Pseudorandomness of Ring-LWE for Any Ring and Modulus

Chris Peikert* Oded Regev† Noah Stephens-Davidowitz‡

March 21, 2017

Abstract

We give a polynomial-time quantum reduction from worst-case (ideal) lattice problems directly to the *decision* version of (Ring-)LWE. This extends to decision all the worst-case hardness results that were previously known for the search version, for the same or even better parameters and with no algebraic restrictions on the modulus or number field. Indeed, our reduction is the first that works for decision Ring-LWE with *any number field* and *any modulus*.

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

†Courant Institute of Mathematical Sciences, New York University. Supported by the Simons Collaboration on Algorithms and Geometry and by the National Science Foundation (NSF) under Grant No. CCF-1320188.

‡Courant Institute of Mathematical Sciences, New York University. Email: noahsd@gmail.com. Supported by the National Science Foundation (NSF) under Grant No. CCF-1320188, and the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-15-C-0236. Part of this work was done while visiting the first author at the University of Michigan.

1 Introduction

In recent years, *Learning With Errors* (LWE) [Reg05] and its more compact variant Ring-LWE [LPR10] have served as foundations for a wide variety of lattice-based cryptographic constructions (e.g., [PW08, GPV08, SS11, BV11, BGV12, GVW13, GSW13, Pei14, ADPS16, BCD+16]). Informally, for a dimension n and integer modulus q, LWE is concerned with "noisy" linear equations $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle / q + e_i$ for a secret $\mathbf{s} \in \mathbb{Z}_q^n$ and public vectors $\mathbf{a}_i \in \mathbb{Z}_q^n$, all uniformly random and independent. The error terms e_i are drawn from some known distribution, usually a Gaussian of width $\alpha \ll 1$, called the *error rate*. Similarly, Ring-LWE deals with noisy ring products $b_i = a_i \cdot s/q + e_i$ for a secret $s \in R^\vee/qR^\vee$ and public $a_i \in R/qR$, where R is usually taken to be the ring of algebraic integers in a number field K, and $R^\vee \subset K$ is the "dual" ideal of R. Here the error terms e_i are drawn from an appropriate n-dimensional (usually Gaussian) distribution, with an analogous notion of error rate.

The average-case *search* (Ring-)LWE problem is to find the secret, given many noisy equations as described above. For cryptographic purposes, however, one usually relies on hardness of the *decision* problem, which is to distinguish such equations from uniformly random ones, i.e., where the b_i are uniform and independent of everything else. In other words, one wants (Ring-)LWE samples to be pseudorandom. A main attraction of these problems is their provable hardness assuming the intractability of *worst-case* lattice problems. Ring-LWE is particularly attractive for its efficiency and compactness, but many of its instantiations still lack such hardness theorems, as we now detail.

Hardness of Ring-LWE. There is a large disparity between known hardness theorems for search and decision Ring-LWE. Following the original template of [Reg05] for plain LWE, prior work [LPR10] gives a sequence of two main reductions. The first is a (quantum) reduction from worst-case problems like the approximate Shortest Independent Vectors Problem (SIVP) on *ideal* lattices in a given ring to search Ring-LWE over that same ring [LPR10, Section 4].[1] The second is a (classical) reduction from search to decision for Ring-LWE [LPR10, Section 5]. However, the first reduction is much more general: it applies to any number field and any modulus that is not too small (relative to the error rate), whereas the search-to-decision reduction works only for prime moduli that "split well" and *cyclotomic* number fields, or more generally, Galois fields (see [EHL14]).[2] The splitting condition can be waived via an additional "modulus switching" reduction [LS15] (which builds on a technique from [Pei09, BV11, BLP+13]), at the cost of somewhat increasing the error rate. However, the restriction to Galois number fields—a rather narrow class—seems unavoidable using prior techniques.

There are a few reasons why one might wish to use Ring-LWE over non-Galois number fields. For one, the algebraic structure of Galois fields, or cyclotomics in particular, might conceivably be used to attack worst-case ideal-lattice problems like approximate SIVP, or Ring-LWE itself. Most cryptographic constructions based on Ring-LWE use inverse-polynomial (or inverse-quasipolynomial) error rates, which correspond to polynomial worst-case approximation factors for ideal-lattice problems. For such parameters, no known algorithm significantly outperforms those for plain LWE or SIVP on general lattices. However, a series of recent works [CGS14, BS16, CDPR16, CDW17] has yielded polynomial-time quantum algorithms that (under plausible heuristics) obtain *subexponential* $\exp(\widetilde{O}(\sqrt{n}))$ approximation factors for SIVP in prime-power cyclotomics. These algorithms are not known to apply to any nontrivial instantiations of Ring-LWE,

[1] An interesting alternative approach to deriving hardness of search Ring-LWE was given in [SSTX09], independently of [LPR10]. However, as explained in [LPR10, Section 1.4], this approach leads to weaker hardness results.

[2] An algebraic field extension K/L of degree n is Galois if there are n automorphisms of K that fix L pointwise.

however. (Indeed, it is unknown whether Ring-LWE reduces to even the exact version of SIVP on ideal lattices in the same ring.)

A second reason to consider alternative rings is that there are families of number fields for which the worst-case reduction to *search* Ring-LWE delivers better approximation factors than for cyclotomics [Roq67, PR07, LPR10]—as small as $\omega(\sqrt{\log n})$, versus small polynomial factors—which suggests stronger hardness for cryptographic applications. However, *decision* Ring-LWE for such fields was not previously supported by any hardness theorems.

Hardness of plain LWE. There is a less-pronounced gap between the known hardness of search and decision for plain LWE, but the state of the art is still unsatisfactory. For search, there is a quantum reduction (from worst-case SIVP on general lattices) that works for any large enough modulus, and obtains the best known parameters [Reg05]. For decision, however, there are various specialized and incomparable reductions [Reg05, Pei09, ACPS09, MM11, MP12, BLP$^+$13], which either impose some number-theoretic constraints on the modulus, or incur some significant loss in the LWE parameters. For example, the modulus-switching reduction from [BLP$^+$13] increases the error rate, which ultimately yields a weaker worst-case approximation factor than for the corresponding search LWE problem.

1.1 Contributions

Our main contribution is a polynomial-time quantum reduction from worst-case (ideal) lattice problems directly to *decision* (Ring-)LWE. This yields a conceptually simpler hardness proof for (Ring-)LWE, and avoids the need for search-to-decision reductions. More specifically, we extend to decision all worst-case hardness results that were previously known for search, for the same or even better parameters and with no algebraic restrictions on the modulus or number field. In particular, our reduction works for *any modulus* in the plain LWE setting (Theorem 5.1), and for any modulus and *any number field* in the ring setting (Theorem 6.2). Our results also appear to be sufficiently general to adapt to "module" lattices and LWE [BGV12, LS15].[3] Finally, our techniques apply entirely to the *classical* (non-quantum) part of the iterative quantum reduction from [Reg05], and can therefore also be applied to the alternative classical reduction for LWE from [Pei09].

Our second contribution, which may be of independent interest, is a random self-reduction for bounded-distance decoding (BDD) with Gaussian error. (See Section 3.) This ultimately yields a $\widetilde{\Theta}(n^{1/4})$-factor improvement in the Ring-LWE error size versus the result of [LPR10, Section 5] (for non-spherical error).

1.2 Which Number Fields To Use?

Our reduction says that decision Ring-LWE with any large enough modulus (and appropriate error distribution) is supported by a worst-case hardness guarantee, where the approximation factor is determined solely by the error rate. This applies to any *fixed* choice of number field, but it gives no guidance about which number fields are preferable for security. Apart from trivial reductions (e.g., from a number field to an extension), it is an open question whether there is any formal connection between different number fields for worst-case ideal lattice problems (like SIVP) or Ring-LWE. It is also unclear what properties of number fields may affect the complexity of such problems.

We do note, however, that the geometry of the dual ideal $R^\vee$ affects the error rates that can be usefully employed in cryptographic applications. These typically need error of rate α to be decodable relative to the

[3] Module-LWE interpolates between the plain and ring variants, providing a smooth tradeoff between the efficiency and compactness of Ring-LWE and the potential additional security of plain LWE.

dual ideal $R^\vee$, which is usually ensured by taking α to be somewhat smaller than the minimum distance $\lambda_1(R^\vee)$. As α decreases, worst-case hardness theorems give weaker guarantees (i.e., larger approximation factors), and known attacks on Ring-LWE become more efficient. Therefore, a smaller $\lambda_1(R^\vee)$ corresponds to worse parameters and thereby appears to provide less security. (A similar phenomenon arises for rings having large "expansion factors." See, e.g., [LM06, Gen09].) Cyclotomic fields have relatively large $\lambda_1(R^\vee)$ (and correspondingly small expansion factors), which is one of the reasons that they are often used in applications.

