}$, $N(f'(\alpha))^{1/n}$ and \sqrt{n} . Therefore, it is enough to compare $\sqrt{q} \cdot q^{2\delta}$ and $M \cdot n \cdot \sqrt{\log n}$, respectively, where M is the hidden constant from Corollary 4.6 and $q = \text{poly}(n)$. An asymptotic comparison shows that Corollary 4.6 yields a better bound than Corollary 4.8. Also, recall that the latter result assumes q to be coprime to the conductor, whereas the earlier one is true for all rational primes q .

4.4 Comparison

Both Corollary 4.3 and Corollary 4.5 relate PLWE to worst-case ideal lattice problems. The former result involves invertible R -ideals, whereas the family of lattices in the latter is the set of invertible \mathcal{O} -ideals. These two families are disjoint, as any ideal can be invertible in at most a single order. In this regard, the two results are incomparable. We note that despite being disjoint, they are known to be related by the conductor ideal, see [Cona] for reference. We leave exploring this connection to future work.

Another parameter for comparison is the increase of the error in both hardness results. In the proposition below we show that the element t from Theorem 4.4 can be chosen to be smaller than the product $t_1 t_2^2$ from Theorem 4.2. Before doing so we give a more precise description of the elements t, t_1 and t_2 in the case where qR is coprime to the conductor.

Let $q \geq 2$ be some rational integer such that $qR + \mathcal{C}_{\mathcal{O}} = R$, and let $qR = \prod_{i=1}^k \mathfrak{p}_i^{e_i}$ be its factorization into prime ideals in R . Notice that also the primes \mathfrak{p}_i are coprime to the conductor. By Corollary 2.30, $q\mathcal{O} = qR \cap \mathcal{O}$ and coprime to the conductor as \mathcal{O} ideal. By Theorem 2.29 and Corollary 2.31, we get that $q\mathcal{O} = (\prod_{i=1}^k \mathfrak{p}_i^{e_i}) \cap \mathcal{O} = \prod_{i=1}^k (\mathfrak{p}_i \cap \mathcal{O})^{e_i}$ is the unique factorization of $q\mathcal{O}$ into prime ideals in \mathcal{O} . So the associated primes of $q\mathcal{O}$, $\tilde{\mathfrak{p}}_i$, are actually $\mathfrak{p}_i \cap \mathcal{O}$. In this setting, the elements t, t_1, t_2 are any elements satisfying the following conditions:

1. $t \in (\mathcal{O} : \mathcal{O}^\vee)$ and $t \notin (\mathfrak{p}_i \cap \mathcal{O})(\mathcal{O} : \mathcal{O}^\vee)$ for all $i \in [k]$.
2. $t_1 \in (R : R^\vee)$ and $t_1 \notin \mathfrak{p}_i(R : R^\vee)$ for all $i \in [k]$.
3. $t_2 \in \mathcal{C}_{\mathcal{O}}$ and $t_2 \notin \mathfrak{p}_i \mathcal{C}_{\mathcal{O}}$ for all $i \in [k]$. Notice that t_2^2 satisfies same properties.

Proposition 4.9. *With the same notations as above, for any t_1 and t_2 satisfying conditions 2 and 3, the product $t^* = t_1 t_2^2$ satisfies condition 1. In particular, letting t to be the shortest satisfying condition 1, and t_1 and t_2 satisfying conditions 2 and 3, respectively, such that $t_1 t_2^2$ is the shortest, we have that $\|t\|_\infty \leq \|t_1 t_2^2\|_\infty$.*

Proof. Recall that $\mathcal{C}_\mathcal{O} = (\mathcal{O} : R) = (R^\vee : \mathcal{O}^\vee)$, by Definition 2.28. Therefore

$$\mathcal{C}_\mathcal{O}^2(R : R^\vee) = (R^\vee : \mathcal{O}^\vee)(\mathcal{O} : R)(R : R^\vee) \subseteq (R^\vee : \mathcal{O}^\vee)(\mathcal{O} : R^\vee) \subseteq (\mathcal{O} : \mathcal{O}^\vee),$$

and so $t^* = t_1 t_2^2 \in (\mathcal{O} : \mathcal{O}^\vee)$.

It is now sufficient to prove that t^* is not in any $(\mathfrak{p}_i \cap \mathcal{O})(\mathcal{O} : \mathcal{O}^\vee)$, which would be equivalent to $t^* \mathcal{O}^\vee$ not being included in any $\mathfrak{p}_i \cap \mathcal{O}$. Assume there exists an $i \in [k]$ for which $t_1 t_2^2 \mathcal{O}^\vee \subseteq \mathfrak{p}_i \cap \mathcal{O}$. Then $t_1 t_2^2 \mathcal{O}^\vee R \subseteq (\mathfrak{p}_i \cap \mathcal{O})R = \mathfrak{p}_i$, according to Theorem 2.29. By Lemma 2.32, we get $t_1 R^\vee \cdot t_2^2 \mathcal{C}_\mathcal{O}^{-1} \subseteq \mathfrak{p}_i$. But since \mathfrak{p}_i is a prime ideal, it should follow that either $t_1 R^\vee \subseteq \mathfrak{p}_i$ or $t_2^2 \mathcal{C}_\mathcal{O}^{-1} \subseteq \mathfrak{p}_i$. This comes in contradiction with the choices of t_1 and t_2 .

This shows that $t^* = t_1 t_2^2$ belongs to $\mathcal{C}_\mathcal{O}^2(R : R^\vee) \setminus \bigcup_i (\mathfrak{p}_i \cap \mathcal{O})(\mathcal{O} : \mathcal{O}^\vee)$. By letting t to be the shortest satisfying condition 1 and t_1, t_2 satisfying conditions 2 and 3, respectively, such that t^* is the shortest, we have:

$$\begin{aligned} \|t\|_\infty &= \min\{\|x\|_\infty \mid x \in (\mathcal{O} : \mathcal{O}^\vee) \setminus \bigcup_i (\mathfrak{p}_i \cap \mathcal{O})(\mathcal{O} : \mathcal{O}^\vee)\} \\ &\leq \min\{\|x\|_\infty \mid x \in \mathcal{C}_\mathcal{O}^2(R : R^\vee) \setminus \bigcup_i (\mathfrak{p}_i \cap \mathcal{O})(\mathcal{O} : \mathcal{O}^\vee)\} \\ &\leq \|t^*\|_\infty \end{aligned}$$

□

5 Sampling Secrets from Orders

In this section, we consider a setting where the RLWE secret s is sampled from a subring of its designated space. For this purpose, it is more convenient to work with the dual version of R -LWE, which is used interchangeably in the literature but according to our notation should be denoted R^\vee -LWE. In this variant a is sampled from R^\vee and the secret s comes from R .

More formally, we assume the following setting. Let $q \geq 2$ be a rational prime that splits completely over R .¹² Then $R_q \simeq \mathbb{Z}_q^n$, as rings, and a subring $\bar{S} \subseteq R_q$ isomorphic to \mathbb{Z}_q^k corresponds to an order \mathcal{O} satisfying $qR \subseteq \mathcal{O} \subseteq R$. We show that this version of the Ring-LWE problem is at least as hard as the \mathcal{O}^\vee -LWE problem, defined in Section 3. In fact, this reduction follows as a corollary of a stronger result that shows that the \mathcal{O}^\vee -LWE problem becomes harder as the order becomes bigger. This result can be viewed as an analogue of [GHPS13, Lemma 3.1], for the \mathcal{O}^\vee -LWE problem, instead of the ring variant.