1.3 Techniques

Here we describe the main new techniques used in our reduction. (See Section 2 for the relevant definitions.) Except where noted, we focus our discussion on plain LWE and general lattices.

Staring at the black box. A part of Regev's original reduction for LWE [Reg05] transforms an instance of Bounded Distance Decoding (BDD, the problem of finding the closest lattice vector to a target $\mathbf{t}$ that is guaranteed to be rather close to a lattice $\mathcal{L}$) into LWE samples whose secret corresponds to the closest lattice vector $\mathbf{v} \in \mathcal{L}$ to $\mathbf{x}$, and whose error rate α is (proportional to) $\mathrm{dist}(\mathbf{t}, \mathcal{L}) = \|\mathbf{t} - \mathbf{v}\|$ (see Lemma 5.3). Regev showed that, as long as $\mathrm{dist}(\mathbf{t}, \mathcal{L})$ is small enough, a suitable oracle for *search* LWE will find the secret, which allows for the recovery of $\mathbf{v}$.

In contrast, our reduction uses an oracle for *decision* LWE for a particular error rate $\widehat{\alpha}$, and works by incrementally moving $\mathbf{t}$ towards the closest lattice vector $\mathbf{v}$ by carefully measuring the behavior of the oracle. To accomplish this, we consider $p(\alpha)$, the probability that the oracle accepts when given LWE samples with error rate α. Notice that we can closely approximate $p(\mathrm{dist}(\mathbf{t}, \mathcal{L}))$ for any $\mathbf{t} \in \mathbb{R}^n$ by repeatedly invoking the oracle on LWE samples generated from $\mathbf{t}$ and measuring the oracle's acceptance probability. Our goal is to use this to detect when a point $\mathbf{t}'$ is "significantly closer" to the lattice than $\mathbf{t}$ is (more precisely, $\mathrm{dist}(\mathbf{t}', \mathcal{L}) \leq (1 - 1/\operatorname{poly}(n)) \mathrm{dist}(\mathbf{t}, \mathcal{L})$). This allows us to solve BDD by repeatedly perturbing $\mathbf{t}$ to a new point $\mathbf{t}'$, testing whether the new point is significantly closer to the lattice, and if so setting $\mathbf{t} = \mathbf{t}'$. (A similar idea is used in [LLM06], in a different context.)

To see how we can use the oracle to detect when we have moved closer to the lattice, suppose for the moment that its acceptance probability $p(\alpha)$ is a monotonically decreasing function of the error rate α, and assume also that it decreases noticeably around $\alpha = \mathrm{dist}(\mathbf{t}, \mathcal{L})$. In this case, $p(\mathrm{dist}(\mathbf{t}', \mathcal{L}))$ will be noticeably larger than $p(\mathrm{dist}(\mathbf{t}, \mathcal{L}))$ if and only if $\mathbf{t}'$ is significantly closer to the lattice than $\mathbf{t}$ is, and we can easily detect this by approximating $p(\mathrm{dist}(\mathbf{t}', \mathcal{L}))$ and $p(\mathrm{dist}(\mathbf{t}, \mathcal{L}))$.

Looking for change in all the wrong places. Let us now drop the assumption that p decreases noticeably around $\alpha = \mathrm{dist}(\mathbf{t}, \mathcal{L})$, but still assume that it is monotonically decreasing. We now need to deal with the possibility that $p(\alpha)$ may be nearly constant for all $\alpha \approx \mathrm{dist}(\mathbf{t}, \mathcal{L})$. In this case, $|p(\mathrm{dist}(\mathbf{t}', \mathcal{L})) - p(\mathrm{dist}(\mathbf{t}, \mathcal{L}))|$ will be negligible for any small perturbation $\mathbf{t}'$ of $\mathbf{t}$. As a result, we cannot hope to tell whether we have moved closer to the lattice by examining $p(\mathrm{dist}(\mathbf{t}', \mathcal{L}))$ and $p(\mathrm{dist}(\mathbf{t}, \mathcal{L}))$ alone.

In order to overcome this, recall that by hypothesis, for a particular error rate $\widehat{\alpha}$ the difference between $p(\widehat{\alpha})$ and, say, $p(\sqrt{n})$ is noticeable, because LWE samples with error rate $\sqrt{n}$ are essentially uniform. Therefore, even though $p(\alpha)$ could be constant for $\alpha \in [0, \widehat{\alpha}]$, it must decrease noticeably somewhere beyond $\widehat{\alpha}$. As a first attempt to exploit this property, we can add some extra error to the LWE samples obtained from $\mathbf{t}$ and $\mathbf{t}'$. More precisely, if $(\mathbf{a}_i, b_i)$ are LWE samples with error rate α, then adding independent Gaussians of width β to the b_i yields samples with error rate $\sqrt{\alpha^2 + \beta^2}$. This lets us approximate $p(\sqrt{\mathrm{dist}(\mathbf{t}, \mathcal{L})^2 + \beta^2})$ for any $\mathbf{t}$ and $\beta \geq 0$, which is quite useful. Indeed, we can use this to search for a region where p decreases noticeably:

given $\mathbf{t}$ and a small perturbation $\mathbf{t}'$, we can approximate $p(\sqrt{\mathrm{dist}(\mathbf{t},\mathcal{L})^2+\beta^2})$ and $p(\sqrt{\mathrm{dist}(\mathbf{t}',\mathcal{L})^2+\beta^2})$ for various values of β. If $\mathrm{dist}(\mathbf{t},\mathcal{L}) - \mathrm{dist}(\mathbf{t}',\mathcal{L})$ is sufficiently large, then we will recognize this when we use a value of β for which p decreases noticeably around parameter $\sqrt{\mathrm{dist}(\mathbf{t},\mathcal{L})^2+\beta^2}$.

Be fruitful and multiply. Unfortunately, the above technique stops working once $\mathbf{t}$ becomes too close to the lattice. To see this, suppose that $p(\alpha)$ is constant for all $\alpha \le \widehat{\alpha}$, and that $\mathrm{dist}(\mathbf{t},\mathcal{L}) \ll \widehat{\alpha}$. Then we may not be able to distinguish between $p(\sqrt{\mathrm{dist}(\mathbf{t},\mathcal{L})^2+\beta^2})$ and $p(\sqrt{\mathrm{dist}(\mathbf{t}',\mathcal{L})^2+\beta^2})$ for any small perturbation $\mathbf{t}'$ of $\mathbf{t}$ and any $\beta \ge 0$. For small $\beta \lesssim \widehat{\alpha}$ the two values are nearly identical by assumption, and for large $\beta \gtrsim \widehat{\alpha}$ the extra error of width β "drowns out" the original error of width $\mathrm{dist}(\mathbf{t},\mathcal{L})$ or $\mathrm{dist}(\mathbf{t}',\mathcal{L})$.

Although it is possible to salvage the technique, we choose to follow a different path, mainly because we know how to extend it to the ring setting. Instead of increasing the error *additively* and working with $p(\sqrt{\mathrm{dist}(\mathbf{t},\mathcal{L})^2+\beta^2})$, we increase the error *multiplicatively* and work with the function $p(r\,\mathrm{dist}(\mathbf{t},\mathcal{L}))$ for $r \ge 1$, i.e., we approximate $p(r\,\mathrm{dist}(\mathbf{t},\mathcal{L}))$ and $p(r\,\mathrm{dist}(\mathbf{t}',\mathcal{L}))$ for many different values of r. As long as $\mathrm{dist}(\mathbf{t},\mathcal{L})/\,\mathrm{dist}(\mathbf{t}',\mathcal{L}) - 1$ is non-negligible, we can find an $r \ge 1$ such that $p(r\,\mathrm{dist}(\mathbf{t}',\mathcal{L}))$ is significantly larger than $p(r\,\mathrm{dist}(\mathbf{t},\mathcal{L}))$. So, at least when $p(\alpha)$ is monotone, we can use this to recognize when $\mathbf{t}'$ is significantly closer to the lattice than $\mathbf{t}$ is.

Of course, in order to implement this idea, we need to generate LWE samples whose error rate is $\alpha = r\,\mathrm{dist}(\mathbf{t},\mathcal{L})$ for any desired $r \ge 1$. Fortunately, the original reduction in [Reg05] already allows for this: in addition to a BDD target $\mathbf{t}$ and lattice $\mathcal{L}$, the reduction also takes samples from a discrete Gaussian of some width r over the dual lattice $\mathcal{L}^*$; the error rate of the resulting LWE samples is (proportional to) $r\,\mathrm{dist}(\mathbf{t},\mathcal{L})$. Regev takes r to be as small as possible to minimize this error rate, but since we wish to increase the error rate multiplicatively, we use larger values of r as well. (Intuitively, sampling from the discrete Gaussian becomes easier as r increases.)

In need of monotony. The above discussion relied on the simplifying assumption that our oracle's acceptance probability $p(\alpha)$ is monotonically decreasing in the error rate α. We now describe how to drop this assumption. The idea is to use our ability to approximate $p(r\,\mathrm{dist}(\mathbf{t},\mathcal{L}))$ for any $r \ge 1$ to allow us to approximate a new monotonic function $P(\mathrm{dist}(\mathbf{t},\mathcal{L}))$.

There are many possible choices for P; we use

$$P(\alpha) := \max_{r \ge 1} (1+\log r)(p(r\alpha) - p(\infty)) \ .$$

(Here $p(\infty) := \lim_{\alpha\to\infty} p(\alpha)$, which is well approximated by, say, $p(\sqrt{n})$.) Clearly, $P(\alpha)$ is monotonically decreasing. Furthermore, it is easy to see that $P(\alpha)$ decreases noticeably in the neighborhood of any $\alpha \le \widehat{\alpha}$. In Lemma 4.2, we show how to efficiently approximate $P(\mathrm{dist}(\mathbf{t},\mathcal{L}))$ well enough to recognize when $\mathbf{t}'$ is significantly closer to the lattice than $\mathbf{t}$.

Putting a ring on it. The above ideas exploit the fact that Regev's reduction converts a BDD instance $(\mathcal{L},\mathbf{t})$ into plain-LWE samples whose error rate $\alpha \ge 0$ is proportional to $\mathrm{dist}(\mathbf{t},\mathcal{L})$. The reduction to (search) *Ring*-LWE due to [LPR10] uses a variant of this procedure that works on ideal lattices. More specifically, it converts a BDD instance $(\mathcal{I},t)$ for some ideal $\mathcal{I}$ into Ring-LWE samples, but the resulting error distribution is specified by a *vector* of error rates, not just a scalar. (See Lemma 6.8.) The error rate α_i in the ith coordinate is proportional to $|\sigma_i(v-t)|$, where $v \in \mathcal{I}$ is the closest ideal element to t, and σ_i is the ith ring embedding from the ring into the complex numbers. (See Section 2.3.1 for details.) In general, a decision oracle's acceptance probability can depend on the vector (α_i) of error parameters in complicated ways, so that it is

not immediately clear how to use the oracle to detect when a perturbation t' is significantly closer to the ideal than t is.

To adapt our reduction to the ring setting, we therefore work with each embedding *separately*, varying $\sigma_i(t)$ while holding $\sigma_j(t)$ fixed for all $j \neq i$, and eventually finding every $\sigma_i(v)$ and hence v itself. More specifically: starting from t with $z = \sigma_i(t)$, we repeatedly:

1. slightly perturb t in the ith embedding to t' with $z' = \sigma_i(t')$;
2. use the reduction from [LPR10] to approximate our oracle's acceptance probabilities for error rates $\alpha_i = r|\sigma_i(v) - z|$ and $\alpha'_i = r|\sigma_i(v) - z'|$ for various values of $r \geq 1$ (and $\alpha_j = \alpha'_j = |\sigma_j(v - t)|$ for all $j \neq i$); and
3. use this information to detect if $|\sigma_i(v) - z'|$ is significantly smaller than $|\sigma_i(v) - z|$, setting $t = t'$ if so.

Note that in order for the above procedure to work, error with parameter $\alpha_i = r|\sigma_i(v) - z|$ (and essentially any other parameters α_j for $j \neq i$) must be close to uniform modulo $R^\vee$ for large enough r. Fortunately, this is indeed the case, as shown in Lemma 6.9. We also emphasize that the search for each $\sigma_i(v)$ starts from the same *initial* target t, and varies only one embedding $\sigma_i(t)$, for the following reason. We know that our oracle has noticeable distinguishing advantage on Ring-LWE samples having error parameters $\alpha_j = |\sigma_j(v - t)|$, but we have no additional guarantees about its advantage for other parameters. It is therefore important that we not "lose hold" of the oracle's advantage as we vary the parameters, i.e., that when t' is closer to the ideal than t is, we can always find some $r \geq 1$ such that our oracle still has noticeable advantage on samples having error parameter $\alpha'_i = r|\sigma_i(v) - \sigma_i(t')|$ (and $\alpha'_j = \alpha_j$ for all $j \neq i$).

Tighter parameters via average-case BDD. As mentioned above, the approximation factor achieved by our reduction for Ring-LWE is tighter than the one from [LPR10] by a factor of $\widetilde{\Theta}(n^{1/4})$. We achieve this by giving a natural random self-reduction for BDD, showing that the ability to solve BDD with even small non-negligible probability for Gaussian-distributed offset vectors implies the ability to solve it with high probability for offset vectors that are distributed according to a slightly narrower Gaussian. (See Section 3.) A morally similar reduction appears in [LPR10], but is lossier (and messier) because it solves *worst-case* BDD, as opposed to solving it with high probability for Gaussian-distributed error.

At a technical level, the proof relies on an elegant theorem due to Borell [Bor85], which says that the "least-correlated" functions under Gaussian error are indicators of half-spaces. This theorem is used extensively in the study of the hardness of approximation and the analysis of boolean functions (see, e.g., [O'D14a, Section 11.3]), but to our knowledge, this is its first application in the study of computational problems on lattices.

2 Preliminaries

Throughout this work we assume for simplicity that for computational purposes, real numbers are specified with sufficiently high precision.

For any real $r > 0$, define the Gaussian function $\rho_r \colon \mathbb{R} \to \mathbb{R}^+$ of parameter (or width) r as $\rho_r(x) = \exp(-\pi(x/r)^2)$, and the continuous Gaussian probability distribution D_r to have density function $\rho_r(x)/r$.

Lemma 2.1 ([Reg05, Claim 2.2]). *For $0 < \alpha < \beta$, the statistical distance between D_α and D_β is at most $10(\beta/\alpha - 1)$.*

2.1 Learning with Errors

Here we recall the Learning With Errors (LWE) distribution and decision problem. We specialize to continuous Gaussian error D_α, which will be the main case of interest in this work. Let n and q be positive integers, and let $\alpha > 0$ be an error rate. The quotient ring of integers modulo q is denoted $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z}$. The quotient group of reals modulo the integers is denoted $\mathbb{T} := \mathbb{R}/\mathbb{Z}$.

Definition 2.2. For $\mathbf{s} \in \mathbb{Z}_q^n$, the LWE distribution $A_{\mathbf{s},\alpha}$ over $\mathbb{Z}_q^n \times \mathbb{T}$ is sampled by independently choosing uniformly random $\mathbf{a} \in \mathbb{Z}_q^n$ and $e \leftarrow D_\alpha$, and outputting $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle / q + e \bmod \mathbb{Z})$.

Definition 2.3. For an integer $q = q(n) \geq 2$ and error parameter $\alpha = \alpha(n) > 0$, the average-case decision problem $\mathsf{LWE}_{q,\alpha}$ is to distinguish between independent samples over $\mathbb{Z}_q^n \times \mathbb{T}$, drawn from either: (1) the LWE distribution $A_{\mathbf{s},\alpha(n)}$ for some uniformly random $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ (which is fixed for all samples), or (2) the uniform distribution.

For simplicity, the number of samples $m = m(n)$ provided in the input is usually left as an unspecified polynomial that may even depend on the algorithm. The *advantage* of an algorithm for the above problem is the (absolute value of) the difference between its acceptance probabilities on the two types of inputs.