Given two orders $\mathcal{O}' \subseteq \mathcal{O}$, their duals satisfy $\mathcal{O}^\vee \subseteq \mathcal{O}'^\vee$, as fractional \mathcal{O}' -ideals, and there exist $\{v_1, v_2, \dots, v_m\} \subseteq \mathcal{O}'$ s.t. $\mathcal{O}^\vee = \sum_i \mathcal{O}'^\vee v_i$.¹³ We will be interested in finding such set with the smallest possible norm (that is, the ℓ_2 of the concatenation of the canonical embeddings of all v_i).

¹²A similar argument can be stated for the general case. However, this leads to a very cumbersome statement, and we prefer to avoid it.

¹³It is even possible to do so with $m = 2$, but we will be interested in v_i with small norm, in which case it is sometimes beneficial to use larger m .

Theorem 5.1. *Let $\mathcal{O}' \subseteq \mathcal{O} \subset K$ be orders, \mathcal{Q}' an integral \mathcal{O}' -ideal and \mathcal{Q} an integral \mathcal{O} -ideal such that $\mathcal{Q} = \mathcal{Q}'\mathcal{O}$. Let $\{v_1, v_2, \dots, v_m\} \subseteq \mathcal{O}'$ be s.t. $\mathcal{O}^\vee = \sum_i \mathcal{O}'^\vee v_i$. Let φ be a distribution over $\mathcal{O}'_{\mathcal{Q}'}$, let Υ be a family of distributions, each over $K_{\mathbb{R}}/\mathcal{O}'^\vee$, and let $u \in (\mathcal{O}' : \mathcal{Q}')$. Then there is a probabilistic polynomial time reduction from \mathcal{O}'^\vee -LWE $_{(\mathcal{Q}', u), \varphi, \Upsilon}$ to \mathcal{O}^\vee -LWE $_{(\mathcal{Q}, u), \varphi, \langle \Upsilon, \vec{v} \rangle}$, where $\langle \Upsilon, \vec{v} \rangle$ is the distribution (over distributions) that samples $\varphi \leftarrow \Upsilon$, and then outputs the distribution that e_1, \dots, e_m from φ and outputs $\sum_i e_i v_i$.*

Proof. We describe an efficient transformation that takes m elements from $\mathcal{O}'_{\mathcal{Q}'} \times K_{\mathbb{R}}/\mathcal{O}'^\vee$ and outputs an element in $\mathcal{O}_{\mathcal{Q}}^\vee \times K_{\mathbb{R}}/\mathcal{O}^\vee$. We show that this transformation maps uniform samples to uniform ones, and $\mathcal{O}'_{s, \psi, u}$ samples to $\mathcal{O}_{s, \langle \psi, \vec{v} \rangle, u}^\vee$ samples for any $s \leftarrow \varphi$ and $\psi \leftarrow \Upsilon$.

Given m samples $\{(a'_i, b'_i)\}_{i \in [m]}$, the transformation outputs $(a = \sum_i a'_i v_i, b = \sum_i b'_i v_i)$. Since $\mathcal{O}^\vee = \sum_i \mathcal{O}'^\vee v_i$ and $\mathcal{Q}\mathcal{O}^\vee = \sum_i \mathcal{Q}'\mathcal{O}'^\vee v_i$, this map is well-defined over the cosets that arise in the distributions and maps uniform distribution over $\mathcal{O}'_{\mathcal{Q}'} \times \mathbb{T}_{\mathcal{O}'^\vee}$ to uniform distribution over $\mathcal{O}_{\mathcal{Q}}^\vee \times \mathbb{T}_{\mathcal{O}^\vee}$, respectively.

Now, assume that $\{(a'_i, b'_i)\}_{i \in [m]}$ are sampled from $\mathcal{O}'_{s, \psi, u}$. Then, for $i \in [m]$, the element $b'_i = u \cdot a'_i \cdot s + e'_i$, where $e'_i \leftarrow \psi$. As

$$b = \sum_i b'_i v_i = u \cdot \sum_i a'_i v_i \cdot s + \sum_i e_i v_i = u \cdot a \cdot s + e,$$

where $e = \sum_i e_i v_i$ is sampled from $\langle \psi, \vec{v} \rangle$, so the tuple (a, b) lies in $\mathcal{O}_{s, \langle \psi, \vec{v} \rangle, u}^\vee$. This concludes the proof. \square

Corollary 5.2. *Let $\mathcal{O} \subset R$ be an order such that $qR \subseteq \mathcal{O}$, and let $\vec{v} = \{v_1, v_2, \dots, v_m\} \subset \mathcal{O}$ be short elements such that they generate R^\vee over \mathcal{O}^\vee , i.e., $R^\vee = \sum_i \mathcal{O}^\vee v_i$. Let Υ be a family of error distributions, each over $K_{\mathbb{R}}/\mathcal{O}^\vee$. Then, there exists a polynomial time reduction from \mathcal{O}^\vee -LWE $_{(qR, 1/q), \Upsilon}$ to R -LWE $_{q, U(\mathcal{O}/qR), \langle \Upsilon, \vec{v} \rangle}$.*

We note that in this case, the elements v_1, v_2, \dots, v_m are generators of the conductor ideal $\mathcal{C}_{\mathcal{O}}$ as an R -ideal.

Proof. The proof follows easily as a special case of Theorem 5.1; take $\mathcal{O}' = \mathcal{O}$ and $\mathcal{O} = R$, $\mathcal{Q}' = \mathcal{Q} = qR$, and $u = 1/q$. \square

Important Special Cases. We now discuss a family of orders \mathcal{O} that give rise to interesting secret distributions. Assume that q splits completely in R . Let $qR = \prod \mathfrak{p}_i$ denote the prime factorization of q in R . Then the Chinese remainder theorem yields the following isomorphism:

$$\begin{aligned} R_q &\xrightarrow{\sim} \prod_i \left(\frac{R}{\mathfrak{p}_i} \right) \simeq \mathbb{Z}_q^n \\ x &\mapsto (x \bmod \mathfrak{p}_i)_{i \in [n]}. \end{aligned}$$

Let $\Omega = (\Omega_1, \dots, \Omega_k)$ be a partition of $[n]$ into k disjoint subsets. Define

$$\overline{\mathcal{S}} := \{\mathbf{x} \in \mathbb{Z}_q^n \mid \mathbf{x}_j = \mathbf{x}_{j'}, \text{ for } j, j' \in \Omega_i, \text{ and } i \in [k]\}.$$

Then, the set $\overline{\mathcal{S}}$ is isomorphic to \mathbb{Z}_q^k and can be written as \mathcal{O}/qR , for an order \mathcal{O} such that $qR \subseteq \mathcal{O} \subseteq R$. Due to Corollary 5.2, we can get hardness of the Ring-LWE with the secret sampled from \mathcal{O}/qR from the hardness of \mathcal{O}^\vee -LWE and therefore, from the hardness of $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ -DGS.

In particular, if we consider K' a subfield of K , then its ring of integers $R' = \mathcal{O}_{K'}$ is a subring of R . Hence we can consider the following order $\mathcal{O} = R' + qR$ in R and see that \mathcal{O}/qR corresponds to some partition of $[n]$. Using the hardness result of Ring-LWE (in K') and the comparison result in [GHPS13, Lemma 3.1], one gets that the Ring-LWE problem (in K) with the secret sampled from $\mathcal{O} = R' + qR$ is at least as hard as $\mathfrak{I}(R')$ -DGS (in K').

On the other hand, using the hardness result of \mathcal{O}^\vee -LWE (Theorem 3.8) and Corollary 5.2, we get that the Ring-LWE problem (in K) with the secret sampled from \mathcal{O} is at least as hard as $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ -DGS (in K). One may wonder about a relation between the sets $\mathfrak{I}(R')$ and $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$. It is not too hard to check that the set of invertible ideals $\mathfrak{I}(R')$ embeds into $\mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee$ as follows:

$$\begin{aligned} \mathfrak{I}(R') &\hookrightarrow \mathfrak{I}(\mathcal{O}) \cdot \mathcal{O}^\vee \\ \mathcal{L}' &\mapsto \mathcal{L}'\mathcal{O} \cdot \mathcal{O}^\vee = (\mathcal{L}' + q\mathcal{L}'R) \cdot \mathcal{O}^\vee. \end{aligned}$$

6 RLWE Secrets From Ideals: High Entropy is Not Enough

In this section we show, perhaps surprisingly, that sampling Ring-LWE secrets from a high-entropy distribution is not necessarily sufficient to guarantee security. Specifically, we investigate the security of Ring-LWE in the case where the distribution of secrets is uniform over an *ideal*. We note that by definition this ideal must be a factor of the ideal qR (i.e. the ideal generated by the modulus q in the number field). In many applications of RLWE it is common to choose a value of q as a prime integer which nevertheless factors (splits completely) as an ideal over R .¹⁴ This means that elements in $R_q = R/qR$ can be represented using the Chinese Remainder Theorem as tuples of elements in $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$, and the factors of qR represent elements where some of the CRT coordinates are fixed to zero. Indeed, this CRT representation allows for more efficient operations over R_q and is the reason why such values of q are chosen in the first place. It is therefore natural to investigate whether setting a subset of the CRT coordinates to 0 has an effect on security.

We show that, as mentioned above, fixing a very small ϵ fraction of the CRT coordinates (thus only eliminating ϵ fraction of entropy) could result in complete loss of security. That is, we consider a RLWE instance with uniform secret, where worst-case to average-case reductions guarantee plausible security under the current state of the art in algorithms. We then show that by fixing any ϵ fraction of the CRT coordinates, the instance becomes insecure. The value of ϵ depends on the noise level of the RLWE instance. We complement this with a positive result, showing that taking ϵ that is slightly smaller than the aforementioned prescribed value is insufficient and worst-case hardness can still be established.

Notation. We use the standard RLWE setting where K is a number field of degree n with R as its ring of integers. We usually omit the asymptotic terminology to reduce clutter of notation.

Letting $\mathcal{P} \supset qR$ be an integral ideal in R , we let $\mathcal{Q} = q\mathcal{P}^{-1}$ denote its complement with respect to qR . We note that \mathcal{Q} is also an integral ideal in R . We further note that as per the above exposition, \mathcal{P} (or more accurately \mathcal{P}/qR) represents a subset of the CRT coordinates defined by the decomposition of qR in R . Since, formally RLWE is defined with secrets distributed over R^\vee/qR^\vee , therefore formally we will sample our secret from $\mathcal{P}R^\vee/qR^\vee$ rather than \mathcal{P}/qR itself. Note that the two spaces are isomorphic, due to the cancellation lemma (Lemma 2.35), and this distinction is

¹⁴Sometimes a product of such primes is used.

mere formalism. We would also like to point out that the dual of $\frac{1}{q}\mathcal{P}R^\vee$ is \mathcal{Q} . To see this observe that $\mathcal{Q}^\vee = (q\mathcal{P}^{-1})^\vee = \frac{1}{q}(\mathcal{P}^{-1})^\vee = \frac{1}{q}\mathcal{P}R^\vee$.

Remark 6.1. *As stated above, the results in this section capture secret distributions, or leakage scenarios, where a fixed subset of the CRT coordinates is known to be 0. We remark that all the results below generalize easily to the case where an ϵ fraction of the CRT coordinates is any fixed-value. As one would expect, fixing to some non-zero value corresponds to sampling the secret from a coset of an ideal.*

6.1 Insecure Instances

Our result is stated below.

Theorem 6.2. *Let K, R be a degree n number field and its ring of integers, respectively. Let $\mathcal{P} \supset qR$ be an integral R -ideal and $\mathcal{Q} = q\mathcal{P}^{-1}$ its complement as described above. There is a non-uniform algorithm such that for any distribution ψ satisfying $\Pr_{e \leftarrow \psi}[\|e\| < 1/(2\lambda_n(\mathcal{Q}))]$ is non-negligible and any distribution φ over $\mathcal{P}R^\vee/qR^\vee$, the algorithm solves search $R\text{-LWE}_{q,\varphi,\{\psi\}}$ with non-negligible probability given a single sample.*

We note that the theorem immediately implies that the same holds for $R\text{-LWE}_{q,\varphi,\Upsilon}$ where Υ is a distribution over distributions ψ so long as the probability to sample ψ as required in the theorem is non-negligible.

Proof. The algorithm will use a non-uniform advice string containing short vectors in \mathcal{Q} that will be used for decoding in the lattice \mathcal{P} . Specifically let $V = \{v_1, \dots, v_n\} \subset \mathcal{Q}$ be a set of \mathbb{Z} -linearly independent vectors satisfying $\|v_i\| \leq \lambda_n(\mathcal{Q})$.

The algorithm executes as follows. Given the input (a, b) , we let \bar{a} denote the inverse of a over R_q . This inverse exists with high probability, and is efficiently computable. It then considers b as an element in $K_{\mathbb{R}}$ by taking an arbitrary representative. It further applies Babai's BDD algorithm (Lemma 2.18) on input b with respect to the lattice $\frac{1}{q}\mathcal{P}R^\vee$, and with V as the decoding basis. The BDD subroutine returns an element b' in $\frac{1}{q}\mathcal{P}R^\vee$. Finally it returns $s' = q\bar{a}b' \pmod{qR^\vee} \in \mathcal{P}R^\vee/qR^\vee$.

We show that the algorithm succeeds whenever a is invertible and e satisfies $\|e\| < 1/(2\lambda_n(\mathcal{Q}))$. These conditions occur concurrently with non-negligible probability. We recall that $b = as/q + e \pmod{R^\vee}$, and note that $as/q \in \frac{1}{q}\mathcal{P}R^\vee/R^\vee$. Therefore when casting b as an element in $K_{\mathbb{R}}$ this element is of the form $y + e$ where $y \in \frac{1}{q}\mathcal{P}R^\vee$ and $y = as/q \pmod{R^\vee}$. We furthermore have that, for all i ,

$$|Tr(e \cdot v_i)| = \left| \langle \sigma(e), \overline{\sigma(v_i)} \rangle \right| \leq \|e\| \cdot \|v_i\| < 1/2.$$

Therefore, recalling that \mathcal{Q} is the dual of $\frac{1}{q}\mathcal{P}R^\vee$, we can apply Lemma 2.18 and deduce that the rounding algorithm recovers the value y .

Finally, the output value will be $s' = q\bar{a}y \pmod{qR^\vee} = q\bar{a}as/q \pmod{qR^\vee} = s \pmod{qR^\vee}$ and the result follows. \square

We exemplify the power of the theorem by considering a specific setting in Section 6.3.