2.2 Lattices and Gaussians

The space H. When dealing with number fields and algebraic number theory, we work with a certain linear subspace $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ for some integers $s_1, s_2 \geq 0$ where $n = s_1 + 2s_2$, defined as

$$H = \{(x_1, \ldots, x_n) : x_{s_1+s_2+i} = \overline{x_{s_1+i}}, \text{ for } 1 \leq i \leq s_2\} \subseteq \mathbb{C}^n. \tag{2.1}$$

One can verify that H, with the induced inner product from $\mathbb{C}^n$, is isomorphic to $\mathbb{R}^n$ as a real inner product space.

Lattices. For our purposes, a *lattice* is a full-rank discrete additive subgroup $\mathcal{L} \subset H$. Any lattice is generated as the set of all integer linear combinations of some (non-unique) n linearly independent *basis* vectors $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$, as $\mathcal{L} = \mathcal{L}(\mathbf{B}) = \{\sum_{i=1}^n z_i \mathbf{b}_i : z_i \in \mathbb{Z}\}$. The *minimum distance* of $\mathcal{L}$ is $\lambda_1(\mathcal{L}) := \min_{\mathbf{0} \neq \mathbf{v} \in \mathcal{L}} \|\mathbf{v}\|$, the length of a shortest nonzero lattice vector. More generally, for $1 \leq i \leq n$, the ith *successive minimum* of $\mathcal{L}$ is

$$\lambda_i(\mathcal{L}) := \inf\{r : \mathcal{L} \text{ has } i \text{ linearly independent vectors of length at most } r\} .$$

The *dual lattice* of a lattice $\mathcal{L} \subset H$ is defined as $\mathcal{L}^* := \{\mathbf{x} \in H : \langle \mathcal{L}, \mathbf{x} \rangle \subseteq \mathbb{Z}\}$.

Gaussians. We generalize D_r to axis-aligned *elliptical* Gaussian distributions over the space $H \subset \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$. Define $G = \{\mathbf{r} \in (\mathbb{R}^+)^n : r_{s_1+s_2+i} = r_{s_1+i}, \text{ for } 1 \leq i \leq s_2\}$; 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-rs vector $r\mathbf{1} \in G$. For $\mathbf{r} \in G$, the elliptical Gaussian distribution $D_\mathbf{r}$ over H is defined to have density function proportional to

$$\rho_\mathbf{r}(\mathbf{x}) := \exp\Big(-\pi \sum_{i=1}^n |x_i / r_i|^2\Big). \tag{2.2}$$

This can be seen as essentially a product distribution of s_1 real and $2s_2$ complex Gaussians, modulo the conjugate symmetry of H, where the ith real Gaussian has parameter r_i, and the ith and (s_2+i)th complex Gaussians have parameter $r_{s_1+i}/\sqrt{2}$, and are complex conjugates.

We recall a generalization of the lattice *smoothing parameter* [MR04] to elliptical Gaussians. In constrast to [MR04], where the smoothing parameter is defined to be a real number, here we capture a more general condition that makes an elliptical Gaussian "smooth" with respect to the lattice. Observe that the notation is consistent with the partial ordering on G defined by $\mathbf{r}' \geq \mathbf{r}$ if $r'_i \geq r_i$ for all i. We sometimes omit the subscript ε when it is an unspecified negligible function in n.

Definition 2.4 (Smoothing Condition). For a lattice $\mathcal{L} \subset H$, real $\varepsilon > 0$ and $\mathbf{r} \in G$, we write $\mathbf{r} \geq \eta_\varepsilon(\mathcal{L})$ if $\rho_{1/\mathbf{r}}(\mathcal{L}^* \setminus \{0\}) \leq \varepsilon$, where $1/\mathbf{r} = (1/r_1, \ldots, 1/r_n)$.

The following lemma justifies the name "smoothing parameter," and is an immediate generalization of [MR04, Lemma 4.1] to elliptical Gaussians.

Lemma 2.5. *For any lattice $\mathcal{L} \subset H$, real $\varepsilon > 0$, and $\mathbf{r} \geq \eta_\varepsilon(\mathcal{L})$, the statistical distance between $D_\mathbf{r} \bmod \mathcal{L}$ and the uniform distribution over $H/\mathcal{L}$ is at most $\varepsilon/2$.*

The following is an immediate implication of [Ban93, Lemma 1.5].

Lemma 2.6. *For any lattice $\mathcal{L} \subset H$ and $c \geq 1$, we have $c\sqrt{n}/\lambda_1(\mathcal{L}^*) \geq \eta_\varepsilon(\mathcal{L})$ where $\varepsilon = \exp(-c^2 n)$.*

The following standard fact can be found, e.g., in [Reg05, Claim 2.13].

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

Proof. Let $s := \sqrt{\log(1/\varepsilon)/\pi}/\lambda_1(\mathcal{L}^*)$. Then,

$$\rho_{1/s}(\mathcal{L}^*) \geq 1 + e^{-\pi s^2 \lambda_1(\mathcal{L}^*)^2} = 1 + \varepsilon\ . \qquad \square$$

For a lattice $\mathcal{L}$ and $\mathbf{r} \in G$, the *discrete Gaussian* probability distribution $D_{\mathcal{L},\mathbf{r}}$ is defined to have support $\mathcal{L}$, and has mass function $D_{\mathcal{L},\mathbf{r}}(\mathbf{x}) := \rho_\mathbf{r}(\mathbf{x})/\rho_\mathbf{r}(\mathcal{L})$ for $\mathbf{x} \in \mathcal{L}$.

Computational problems. In the following computational problems, a lattice $\mathcal{L}$ is represented by an arbitrary basis $\mathbf{B}$, and a lattice coset $\mathbf{e}+\mathcal{L}$ is represented by its distinguished representative $\overline{\mathbf{e}} = (\mathbf{e}+\mathcal{L}) \cap \mathcal{P}(\mathbf{B})$, where $\mathcal{P}(\mathbf{B}) := \mathbf{B} \cdot [-\frac{1}{2}, \frac{1}{2})^n$ is the fundamental parallelepiped of $\mathbf{B}$.

Definition 2.8 (Gap Shortest Vector Problem). For an approximation factor $\gamma = \gamma(n) \geq 1$, the GapSVP_γ is: given a lattice $\mathcal{L}$ and length $d > 0$, output YES if $\lambda_1(\mathcal{L}) \leq d$ and NO if $\lambda_1(\mathcal{L}) > \gamma d$.

Definition 2.9 (Shortest Independent Vectors Problem). For an approximation factor $\gamma = \gamma(n) \geq 1$, the SIVP_γ is: given a lattice $\mathcal{L}$, output n linearly independent lattice vectors of length at most $\gamma(n) \cdot \lambda_n(\mathcal{L})$.

Definition 2.10 (Discrete Gaussian Sampling). For a function γ that maps lattices to nonnegative reals, the DGS_γ problem is: given a lattice $\mathcal{L}$ and a parameter $r \geq \gamma(\mathcal{L})$, output an independent sample from a distribution that is within negligible statistical distance of $D_{\mathcal{L},r}$.

Definition 2.11 (Bounded Distance Decoding). For a function δ that maps lattices to nonnegative reals, the BDD_δ problem is: given a lattice $\mathcal{L} \subset H$, a distance bound $d \leq \delta(\mathcal{L})$, and a coset $\mathbf{e} + \mathcal{L}$ where $\|\mathbf{e}\| \leq d$, output $\mathbf{e}$.

Lemma 2.12 ([LLL82, Bab85]). *There is an efficient algorithm that solves* BDD_δ *for* $\delta(\mathcal{L}) = 2^{-n/2} \cdot \lambda_1(\mathcal{L})$.

Finally, we define a new average-case problem, which is essentially BDD where the offset is drawn from a Gaussian.

Definition 2.13 (Gaussian Decoding Problem). For a lattice $\mathcal{L} \subset H$ and a Gaussian parameter $g > 0$, the $\mathsf{GDP}_{\mathcal{L},g}$ 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

Here we briefly review the requisite concepts and notation from algebraic number theory. Our presentation is an abridged version of [LPR10, Section 2.3]; see that work and references therein for many more details.

2.3.1 Number Fields and Their Geometry

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 the *minimal polynomial* of ζ, which is monic without loss of generality. The *degree* n of the number field is the degree of f.