6.2 Secure Instances

We now show that the vulnerability exposed in Theorem 6.2 can be mitigated by increasing the noise rate of the instance. Indeed, we show that sampling the secret from a distribution with lower entropy preserves worst-case hardness so long as the noise level is sufficiently high. To this end, we use our definition of Order-LWE (Definition 3.2), but with the order \mathcal{O} being the ring of integers R . This definition still generalizes the classical R -LWE since it allows us to consider a “modulus” which is not necessarily an integer q but an ideal \mathcal{Q} . In terms of terminology, \mathcal{O} -LWE with $\mathcal{O} = R$ will still be denoted R -LWE, so we overload the notation of the standard RLWE problem.

Theorem 6.3. *Let K, R be a degree n number field and its ring of integers respectively. Let $\mathcal{P} \supset qR$ be an integral R -ideal and $\mathcal{Q} = q\mathcal{P}^{-1}$ its complement as described above. Let $u \in \mathcal{Q}^{-1}$ and Υ be arbitrary. Then there is a polynomial time reduction from $R\text{-LWE}_{(\mathcal{Q}, u), \Upsilon}$ to $R\text{-LWE}_{q, U(\mathcal{P}R^\vee/qR^\vee), \Upsilon}$.*

Proof. We prove the theorem by showing a (randomized) transformation T that takes as input $a \in R_{\mathcal{Q}}$ and outputs $\tilde{a} = T(a) \in R_q$ such that

1. If a is uniform over its domain, then so is $T(a)$ over its domain.
2. For all $s \in R^\vee/\mathcal{Q}R^\vee$, there exists $\tilde{s} \in \mathcal{P}R^\vee/qR^\vee$ s.t. $uas = \tilde{a}\tilde{s}/q \pmod{R^\vee}$, for all $a \in R_{\mathcal{Q}}$.

If indeed such a transformation exists, then the reduction works as follows. Start by sampling s_0 uniformly from $\mathcal{P}R^\vee/qR^\vee$. Then, given a sequence of samples (a, b) for $R\text{-LWE}_{(\mathcal{Q}, u), \Upsilon}$, apply the transformation $(a, b) \rightarrow (\tilde{a}, \tilde{b}) = (T(a), b + \tilde{a}s_0/q)$ on each sample and output the resulting samples as $R\text{-LWE}_{q, U(\mathcal{P}R^\vee/qR^\vee), \Upsilon}$ samples. By the properties of the transformation indeed \tilde{a} is uniform, and $\tilde{b} = uas + e + \tilde{a}s_0/q = \tilde{a}(\tilde{s} + s_0)/q + e \pmod{R^\vee}$. Since s_0 is uniform over $\mathcal{P}R^\vee/qR^\vee$ then so is $(\tilde{s} + s_0)$ and indeed the output samples are distributed as required.

The transformation T is as follows. Given a as input, sample a random a' from \mathcal{Q}/qR and output $\tilde{a} = a + a' \pmod{qR}$. The first property holds since a is uniformly distributed over all cosets of a' . As for the second property, define $\tilde{s} = qus \pmod{qR^\vee}$. Since $u \in \mathcal{Q}^{-1}$, it holds that $qu \in \mathcal{P}$ and therefore indeed $\tilde{s} \in \mathcal{P}R^\vee/qR^\vee$. We have

$$\tilde{a}\tilde{s}/q = (a + a')qus/q = uas + ua's \pmod{R^\vee}.$$

Since $a' \in \mathcal{Q}/qR$, $u \in \mathcal{Q}^{-1}$ and $s \in R^\vee/\mathcal{Q}R^\vee$ we have that $ua's = 0 \pmod{R^\vee}$ and the result follows. \square

We exemplify the power of this theorem and its relation to Theorem 6.2 by considering a specific setting in Section 6.3.

6.3 A Threshold Phenomenon

Combining the results from Theorems 6.2 and 6.3, we show a (commonly used) setting where reducing the entropy of the secret results in tractability of the RLWE instance on one hand, but either using fully uniform secret (with the same noise level) or a modest increase in the noise level (with the same imperfect secret) results in the problem’s intractability being resumed.

We consider the setting where K is a cyclotomic number field (so that we have a good bound on the discriminant Δ_K) and where q is a prime for which the ideal qR splits completely as an ideal over R . The latter condition is the formal description of the fact that elements in R_q (and also in

R_q^\vee , due to Lemma 2.35) can be written in CRT form as tuples in \mathbb{Z}_q^n . We can thus consider secret distributions where k out of the n CRT coordinates are set to be 0, and the remaining $(n - k)$ coordinates are uniform. Naturally, this distribution has entropy $(1 - \frac{k}{n})n \log q$, i.e. $(1 - \frac{k}{n})$ -fraction of the full entropy. Formally, this corresponds to sampling the RLWE secret from an ideal $\mathcal{P} \supset qR$ with algebraic norm $N(\mathcal{P}) = q^k$. The formal statement is as follows.

Corollary 6.4. *Let K , R be a degree n number field and its ring of integers respectively, and assume furthermore that K is cyclotomic. Then for every integer $k \in [0, n]$, letting $\epsilon = k/n$, there exist $q = q_\epsilon = n^{O(1/\epsilon)}$, $\alpha = \alpha_\epsilon = \text{poly}(n)/q$ and a distribution φ over R_q^\vee with entropy $(1 - \epsilon)n \log q$ s.t. $R\text{-LWE}_{q,\varphi,\Upsilon_\alpha}$ is solvable in polynomial time.*

On the other hand, solving the problems $R\text{-LWE}_{q,\Upsilon_\alpha}$ and $R\text{-LWE}_{q,\varphi,\Upsilon_\beta}$ for any $\beta = \alpha \cdot \omega(n^{5/2})$ is as hard as solving $\mathfrak{I}(R)\text{-DGS}_\gamma$ for $\gamma = \eta \cdot \text{poly}(n^{1/\epsilon})$.

Solving DGS with γ as above corresponds to approximating the Shortest Independent Vector Problem (SIVP) to within $\text{poly}(n^{1/\epsilon})$ factor. At least for constant ϵ , achieving such DGS/SIVP approximation is intractable using current state of the art algorithmic techniques. Therefore, we show a threshold effect in two different aspects. First, in terms of entropy, we show a RLWE problem which is plausibly intractable if the secret is uniformly random, becomes tractable when the entropy is slightly reduced. Second, even if the entropy is reduced, a relatively modest increase in the noise level restores intractability.

Proof. Starting with k (and ϵ) as above, the values of q, α are set as follows. Let $p(n)$ be a polynomial, so that our choice of α satisfies $\alpha = p(n)/q$. We set $q \in (2p(n)n^{3/2})^{1/\epsilon} \cdot [1, 2] = n^{O(1/\epsilon)}$, and furthermore require that qR splits completely as an ideal over R . Setting $\alpha = p(n)/q$, it follows that

$$\alpha = p(n)/q = p(n)q^{-\epsilon}/q^{1-\epsilon} \leq 1/(2n^{3/2}q^{1-\epsilon}).$$

As explained in the outline above, once q is determined, we will consider an ideal $\mathcal{P} \supset qR$ of volume q^k . The distribution φ is simply $U(\mathcal{P}R^\vee/qR^\vee)$, and since $|\mathcal{P}R^\vee/qR^\vee| = |\mathcal{P}/qR| = |R/qR|/|R/\mathcal{P}| = q^n/N(\mathcal{P}) = q^{n-k}$, this has the stated entropy.

For the ideal \mathcal{Q} (the complement of \mathcal{P}) in the cyclotomic field K , Lemma 2.21 shows that $\lambda_1(\mathcal{Q}) = \lambda_n(\mathcal{Q}) \leq \sqrt{n}N(\mathcal{Q})^{1/n}\sqrt{\Delta_K}^{1/n}$. We have $N(\mathcal{Q}) = N(q\mathcal{P}^{-1}) = q^n/N(\mathcal{P}) = q^{n-k}$. In addition, since K is cyclotomic then $\Delta_K \leq n^n$. Therefore $\lambda_n(\mathcal{Q}) \leq nq^{1-\epsilon}$.

Given the value for α set above, we have $\alpha \leq 1/(2\sqrt{n}\lambda_n(\mathcal{Q}))$, which means that $\|e\| \leq 1/(2\lambda_n(\mathcal{Q}))$ with all but negligible probability, which allows us to apply Theorem 6.2 to conclude that $R\text{-LWE}_{q,\varphi,\Upsilon_\alpha}$ is solvable in polynomial time.

The intractability of $R\text{-LWE}_{q,\Upsilon_\alpha}$ follows directly from Theorem 2.40. To show the intractability of $R\text{-LWE}_{q,\varphi,\Upsilon_\beta}$, we first use Theorem 6.3 for a reduction from $R\text{-LWE}_{(\mathcal{Q},u),\Upsilon_\beta}$. We then apply Theorem 3.7 in the special case of $\mathcal{O} = R$, which is analyzed in Remark 3.10. We can apply the theorem if $\beta \geq \|u\|_\infty \cdot \omega(1)$, where by Lemma 2.21,

$$\|u\|_\infty = \lambda_1^\infty(\mathcal{Q}^{-1}) \leq N(\mathcal{Q}^{-1})^{1/n}\sqrt{\Delta_K}^{1/n} \leq n/q^{1-\epsilon}.$$

It is therefore sufficient to set $\beta = n/q^{1-\epsilon}\omega(1)$, which is indeed equal to $\alpha \cdot \omega(n^{5/2})$. As analyzed in Remark 3.10, this translates to γ as desired. \square

7 k -Wise Independent Secrets and Hidden Lattice BDD

In this section, we propose a class of high-entropy distributions for RLWE secrets for which we believe worst-case hardness should hold. As evidence, we show how to prove hardness based on a new average-case lattice problem. This new problem is a decision variant of the bounded distance decoding (BDD) problem. In our variant, BDD is to be solved on an ideal lattice which is sampled from a large family of ideals. It allows us to prove the hardness for distributions of RLWE secrets with norm bounded away from q and whose marginal distribution over this family of ideals (i.e. sampling an element from this distribution and taking its product with the ideal) is indistinguishable from uniform.

The Setting and Notation. In this section, we choose to use a simpler notation at the cost of some restriction on the generality of our discussion. This will allow us to present our results in a more digestible manner. In particular, we limit the discussion to cyclotomic number fields, our modulus to completely splitting, and the regime of the RLWE samples to be over $R_q \times R_q$, i.e. discrete and integral, instead of $R_q \times K_{\mathbb{R}}/R^{\vee}$.

Formally, we let K be a cyclotomic number field of degree n , and denote its ring of integers by R . In this case, the ideal R^{\vee} is just a scalar multiple of R , $R^{\vee} = tR$, for $t \in K$ [Conb, Theorem 3.7]. Therefore, we can assume that the RLWE distribution is obtained by sampling s from R_q instead of from R_q^{\vee} . Moreover, we consider that the error e from a discrete Gaussian over R . This setting is quite commonly used (and is perhaps the most popular use of RLWE) and its worst case hardness is presented in [LPR13, Lemma 2.23].

Defining k -Wise Independent Distributions. As explained above, we consider the case where q is an integer prime which splits completely, $qR = \prod_{i=1}^n \mathfrak{p}_i$, where each $\mathfrak{p}_i \subseteq R$ is a prime ideal. For $k \in [n]$, we define the following family of ideals

$$\mathfrak{P}_k := \left\{ \prod_{i \in T} \mathfrak{p}_i \mid T \subseteq [n], |T| = k \right\}.$$

Our class of (perfect/statistical/computational) k -wise independent distributions are those whose marginals are (perfectly/statistically/computationally) indistinguishable from uniform modulo any product of k prime ideals from q , so modulo any $\mathcal{P} \in \mathfrak{P}_k$. Recalling the CRT representation of R_q , this is equivalent to any k -tuple of CRT coordinates being indistinguishable from uniform.

Definition 7.1. *A distribution φ over R_q is (perfectly/statistically/computationally) k -wise independent if the random variables $(s \bmod \mathcal{P})$ and $(z \bmod \mathcal{P})$ are (perfectly/statistically/computationally) indistinguishable, where $s \leftarrow \varphi$, $z \leftarrow R_q$ and $\mathcal{P} \leftarrow \mathfrak{P}_k$. The asymptotics are over the dimension n and $k = k(n)$ is some integer function.*

Lemma 7.2. *Let φ be k -wise independent, and consider the following probability space. Sample $\mathcal{P} \leftarrow \mathfrak{P}_k$ and let $\mathcal{Q} = \mathcal{P}^{-1}q \in \mathfrak{P}_{n-k}$. Sample $x_1, x_2 \leftarrow \mathcal{Q}/qR$ conditioned on x_1 being invertible modulo \mathcal{P} , and $s \xleftarrow{\$} \varphi$. Then the distributions $(x_1, x_1 \cdot s)$ and (x_1, x_2) are indistinguishable.*

Proof. Let s' be any representative of $s \bmod \mathcal{P}$. Then $(x_1, x_1 \cdot s) = (x_1, x_1 \cdot s')$ since $x_1 \in \mathcal{Q}$ and $\mathcal{P} + \mathcal{Q} = qR$. Thus, Definition 7.1 implies that $(x_1, x_1 \cdot s)$ is indistinguishable from $(x_1, x_1 \cdot z)$ where z is uniform in R_q .

Now fix any \mathcal{P} , x_1 , we will show that $x_1 z$ and x_2 are identically distributed. Since x_1 is invertible modulo \mathcal{P} , then $x_1 z$ is uniform modulo \mathcal{P} . Since $x_1 \in \mathcal{Q}/qR$ it follows that $x_1 z = 0 \pmod{\mathcal{Q}}$. Therefore $x_1 z$ is uniform in \mathcal{Q}/qR . \square

7.1 Hidden-Lattice Decision Bounded Distance Decoding

We first define the hidden lattice BDD (HLBDD) distribution, and then the decisional problem associated with it. We use Gaussian noise but other noise distributions can be considered as well, the property that we use in our proof is that the distribution is *bounded*.