A number field $K = \mathbb{Q}(\zeta)$ of degree n has exactly n ring embeddings (i.e., injective ring homomorphisms) $\sigma_i \colon K \to \mathbb{C}$. Each embedding sends ζ to one of the roots of its minimal polynomial f; the embedding is said to be *real* if that root is real (in which case the image of the embedding is contained in $\mathbb{R}$), otherwise it is *complex*. Because the roots of f come in conjugate pairs, so too do the complex embeddings. We let s_1 and s_2 respectively be the number of real embeddings and *pairs* of complex embeddings, so $n = s_1 + 2s_2$, with σ_i for $1 \le i \le s_1$ being the real embeddings and $\sigma_{s_1+s_2+i} = \overline{\sigma_{s_1+i}}$ for $1 \le i \le s_2$ being the conjugate pairs of complex embeddings.

The *(field) norm* (or *algebraic norm*) of an element $a \in K$ can be defined as $\mathrm{N}(a) := \prod_{i=1}^n \sigma_i(a)$; clearly, the norm is multiplicative. Similarly, the *trace* can be defined as $\mathrm{Tr}(a) = \sum_{i=1}^n \sigma_i(a)$. The field norm and trace of a number field element is always rationals.

The *canonical embedding* $\sigma \colon K \to H$, where $H \subset \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ is as defined in Section 2.2 above, is defined as $\sigma(a) = (\sigma_1(a), \ldots, \sigma_n(a))$. Observe that this is a ring homomorphism from K to H, where multiplication and addition in the latter are both component-wise. We define norms and other geometric quantities on K simply by identifying field elements $a \in K$ with their canonical embeddings $\sigma(a) \in H$, e.g., the ℓ_2 norm is $\|a\|_2 := \|\sigma(a)\|_2 = (\sum_{i=1}^n |\sigma_i(a)|^2)^{1/2}$ and the ℓ_∞ norm is $\|a\|_\infty = \max_i |\sigma_i(a)|$.

The canonical embedding also allows us to view Gaussian distributions $D_{\mathbf{r}}$ over H (for $\mathbf{r} \in G$), or their discrete analogues over a lattice $\mathcal{L} \subset H$, as distributions over K. Formally, the continuous distribution $D_{\mathbf{r}}$ is actually over the field tensor product $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$, which is in bijective correspondence with H via the natural extension of σ.

2.3.2 Ring of Integers and Ideals

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

An *(integral) ideal* $\mathcal{I} \subseteq \mathcal{O}_K$ is a nontrivial additive subgroup that is also closed under multiplication by $\mathcal{O}_K$, i.e., $r \cdot a \in \mathcal{I}$ for any $r \in \mathcal{O}_K$ and $a \in \mathcal{I}$. Any ideal $\mathcal{I}$ is a free $\mathbb{Z}$-module of rank n, i.e., it is the set of all $\mathbb{Z}$-linear combinations of some basis $\{b_1, \ldots, b_n\} \subset \mathcal{I}$ of linearly independent (over $\mathbb{Z}$) elements b_i.

The *norm* of an ideal $\mathcal{I}$ is its index as a subgroup of $\mathcal{O}_K$, i.e., $\mathrm{N}(\mathcal{I}) := |\mathcal{O}_K/\mathcal{I}|$. The product $\mathcal{I}\mathcal{J}$ of two ideals $\mathcal{I}, \mathcal{J}$ is the set of all sums of terms xy for $x \in \mathcal{I}, y \in \mathcal{J}$. The norm for ideals is consistent with the norm for field elements, in that $\mathrm{N}(a\mathcal{O}_K) = |\mathrm{N}(a)|$ for any $a \in \mathcal{O}_K$, and $\mathrm{N}(\mathcal{I}\mathcal{J}) = \mathrm{N}(\mathcal{I})\,\mathrm{N}(\mathcal{J})$.

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

2.3.3 Ideal Lattices and Duality

Under the canonical embedding σ, any fractional ideal $\mathcal{I} \subset K$ yields a lattice $\mathcal{L} = \sigma(\mathcal{I}) \subset H$, which is called an *ideal lattice*. Recalling that $\mathcal{I}$ has a $\mathbb{Z}$-basis $B = \{b_1, \ldots, b_n\}$, we see that $\mathbf{B} = \{\sigma(b_1), \ldots, \sigma(b_n)\} \subset \mathcal{L}$ is a basis of $\mathcal{L}$. We often identify an ideal with its embedded lattice, and refer to the minimum distance $\lambda_1(\mathcal{I})$ of an ideal, etc. The *(absolute) discriminant* Δ_K of a number field K is the squared volume of the ring of integers lattice $\sigma(\mathcal{O}_K)$, i.e., $\Delta_K = |\det(\mathrm{Tr}(b_i \cdot b_j))|$ where $\{b_1, \ldots, b_n\}$ is any $\mathbb{Z}$-basis of $\mathcal{O}_K$.

For any fractional ideal $\mathcal{I} \subset K$, its *dual ideal* is defined as $\mathcal{I}^\vee := \{a \in K : \mathrm{Tr}(x\mathcal{I}) \subseteq \mathbb{Z}\}$. Observe that because $\mathrm{Tr}(ab) = \sum_{i=1}^n \sigma_i(a)\sigma_i(b) = \langle \sigma(a), \overline{\sigma(b)} \rangle$, dual ideals and dual lattices are related by $\sigma(\mathcal{I}^\vee) = \overline{\sigma(\mathcal{I})^*}$. An important canonical fractional ideal in a number field K is the *codifferent* ideal $\mathcal{O}_K^\vee$, i.e., the dual ideal of the ring of integers. This ideal has norm $\mathrm{N}(\mathcal{O}_K^\vee) = \Delta_K^{-1}$, and provides a link between dual and inverse ideals: $\mathcal{I}^\vee = \mathcal{I}^{-1} \cdot \mathcal{O}_K^\vee$ for any fractional ideal $\mathcal{I} \subset K$.

All of the computational problems on lattices defined in Section 2.2 can be specialized by restricting them to (fractional) ideal lattices in a number field K. We refer to these specialized problems by prefixing them by K, e.g., K-$\mathsf{GDP}_{\mathcal{I},r}$, K-DGS_γ, etc.

2.4 Ring-LWE

Let K be a number field with ring of integers $R = \mathcal{O}_K$. Recall that $R^\vee$ is the (fractional) codifferent ideal of K, and let $\mathbb{T} = K_\mathbb{R}/R^\vee$. Let $q \geq 2$ be a (rational) integer modulus, and for any fractional ideal $\mathcal{I}$ of K, let $\mathcal{I}_q = \mathcal{I}/q\mathcal{I}$.

Definition 2.14 (Ring-LWE Distribution). For $s \in R_q^\vee$ and an error distribution ψ over $K_\mathbb{R}$, the R-LWE distribution $A_{s,\psi}$ over $R_q \times \mathbb{T}$ is sampled by independently choosing a uniformly random $a \leftarrow R_q$ and an error term $e \leftarrow \psi$, and outputting $(a, b = (a \cdot s)/q + e \bmod R^\vee)$.

Definition 2.15 (Ring-LWE, Average-Case Decision). 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-$\mathsf{LWE}_{q,\Upsilon}$, is to distinguish (with non-negligible advantage) between independent samples from $A_{s,\psi}$ for a *random* choice of $(s, \psi) \leftarrow U(R_q^\vee) \times \Upsilon$, and the same number of uniformly random and independent samples from $R_q \times \mathbb{T}$.

2.5 Probability

Lemma 2.16 (Chernoff-Hoeffding bound [Hoe63]). *Let $X_1, \ldots, X_N$ be independent and identically distributed random variables with $0 \leq X_i \leq 1$ and $\overline{X} := \mathbb{E}[X_i]$. Then, for any $s > 0$*

$$\Pr\left[\left|N\overline{X} - \sum X_i\right| \geq s\right] \leq 2\exp(-s^2/N)\ .$$

3 Random Self-Reduction for GDP

Our goal in this section is to prove Proposition 3.3, a random self-reduction for GDP, showing that the ability to solve GDP even with some small non-negligible probability over the offset vector implies the ability to solve GDP for a slightly smaller parameter with all but negligible probability.

The proof is based on the following theorem by Borell [Bor85]. It says that the functions that are "least correlated" under Gaussian error are indicators of two opposing half-spaces. We let $D^n = D_1^n$ be the Gaussian distribution D_1 in n-dimensional space. Let $\Phi^{-1} : (0,1) \to \mathbb{R}$ denote the inverse of the cumulative distribution function of the normal distribution D^1 (and notice that $\Phi^{-1}(\mu) < 0$ for $\mu < 1/2$).