Definition 7.3 (Hidden Lattice BDD Distribution). *Let \mathcal{L}_1 and \mathcal{L}_2 be two given lattices, and let \mathcal{L} be a finite family of lattices, where each member $\mathcal{L}' \in \mathcal{L}$ satisfies $\mathcal{L}_1 \subseteq \mathcal{L}' \subseteq \mathcal{L}_2$. Let $\mathbf{r} \in G$ be a Gaussian parameter. The Hidden Lattice BDD Distribution over $\mathcal{L}_2/\mathcal{L}_1$, denoted by $C_{\mathcal{L}_1, \mathcal{L}, D_{\mathcal{L}_2}, \mathbf{r}}$, is sampled by choosing uniformly at random a lattice $\mathcal{L}' \xleftarrow{\$} \mathcal{L}$, an element $x \xleftarrow{\$} \mathcal{L}'/\mathcal{L}_1$, and an error term $e \leftarrow D_{\mathcal{L}_2, \mathbf{r}}$ and outputting $y = x + e \pmod{\mathcal{L}_1}$.*

One should think of the lattice \mathcal{L}_2 as the “ambient space”, i.e. \mathbb{Z}^n for general Euclidean setting or the ring of integer R in the algebraic setting. Note that it is possible to define the distribution with a continuous noise term e . The usual connection between discrete and continuous distribution from LWE / RLWE apply here as well (see, e.g., [LPR13, Lemma 2.23]).

Definition 7.4 (HLBDD Problem). *Let $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}, \mathbf{r}$ be as in Definition 7.3. The HLBDD Problem, denote by $\text{HLBDD}_{\mathcal{L}_1, \mathcal{L}, D_{\mathcal{L}_2}, \mathbf{r}}$ is to distinguish between two samples from the distribution $C_{\mathcal{L}_1, \mathcal{L}, D_{\mathcal{L}_2}, \mathbf{r}}$, and two samples from the uniform distribution over $\mathcal{L}_2/\mathcal{L}$.*

For the purpose of this section we will set $\mathcal{L}_2 = R$ (the ring of integers of our number field) and $\mathcal{L} = \mathfrak{P}_k$, for $k = n^{\Omega(1)}$.

Hardness and Variants. We defined HLBDD as the problem where the distinguisher only gets 2 samples from the HLBDD distribution. This is the minimal definition that is needed for our application. However, we note that we do not know of polynomial time algorithms even for weaker variants. For example, one where polynomially many samples are given to the distinguisher instead of only 2, or one where the distinguisher is provided with a (canonical) \mathbb{Z} -basis of the lattice \mathcal{P} in addition to the samples. We note that the latter variant is at least as easy as the former since it is possible to use a hybrid argument to show that if \mathcal{P} is known then indistinguishability for one sample implies indistinguishability for polynomially many samples. The connections to other problems in the literature, e.g. [HPS⁺14], is described in the introduction.

7.2 Stating and Proving Hardness

The reduction from HLBDD to RLWE with s from a k -wise distribution consists of two main steps.

1. A reduction from RLWE where the adversary gets only one RLWE sample, to the version with polynomially many samples. This reduction applies to any distribution of secrets which is bounded (and is the same on both the initial and final instances). The reduction assumes in addition the hardness of the standard RLWE problem (with the usual noise distribution).

The reduction follows using a rerandomization technique from [LPR13, Section 8.2], [BV11a, Lemma 4]. This transformation unfortunately also requires “noise swallowing”, a technique

that uses the fact that adding a Gaussian with super-polynomial Gaussian parameter will mask any random variable with polynomial amplitude.

2. A reduction from HLBDD to RLWE with a single sample. For this we assume that there is an adversary that can distinguish between a single RLWE sample $(a, b = as + e)$ and a uniform one.

We begin by replacing a with a decisional hidden-lattice BDD sample $(v_1 + e_1)$, where v_1 only has k nonzero CRT coordinates (randomly chosen) and e_1 is small. The decisional hidden-lattice BDD assumption asserts that this distribution will be indistinguishable from the original one. Namely, we now have $(v_1 + e_1, b = (v_1 + e_1)s + e)$. Opening the parenthesis, we have $b = v_1s + e_1s + e$.

We again use noise swallowing to argue that b is statistically close to $b = v_1s + e$, i.e. we use e to swallow e_1s , which can be done so long as s is small enough and e is large enough. Now we observe that since v_1 is zero on all but k CRT coordinates, and s is close to uniform in any subset of k coordinates, it follows that v_1s is statistically close to a fresh v_2 that is sampled from the same distribution as v_1 (i.e. has the same set of nonzero coordinates, but the value in each coordinate is randomly chosen). We get $b = v_2 + e$. We now apply decisional hidden-lattice BDD again to claim that $(a, b) = (v_1 + e_1, v_2 + e)$ is indistinguishable from uniform, which completes the proof.

We now state the main theorem of the section: the hardness of RLWE with k -wise independent secrets.

For $\sigma > 0$ and a Gaussian parameter $\mathbf{r} \in G$ such that $\mathbf{r} \geq \sigma$, recall from Lemma 2.11 \mathbf{r}_σ as defined by $r_{\sigma,i}^2 = r_i^2 - \sigma^2$.

Theorem 7.5. *Let $q \geq 2$ be a integer prime that splits completely over R . Let $k(n) \in [n]$. Let φ be a distribution over R_q that is B -bounded, i.e. $\Pr_{s \leftarrow \varphi}[\|s\| < B]$ is non-negligible. and k -wise independent as per Definition 7.1.*

Let $\sigma \geq \eta_\varepsilon(R)$ and $\mathbf{r}, \mathbf{r}', \mathbf{t}, \tilde{\mathbf{r}}, \mathbf{r}^ \in G$ be Gaussian parameters such that \mathbf{r}^* defined as $(\mathbf{r}_{\sigma,i}^*)^2 = r_{\sigma,i}$ and swallows \mathbf{t} and $\tilde{\mathbf{r}}$, \mathbf{r} swallows \mathbf{t} and $\tilde{\mathbf{r}}$ swallows \mathbf{r}' as in Definition 2.13.*

Then, there is a polynomial time reduction from $\text{HLBDD}_{qR, \mathfrak{P}_{n-k}, D_{R, \mathbf{r}'}}$ to $R\text{-LWE}_{q, \varphi, D_{R, \mathbf{r}}}$, assuming the hardness of $R\text{-LWE}_{q, \varphi, D_{R, \mathbf{t}}}$.

Proof. It follows from combining Lemma 7.6 and Lemma 7.7. □

Now we go to the first step of the theorem: we reduce from R -LWE with only one sample to R -LWE with polynomially many samples.

Lemma 7.6. *Let φ be a distribution over R_q that is B -bounded, i.e. $\Pr_{s \leftarrow \varphi}[\|s\| < B]$ is non-negligible. Let $\sigma \geq \eta_\varepsilon(R)$ and $\mathbf{r}, \tilde{\mathbf{r}}, \mathbf{t}, \mathbf{r}^* \in G$ be Gaussian parameters such that \mathbf{r}^* defined as $(\mathbf{r}_{\sigma,i}^*)^2 = r_{\sigma,i}$, such that it swallows \mathbf{t} and $\tilde{\mathbf{r}}$ and \mathbf{r} swallows \mathbf{t} as in Definition 2.13. Assuming the hardness of $R\text{-LWE}_{q, D_{R, \mathbf{t}}}$, there is a reduction from $R\text{-LWE}_{q, \varphi, D_{R, \tilde{\mathbf{r}}}}$ where the adversary gets only one sample, to $R\text{-LWE}_{q, \varphi, D_{R, \mathbf{r}}}$ with polynomially many samples.*

Proof. Given a single sample $(\tilde{a}, \tilde{b}) \in R_q \times R_q$, we do the following for each requested sample. Sample an element $z \leftarrow D_{R, \mathbf{t}}$, and errors $e' \leftarrow D_{R, \mathbf{t}}$ and $e'' \leftarrow D_{R, \mathbf{r}}$, and output

$$(a = \tilde{a}z + e', b = \tilde{b}z + e'') .$$

We claim that if (\tilde{a}, \tilde{b}) is distributed uniformly, then so is (a, b) . View (\tilde{a}, a) as a $R\text{-LWE}_{q, D_{R,t}, D_{R,t}}$ sample. Since the hardness of $R\text{-LWE}_{q, D_{R,t}}$ implies the hardness of $R\text{-LWE}_{q, D_{R,t}, D_{R,t}}$ (see [LPR13]), which in turn implies the hardness of $R\text{-LWE}_{q, D_{R,t}, D_{R,r}}$, it follows that a is indistinguishable from uniform. Similarly, considering (\tilde{b}, b) as a $R\text{-LWE}_{q, D_{R,t}, D_{R,r}}$, we get that b is uniform in R_q .

Now, let (\tilde{a}, \tilde{b}) be a single sample from the $R\text{-LWE}_{q, \varphi, D_{R,r}}$ distribution, so $\tilde{b} = \tilde{a}s + \tilde{e}$ for some $s \leftarrow \varphi$ and $\tilde{e} \leftarrow D_{R,\tilde{r}}$. The error term in this sample is as follows:

$$b - as = (\tilde{b}z + e'') - (\tilde{a}zs + e's) = (\tilde{a}sz + \tilde{e}z + e'') - (\tilde{a}zs + e's) = \tilde{e}z + e'' - e's =: e.$$

Let us analyze the distribution of $z\tilde{e} + e''$. Notice that with non-negligible probability, $\|z\tilde{e}\| \leq \|z\| \cdot \|\tilde{e}\| \leq \|\mathbf{t}\|_\infty \sqrt{n} \cdot \|\tilde{\mathbf{r}}\|_\infty \sqrt{n} = n \cdot \|\mathbf{t}\|_\infty \cdot \|\tilde{\mathbf{r}}\|_\infty$, by Lemma 2.8. Recall that \mathbf{r}^* swallows \mathbf{t} and $\tilde{\mathbf{r}}$, so by Definition 2.13, it follows that $n \cdot \|\mathbf{t}\|_\infty \cdot n \cdot \|\tilde{\mathbf{r}}\|_\infty / \min r_{\sigma,i} = \text{negl}(n)$. Hence $\sqrt{n} \cdot n \cdot \|\mathbf{t}\|_\infty \|\tilde{\mathbf{r}}\|_\infty / \min r_{\sigma,i} = \text{negl}(n)$. Also, since $z\tilde{e}$ is in R , according to Lemma 2.11, it follows that $z\tilde{e} + D_{R,\mathbf{r}}$ and $D_{R,\mathbf{r}}$ are statistically close. Therefore it follows that the distribution of $z\tilde{e} + e''$ is statistically close to the one of e'' .

Now we analyze the distribution of $e's + e''$. Notice that by Lemma 2.8, with non-negligible probability $\|e's\| \leq \|e'\| \cdot \|s\| \leq \|\mathbf{t}\|_\infty \sqrt{n} \cdot B$, since s is drawn from φ , a B -bounded distribution. Recall that \mathbf{r} swallows \mathbf{t} , so by Definition 2.13 it follows that $B \cdot n \cdot \|\mathbf{t}\|_\infty / \min r_{\sigma,i} = \text{negl}(n)$. Since $e's \in R$, by Lemma 2.11, it follows that the distribution of $e's + e''$ is statistically close to that of e'' . We conclude that e is distributed statistically close to $D_{R,\mathbf{r}}$, and therefore (a, b) is distributed as in the distribution of $R\text{-LWE}_{q, \varphi, D_{R,\mathbf{r}}}$. \square

Now we prove the second step of the theorem: we show a reduction from HLBDD to $R\text{-LWE}$ with only one sample.

Lemma 7.7. *Let φ be a distribution over R_q that is B -bounded, i.e. $\Pr_{s \leftarrow \varphi}[\|s\| < B]$ is non-negligible, and k -wise independent. Let $\sigma \geq \eta_\varepsilon(R)$ and $\tilde{\mathbf{r}}, \tilde{\mathbf{r}}, \mathbf{r}' \in G$ be Gaussian parameters such that $\tilde{\mathbf{r}}$ is as in Lemma 2.11, $\tilde{\mathbf{r}} \geq \max\{\mathbf{r}', \sigma\}$ and $\tilde{\mathbf{r}}$ swallows \mathbf{r}' . Then, there is a reduction from $\text{HLBDD}_{qR, \mathfrak{P}_{n-k}, D_{R,\mathbf{r}'}}$ to $R\text{-LWE}_{q, \varphi, D_{R,\tilde{\mathbf{r}}}}$ with a single sample.*

Proof. We show that under the $\text{HLBDD}_{qR, \mathfrak{P}_{n-k}, D_{R,\mathbf{r}'}}$ assumption, a sample from $R\text{-LWE}_{q, \varphi, D_{R,\tilde{\mathbf{r}}}}$ is indistinguishable from a sample from $U(R_q \times R_q)$. We prove this using a hybrid argument. The hybrids used are defined here and followed by a detailed description.

$$\begin{aligned} \mathcal{H}_0 &: (a, b = as + e) \in R_q \times R_q \text{ where } s \leftarrow \varphi, e \leftarrow D_{R,\tilde{\mathbf{r}}} \\ \mathcal{H}_1 &: (x_1 + e_1, (x_1 + e_1)s + e) \in (\mathcal{Q} + R) \bmod q \times R_q \text{ where } \mathcal{Q} \xleftarrow{\$} \mathfrak{P}_{n-k}, e_1 \leftarrow D_{R,\mathbf{r}'} \\ \mathcal{H}_2 &: (x_1 + e_1, x_1s + e_2) \in R_q \times R_q, \text{ where } e_2 \leftarrow D_{R,\tilde{\mathbf{r}}} \\ \mathcal{H}_3 &: (x_1 + e_1, x_2 + e_2) \in R_q \times R_q, \text{ where } x_2 \leftarrow \mathcal{Q}/qR \\ \mathcal{H}_4 &: U(R_q \times R_q) \end{aligned}$$

Hybrid \mathcal{H}_0 . In this hybrid, we consider the distribution $A_{s, D_{R,\tilde{\mathbf{r}}}}$, where $s \leftarrow \varphi$.

Hybrid \mathcal{H}_1 . Here, instead of sampling $a \xleftarrow{\$} R_q$, we do the following;

1. sample an ideal $\mathcal{Q} \xleftarrow{\$} \mathfrak{P}_{n-k}$,
2. sample an element $x_1 \xleftarrow{\$} \mathcal{Q}/qR$ and an error term $e_1 \leftarrow D_{R,\mathbf{r}'}$, and

3. output $a := x_1 + e_1 \pmod{qR}$.

Thus, this hybrid has samples of the form

$$(a = x_1 + e_1 \pmod{qR}, b = x_1 \cdot s + e_1 \cdot s + e \pmod{qR}) .$$

Assuming the hardness of $\text{HLBDD}_{qR, \mathfrak{P}_{n-k}, D_{R, \mathbf{r}'}}$, we get that $\mathcal{H}_0 \approx \mathcal{H}_1$.