Theorem 3.1 ([Bor85]; see also [O'D14b, Exercise 11.26]). *Let $f, g\colon \mathbb{R}^n \to [0,1]$ be two functions and let $\mu := \mathbb{E}_{\mathbf{x}\sim D^n}[f(\mathbf{x})], \nu := \mathbb{E}_{\mathbf{x}\sim D^n}[g(\mathbf{x})]$. Then, for any $\delta > 0$,*

$$\mathop{\mathbb{E}}_{\mathbf{x},\mathbf{z}\sim D^n}[f(\mathbf{x})g(\delta\mathbf{x} + (1-\delta^2)^{1/2}\mathbf{z})] \geq \Pr_{x,z\sim D^1}[x \leq \Phi^{-1}(\mu) \text{ and } \delta x + (1-\delta^2)^{1/2}z \geq -\Phi^{-1}(\nu)] \,.$$

Corollary 3.2. *Let $\delta = \delta(n)$ be any $o(1)$ function, and let $g\colon \mathbb{R}^n \to [0,1]$ have expectation $\nu := \mathbb{E}_{\mathbf{x}\sim D^n}[g(\mathbf{x})] \geq 1/\operatorname{poly}(n)$. Assume $\mathbf{x}$ is chosen according to D^n. Then, there exists a $c > 0$ such that with all but negligible probability (over the choice of $\mathbf{x}$), it holds that*

$$\mathop{\mathbb{E}}_{\mathbf{z}\sim D^n}[g(\delta\mathbf{x} + (1-\delta^2)^{1/2}\mathbf{z})] > n^{-c} \,. \tag{3.1}$$

Proof. Let $f : \mathbb{R}^n \to \{0,1\}$ be the indicator function of the points for which Equation (3.1) does *not* hold, and denote its expectation by $\mu := \mathbb{E}_{\mathbf{x}\sim D^n}[f(\mathbf{x})]$. Then, clearly

$$\mathop{\mathbb{E}}_{\mathbf{x},\mathbf{z}\sim D^n}[f(\mathbf{x})g(\delta\mathbf{x} + (1-\delta^2)^{1/2}\mathbf{z})] \leq \mu \cdot n^{-c} \,.$$

By Theorem 3.1, we get that

$$\Pr_{x,z\sim D^1}[x \leq \Phi^{-1}(\mu) \text{ and } \delta x + (1-\delta^2)^{1/2}z \geq -\Phi^{-1}(\nu)] \leq \mu \cdot n^{-c} \,,$$

or equivalently,

$$\mathop{\mathbb{E}}_{x\sim D^1}\big[\Pr_{z\sim D^1}[\delta x + (1-\delta^2)^{1/2}z \geq -\Phi^{-1}(\nu)] \mid x \leq \Phi^{-1}(\mu)\big] \leq n^{-c} \,.$$

By Markov's inequality, with probability at least $1/2$ over the choice of x conditioned on $x \leq \Phi^{-1}(\mu)$,

$$\Pr_{z\sim D^1}[\delta x + (1-\delta^2)^{1/2}z \geq -\Phi^{-1}(\nu)] \leq 2n^{-c} \,,$$

but since $\delta = o(1)$, and assuming c is large enough, this implies that $x \leq -\omega(\sqrt{\log n})$. As a result, $\Phi^{-1}(\mu) \leq -\omega(\sqrt{\log n})$ which in turn implies that μ is negligible, as desired. □

Proposition 3.3. *Assume we are given access to an oracle that solves $\mathsf{GDP}_{\mathcal{L},r}$ with some* non-negligible *probability over the choice of the coset. Then we can efficiently solve $\mathsf{GDP}_{\mathcal{L},\delta r}$ with* all but negligible *probability, where $\delta = \delta(n)$ is any $o(1)$ function.*

Proof. Assume without loss of generality that $r = 1$. Given an input coset $\mathbf{x} + \mathcal{L}$, the algorithm repeats the following a polynomial number of times. It chooses a vector $\mathbf{z}$ from $D_{r'}$, where $r' := (1-\delta^2)^{1/2}r$, and calls the oracle with $\mathbf{z} + \mathbf{x} + \mathcal{L}$. At the end it returns the shortest solution among all the oracle responses, minus $\mathbf{z}$. Correctness follows from Corollary 3.2, where we take g to be the acceptance probability of the oracle. □

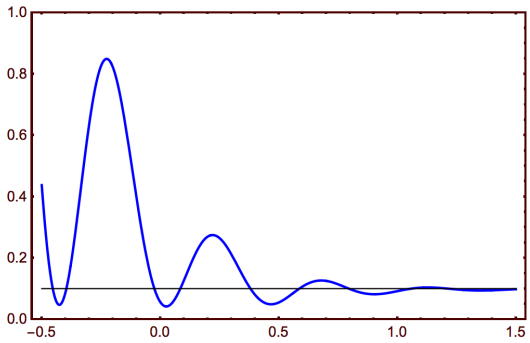

(a) The probability function $p_{\mathcal{O}}(t)$ of an oracle $\mathcal{O}$ with a horizontal line at $p(\infty) := \lim_{t\to\infty} p_{\mathcal{O}}(t) = 1/10$.

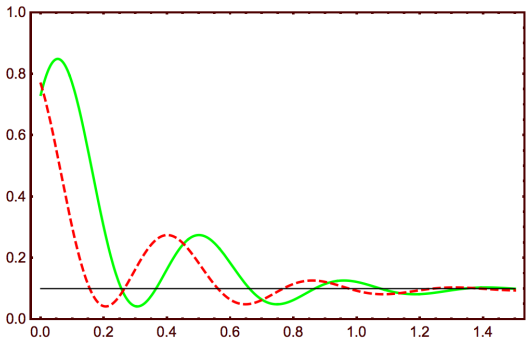

(b) The probability functions $p_{\mathcal{O}_{-0.3}}(t)$ and $p_{\mathcal{O}_{-0.2}}(t)$ of two different shifts of $\mathcal{O}$.

Figure 1

4 Finding Your (Oracle's) Center

For a (randomized) oracle $\mathcal{O}\colon S \to \{0,1\}$ with some domain S, let $p_{\mathcal{O}}(t) = \Pr[\mathcal{O}(t) = 1]$ for any $t \in S$. Usually, $\mathcal{O}$ will be clear from context, so we often omit the subscript and simply write $p(t)$. If $S = \mathbb{R}$, then for any $s \in \mathbb{R}$, let $\mathcal{O}_s\colon \mathbb{R}^{\geq 0} \to \{0,1\}$ be the "suffix" oracle $\mathcal{O}_s(t) = \mathcal{O}(s+t)$. (See Figure 1.) We wish to show that, if $\mathcal{O}$ has certain properties, then we can tell the difference between suffix oracles having sufficiently different shifts. (In our only application, Proposition 4.4, the shift s corresponds to $\log \|\mathbf{z}^* - \mathbf{z}\|$, where $\mathbf{z}^* \in \mathbb{R}^k$ is some unknown target vector and we choose $\mathbf{z} \in \mathbb{R}^k$.)

Definition 4.1 (Oracle Comparison Problem). For error parameter $\varepsilon \geq 0$ and $r > 0$, (ε, r)-OCP is a promise problem defined as follows. Given access to two (randomized) oracles $\mathcal{O}_1 := \mathcal{O}_{s_1}$ and $\mathcal{O}_2 := \mathcal{O}_{s_2}$ for some unknown shifts $s_1, s_2 \in [-r, r]$ and some underlying oracle $\mathcal{O}\colon \mathbb{R} \to \{0,1\}$, the goal is to output YES if $s_2 \leq s_1 - \varepsilon$ and NO if $s_2 > s_1$. (If neither is the case, then any output is considered correct.)