Hybrid \mathcal{H}_2 . Here, we replace the term $e_1 \cdot s + e$ with e_2 , where $e_2 \xleftarrow{\$} D_{R, \tilde{\mathbf{r}}}$. This yields,

$$(a = x_1 + e_1 \pmod{qR}, b = x_1 \cdot s + e_2 \pmod{qR}) .$$

Recall that φ is B -bounded. Let us check that $e_1 \cdot s + e$ is statistically close to e_2 . Indeed, by Lemma 2.8, with non-negligible probability $\|e_1 s\| \leq \|e_1\| \cdot \|s\| \leq \|\mathbf{r}'\|_\infty \cdot \sqrt{n} \cdot B$. Recall that $\tilde{\mathbf{r}}$ swallows \mathbf{r}' , so by Definition 2.13, it follows that $B \cdot n \cdot \|\mathbf{r}'\|_\infty / \min \tilde{r}_{\sigma, i} = \text{negl}(n)$. Since $e_1 s \in R$, by Lemma 2.11 we get that $e_1 \cdot s + e$ is distributed statistically close to the distribution of e_2 . We conclude that the hybrids \mathcal{H}_1 and \mathcal{H}_2 are statistically close.

Hybrid \mathcal{H}_3 . We note that since q is super-polynomial, x_1 is invertible with all but negligible probability. Therefore, Lemma 7.2 implies that the hybrids \mathcal{H}_2 and \mathcal{H}_3 are indistinguishable.

Hybrid \mathcal{H}_4 . Observe that the resulting distribution from the previous hybrid is

$$(a = x_1 + e_1 \pmod{qR}, b = x_2 + e_2 \pmod{qR}) ,$$

where $x_1, x_2 \xleftarrow{\$} \mathcal{Q}/qR$, $e_1 \leftarrow D_{R, \mathbf{r}'}$ and $e_2 \leftarrow D_{R, \tilde{\mathbf{r}}}$. In this hybrid, we sample a and b uniformly at random. Assuming the hardness of $\text{HLBDD}_{qR, \mathfrak{P}_{n-k}, D_{R, \mathbf{r}'}}$, which implies the hardness of $\text{HLBDD}_{qR, \mathfrak{P}_{n-k}, D_{R, \tilde{\mathbf{r}}}}$ as long as $\tilde{\mathbf{r}} \geq \mathbf{r}'$, we replace $x_1 + e_1 \pmod{qR}$ and $x_2 + e_2$ with elements from $U(R_q)$. Thus $\mathcal{H}_3 \approx \mathcal{H}_4$.

Putting it all together, we get that $A_{s, D_{R, \tilde{\mathbf{r}}}} = \mathcal{H}_0 \approx \mathcal{H}_4 = U(R_q \times R_q)$. \square

Acknowledgments

We wish to thank the anonymous referees for their insightful comments and useful suggestions.

References

- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.
- [ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, pages 327–343. USENIX Association, 2016.

- [AGHS13] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete gaussian leftover hash lemma over infinite domains. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, volume 8269 of *Lecture Notes in Computer Science*, pages 97–116. Springer, 2013.
- [AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *Proceedings of the 6th Theory of Cryptography Conference*, pages 474–495, 2009.
- [AP13] Jacob Alperin-Sheriff and Chris Peikert. Practical bootstrapping in quasilinear time. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2013.
- [Bab86] László Babai. On lovász lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [BF11] Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 149–168. Springer, 2011.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 533–556, 2014.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS*, pages 309–325. ACM, 2012. Invited to ACM Transactions on Computation Theory.
- [BKLP15] Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl, editors, *Computer Security - ESORICS 2015 - 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21-25, 2015, Proceedings, Part I*, volume 9326 of *Lecture Notes in Computer Science*, pages 305–325. Springer, 2015.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584. ACM, 2013.

- [BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *Advances in Cryptology—EUROCRYPT 2012*, pages 719–737. Springer, 2012.
- [BV11a] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In Phillip Rogaway, editor, *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 501–521. Springer, 2011.
- [BV11b] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *FOCS*, pages 97–106. IEEE, 2011.
- [BVWW16] Zvika Brakerski, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Obfuscating conjunctions under entropic ring LWE. In Madhu Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 147–156. ACM, 2016.
- [Cona] Keith Conrad. The conductor ideal. Expository papers/Lecture notes. Available at: <http://www.math.uconn.edu/~kconrad/blurbs/gradnumthy/conductor.pdf>.
- [Conb] Keith Conrad. The different ideal. Expository papers/Lecture notes. Available at: <http://www.math.uconn.edu/~kconrad/blurbs/gradnumthy/different.pdf>.
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 40–56. Springer, 2013.
- [DGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In *EUROCRYPT*, pages 24–43, 2010.
- [DGK⁺10] Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Public-key encryption schemes with auxiliary inputs. In Daniele Micciancio, editor, *TCC*, volume 5978 of *Lecture Notes in Computer Science*, pages 361–381. Springer, 2010.
- [GHPS13] Craig Gentry, Shai Halevi, Chris Peikert, and Nigel P Smart. Field switching in bgv-style homomorphic encryption. *Journal of Computer Security*, 21(5):663–684, 2013.
- [GHS12] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the AES circuit. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 850–867. Springer, 2012.
- [GKPV10] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In Andrew Chi-Chih Yao, editor, *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 230–240. Tsinghua University Press, 2010.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *Proceedings of the*

- 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 545–554. ACM, 2013.
 - [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS98*, pages 267–288, 1998.
 - [HPS⁺14] Jeff Hoffstein, Jill Pipher, John M Schanck, Joseph H Silverman, and William Whyte. Practical signatures from the partial fourier recovery problem. In *International Conference on Applied Cryptography and Network Security*, pages 476–493. Springer, 2014.
 - [HS14] Shai Halevi and Victor Shoup. Algorithms in helib. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 554–571. Springer, 2014.
 - [HS15] Shai Halevi and Victor Shoup. Bootstrapping for helib. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 641–670. Springer, 2015.
 - [Lan14] Adeline Langlois. *Lattice-Based Cryptography: Security Foundations and Constructions*. PhD thesis, ENS Lyon, 2014.
 - [LD18] Tommy Reddad Luc Devroye, Abbas Mehrabian. The total variation distance between high-dimensional gaussians. Available at: <https://arxiv.org/abs/1810.08693>, 2018.
 - [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155. Springer, 2006.
 - [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.
 - [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-lwe cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 35–54. Springer, 2013.

- [MP13] Daniele Micciancio and Chris Peikert. Hardness of sis and lwe with small parameters. In Ran Canetti and Juan A. Garay, editors, *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2013.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *STOC*, pages 333–342. ACM, 2009.
- [Pei10] Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *Annual Cryptology Conference*, pages 80–97. Springer, 2010.
- [Pei16] Chris Peikert. How (not) to instantiate ring-lwe. In Vassilis Zikas and Roberto De Prisco, editors, *Security and Cryptography for Networks - 10th International Conference, SCN 2016, Amalfi, Italy, August 31 - September 2, 2016, Proceedings*, volume 9841 of *Lecture Notes in Computer Science*, pages 411–430. Springer, 2016.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006.
- [PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 478–487. ACM, 2007.
- [PRSD17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. *IACR Cryptology ePrint Archive*, 2017:258, 2017.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005. Full version in [Reg09].
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-lwe and polynomial-lwe problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 146–173. Springer, 2018.
- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of*

Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635. Springer, 2009.

- [Ste08] Peter Stevenhagen. The arithmetic of number rings. *Algorithmic number theory: lattices, number fields, curves and cryptography*, 44:209–266, 2008.
- [Was83] Lawrence C. Washington. *Introduction to Cyclotomic Fields*. 1983.