Lemma 4.2. *There is a* $\mathrm{poly}(\kappa)$*-time algorithm that takes as input some confidence parameter* $\kappa \geq 200$ *and solves* $(1/\kappa, \kappa)$*-OCP except with probability at most* $\exp(-\kappa)$*, provided that there exists a* $p(\infty) \in [0,1]$ *and* $t^* \geq s_1$ *such that*

1. $p(t^*) - p(\infty) \geq 1/\kappa$ *;*
2. $|p(t) - p(\infty)| \leq 2\exp(-t/\kappa)$ *for all* t*; and*
3. $p(t)$ *is* κ*-Lipschitz.*

Furthermore, each of the algorithm's oracle calls takes the form $\mathcal{O}_j(i\Delta)$ *for some* $\Delta < 1$ *that depends only on* κ *and integer* $0 \leq i \leq \mathrm{poly}(\kappa)$.

Proof. On input $\kappa \geq 200$, the algorithm behaves as follows. For $i = 0, \ldots, T := 1000\kappa^{10}$, it calls $\mathcal{O}_1$ and $\mathcal{O}_2$ repeatedly, $N := \lceil 200 \log T/\Delta^2 \rceil$ times each with $\Delta := 1/(200\kappa^8)$, on input $i\Delta$, and sets $\overline{p}_i^{(1)}, \overline{p}_i^{(2)}$ to be the respective empirical probabilities that the oracles output 1. The algorithm then computes

$$h_j := \max_i (1 + i\Delta)(\overline{p}_i^{(j)} - \overline{p}_T^{(1)}) \ .$$

for $j \in \{1,2\}$. Finally, it returns YES if $h_2 - h_1 > 1/(20\kappa^4)$ and NO otherwise.

The running time is clear, as is the fact that the oracle queries are polynomially bounded and lie in an arithmetic progression.

To prove correctness, we define $m := s_1 - s_2$ and $q(t) := p(s_1 + t) - p(\infty)$. Let

$$h'_1 := \max_{0 \le t \le T\Delta} (1+t)q(t) \ ,$$

and

$$h'_2 := \max_{0 \le t \le T\Delta} (1+t)q(t-m) = \max_{-m \le t \le T\Delta - m} (1+t+m)q(t) \ .$$

We first wish to argue that $\overline{p}_i^{(j)} \approx p(s_j + i\Delta)$, $p(s_j + (i+\chi)\Delta) \approx p(s_j + i\Delta)$ for any $\chi \in [0,1]$, and $\overline{p}_T^{(1)} \approx p(\infty)$. This will allow us to argue that $h_j \approx h'_j$. Indeed, by the Chernoff-Hoeffding bound (Lemma 2.16), we have

$$|\overline{p}_i^{(j)} - p(s_j + i\Delta)| \le \frac{\kappa\Delta}{10} \tag{4.1}$$

for all i, j except with probability at most $10T\exp(-N(\kappa\Delta)^2/100) \le \exp(-\kappa)$. So, we may assume that this holds. By Item 3, we have

$$|p(s_j + (i+\chi)\Delta) - p(s_j + i\Delta)| \le \kappa\chi\Delta \le \kappa\Delta \tag{4.2}$$

for any $\chi \in (0,1)$. Furthermore, by Item 2, we have

$$|p(s_1 + T\Delta) - p(\infty)| \le 2\exp(-(s_1 + T\Delta)/\kappa) \ll \frac{\kappa\Delta}{10} \ . \tag{4.3}$$

Combining Eqs. (4.1), (4.2), and (4.3) and recalling the definitions of h_j, h'_j gives

$$|h_2 - h_1 - (h'_2 - h'_1)| < 5\kappa T\Delta^2 \le \frac{1}{20\kappa^4} \ . \tag{4.4}$$

So, we move to studying $h'_2 - h'_1$.

Suppose $s_1 - s_2 = m \le 0$. Let $t \ge -m$ such that $h'_2 = (1+t+m)q(t)$. If $t > T\Delta$, then by Item 2, $h'_2 \ll 1/(2\kappa)$, but by Item 1, $h'_1 \ge 1/\kappa \ge h'_2$. Otherwise, $t \le T\Delta$ and therefore

$$h'_1 \ge (1+t)q(t) \ge (1+t+m)q(t) = h'_2 \ .$$

So, in both cases, $h'_1 \ge h'_2$. Combining this with Equation (4.4), we see that the algorithm will correctly output NO.

Suppose, on the other hand, that $s_1 - s_2 = m \ge 1/\kappa$. Let $t \ge 0$ be such that $h'_1 = (1 + \log t)q(t)$. By Item 1, we have that $h'_1 \ge 1/\kappa$. Therefore,

$$q(t) \ge \frac{1}{\kappa(1+t)} \ , \tag{4.5}$$

and by Item 2, we see that

$$t + 1 \le 10\kappa^2 \le T\Delta - m \ . \tag{4.6}$$

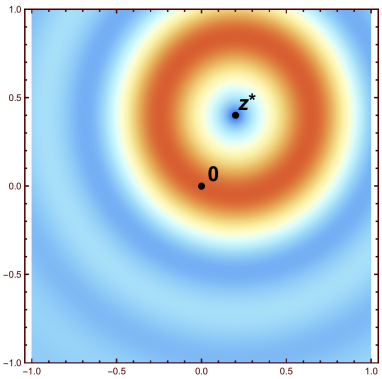

Figure 2: A depiction of an instance of OHCP in $\mathbb{R}^2$. The goal is to find a good approximation to $\mathbf{z}^*$ using calls to the oracle $\mathcal{O}(\mathbf{z}, t)$.

Putting everything together, we have

$$\begin{aligned} h_2' - h_1' &\geq (1 + t + m)q(t) - h_1' \\ &= m \cdot q(t) \\ &\geq \frac{m}{\kappa(1+t)} && \text{(Equation (4.5))} \\ &\geq \frac{1}{\kappa^2(1+t)} \\ &\geq \frac{1}{10\kappa^4} && \text{(Equation (4.6))} . \end{aligned}$$

It follows from Equation (4.4) that the algorithm correctly outputs YES. □

Definition 4.3 (Oracle Hidden Center Problem). For any parameters $\varepsilon, \delta \in [0, 1)$, the (ε, δ)-OHCP is an approximate search problem defined as follows. Given a scale parameter $d > 0$ and access to a (randomized) oracle $\mathcal{O}\colon \mathbb{R}^k \times \mathbb{R}^{\geq 0} \to \{0, 1\}$ such that its acceptance probability $p(\mathbf{z}, t)$ depends only on $\exp(t)\|\mathbf{z} - \mathbf{z}^*\|$ for some (unknown) "hidden center" $\mathbf{z}^* \in \mathbb{R}^k$ with $\delta d \leq \|\mathbf{z}^*\| \leq d$, the goal is to output $\mathbf{z} \in \mathbb{R}^k$ such that $\|\mathbf{z} - \mathbf{z}^*\| \leq \varepsilon d$.

Proposition 4.4. *There is a $\mathrm{poly}(\kappa, k)$-time algorithm that takes as input a confidence parameter $\kappa \geq 20\log(k+1)$ (and the scale parameter $d > 0$) and solves $(\exp(-\kappa), \exp(-\kappa))$-OHCP in dimension k except with probability $\exp(-\kappa)$, provided that there exists a $p(\infty) \in [0, 1]$ such that*

1. *$p(\mathbf{0}, t^*) - p(\infty) \geq 1/\kappa$ for some $t^* \geq 0$;*
2. *$|p(\mathbf{0}, t) - p(\infty)| \leq 2\exp(-t/\kappa)$ for any t; and*
3. *$p(\mathbf{z}, t)$ is κ-Lipschitz in t for any $\mathbf{z} \in \mathbb{R}^k$.*

Furthermore, each of the algorithm's oracle calls takes the form $\mathcal{O}(\cdot, i\Delta)$ for some $\Delta < 1$ that depends only on κ and k and $0 \leq i \leq \mathrm{poly}(\kappa, k)$.

Proof. The idea behind the algorithm is to take a "guided random walk" towards the center $\mathbf{z}^*$. In particular, starting with $\mathbf{z} = \mathbf{0}$, we repeatedly randomly perturb $\mathbf{z}$ to $\mathbf{z}'$, use our oracle and Lemma 4.2 to check whether

$\|\mathbf{z}^* - \mathbf{z}'\| \le \|\mathbf{z}^* - \mathbf{z}\|$, and update $\mathbf{z}$ to $\mathbf{z}'$ if it is. Since we do not know $\|\mathbf{z}^* - \mathbf{z}\|$, we choose the size of the perturbation randomly on an exponential scale (and we use axis-aligned perturbations for simplicity). We proceed with the technical details.

We may assume that $d = 1$, since the problem is invariant under scalings. On input $\kappa \ge 20 \log(k+1)$, the algorithm first sets $\mathbf{z}_0 = \mathbf{0}$. It then does the following for $i = 0, \dots, T := \lceil 100\kappa^2 k^2 \rceil$. It samples a coordinate $j \in \{1, \dots, k\}$, $x \in [0, 1]$, and $\sigma \in \{\pm 1\}$ uniformly at random. Let $\mathbf{v}_i := \sigma \exp(-2\kappa x)\mathbf{e}_j/\sqrt{k}$. The algorithm then simulates the procedure from Lemma 4.2 with oracles $\mathcal{O}_1 := \mathcal{O}(\mathbf{z}_i, \cdot)$ and $\mathcal{O}_2 := \mathcal{O}(\mathbf{z}_i + \mathbf{v}_i, \cdot)$ and parameter $\kappa' := 100\kappa k \log \kappa$. If the oracle outputs YES, then it sets $\mathbf{z}_{i+1} = \mathbf{z}_i + \mathbf{v}_i$. Otherwise, it sets $\mathbf{z}_{i+1} = \mathbf{z}_i$. Finally, it outputs $\mathbf{z}_{T+1}$.

The running time is clear, as is the fact that the values of t used in the oracle queries lie in a polynomially bounded arithmetic progression. To prove correctness, let $\mathcal{O}^*(\cdot) := \mathcal{O}(\mathbf{z}^* + \mathbf{e}_1, \cdot)$. Note that for any $\mathbf{z} \ne \mathbf{z}^*$, the oracle $\mathcal{O}(\mathbf{z}, \cdot)$ is $\mathcal{O}^*_{\log \|\mathbf{z} - \mathbf{z}^*\|}$. I.e., $\mathcal{O}(\mathbf{z}, \cdot)$ is a shift of $\mathcal{O}^*$ by $\log \|\mathbf{z} - \mathbf{z}^*\|$.

Therefore, the input to the subprocedure in Lemma 4.2 will be a valid instance to $(-(\kappa')^2, 1/\kappa')$-OCP as long as $\log \|\mathbf{z}_i - \mathbf{z}^*\|$ and $\log \|\mathbf{z}_i + \mathbf{v}_i - \mathbf{z}^*\|$ lie in the interval $[-\kappa', \kappa']$. The upper bound is trivial, since in any given step the step size is at most $1/\sqrt{k}$ and there are only $T + 1$ steps, so that $\|\mathbf{z}_{T+1}\| \le (T+1)/\sqrt{k} \ll \exp(\kappa')$. To prove the lower bound, we must show that $\|\mathbf{z}_i - \mathbf{z}^*\| \ge \exp(-\kappa')$ for all i except with small probability. Note that this is true by assumption for $\mathbf{z}_0$. So, it suffices to show that $\|\mathbf{z}_i + \mathbf{v}_i - \mathbf{z}^*\| \ge \exp(-\kappa')$ for all i, except with negligible probability. Indeed, recall that $\|\mathbf{v}_i\| = \exp(-2\kappa x)/\sqrt{k}$ for a uniformly random $x \in [0, 1]$. If $\|\mathbf{z}_i + \mathbf{v}_i - \mathbf{z}^*\| \le -\exp(\kappa')$, then by triangle inequality, we must have

$$\|\mathbf{z}_i - \mathbf{z}^*\| - \exp(-\kappa') \le \|\mathbf{v}_i\| \le \|\mathbf{z}_i - \mathbf{z}^*\| + \exp(-\kappa') \; .$$

But, the probability of this happening is at most, say, $\exp(-\kappa'/2)$. By union bound, this will not happen for any i except with probability at most, say, $\exp(-10\kappa)$.

Now, we show that "we never get farther from $\mathbf{z}^*$." Assume for induction that $\|\mathbf{z}_j - \mathbf{z}^*\| \le \|\mathbf{z}_{j-1} - \mathbf{z}^*\|$ for all $1 \le j \le i$ except with probability $i \exp(-\kappa')$. Note in particular that this implies that $\|\mathbf{z}_i - \mathbf{z}^*\| \le \|\mathbf{z}_0 - \mathbf{z}^*\|$. This shows that Item 1 of Lemma 4.2 holds for the ith instance of OCP. (Here, we take $(t^*)' := t^* + \log \|\mathbf{z}^*\|$ and $s_1 := \log \|\mathbf{z}_i - \mathbf{z}^*\| \le \log \|\mathbf{z}^*\|$.) The other items of Lemma 4.2 are immediate. Therefore, except with probability $\exp(-\kappa')$, the oracle only outputs YES when $\|\mathbf{z}_i + \mathbf{v}_i - \mathbf{z}^*\| \ge \|\mathbf{z}_i - \mathbf{z}^*\|$. It follows immediately that $\|\mathbf{z}_{i+1} - \mathbf{z}^*\| \le \|\mathbf{z}_i - \mathbf{z}^*\|$ except with probability at most $(i+1)\exp(-\kappa')$, as needed. So, we may assume that the subprocedure from Lemma 4.2 always returns a valid response to its input OCP instance.

Finally, we show that the output of the algorithm is correct. During the ith step, if $\log \|\mathbf{z}_i - \mathbf{z}^*\| \le -\kappa$, then by the above argument, the same will be true of $\mathbf{z}_{T+1}$, and we are done. Otherwise, let $\mathbf{y} := \mathbf{z}^* - \mathbf{z}_i$, and note that there exists some coordinate j such that $|y_j| \ge \|\mathbf{y}\|/\sqrt{k} \ge \exp(-\kappa)/\sqrt{k}$. We also have $|y_j| \le \|\mathbf{y}\| \le 1$. So, with probability at least $\log 2/(4\kappa k)$, the algorithm will select the coordinate j, $\sigma = \operatorname{sign}(y_j)$, and $x \in [-\log |y_j|/(2\kappa), -\log(|y_j|/2)/(2\kappa)]$. If this happens, then

$$\|\mathbf{z}_i + \mathbf{v}_i - \mathbf{z}^*\|^2 \le (1 - 1/(2\sqrt{k}))^2 y_j^2 + \sum_{j' \ne j} y_k^2 \le \|\mathbf{y}\|^2 \cdot (1 - 1/(2k)) \; .$$

In particular, $\log \|\mathbf{z}_i + \mathbf{v}_i - \mathbf{z}^*\| \le \log \|\mathbf{z}_i - \mathbf{z}^*\| + \log(1 - 1/(2k))/2 \le \log \|\mathbf{z}_i - \mathbf{z}^*\| - 1/\kappa'$, and the subprocedure from Lemma 4.2 must output YES. So, if this is the case, we have $\mathbf{z}_{i+1} = \mathbf{z}_i + \mathbf{v}_i$ and therefore

$$\log \|\mathbf{z}_i - \mathbf{z}^*\| - \log \|\mathbf{z}_{i+1} - \mathbf{z}^*\| \ge -\log(1 - 1/(2k)) \ge \frac{1}{2k} \; .$$

We conclude that, unless $\log \|\mathbf{z}_i - \mathbf{z}^*\| \leq -\kappa$, we have

$$\Pr\Big[\log \|\mathbf{z}_i - \mathbf{z}^*\| - \log \|\mathbf{z}_{i+1} - \mathbf{z}^*\| \geq \frac{1}{2k}\Big] \geq \frac{\log 2}{4\kappa k} \ .$$

Therefore, after running the protocol $T+1$ times, either $\log \|\mathbf{z}_i - \mathbf{z}^*\| \leq -\kappa$ for some i, or by Lemma 2.16, the log-distance will drop by at least $1/(2k)$ at least $T \log 2/(4\kappa k) - \kappa^2 k \geq 10\kappa k$ times except with probability at most $\exp(-\kappa^4 k^2/(10T)) \leq \exp(-10\kappa)$. But, if this happens, then $\log \|\mathbf{z}_{T+1} - \mathbf{z}^*\| \leq -10\kappa k/(2k) \leq -\kappa$. So, it follows that $\log \|\mathbf{z}_{T+1} - \mathbf{z}^*\| \leq \exp(-\kappa)$, as needed. □

5 Hardness of plain LWE

Before we present our main result, a worst-case to average-case reduction for Ring-LWE for any ring, we show how to use the same technique to derive a worst-case to average-case reduction for plain LWE. This reduction is significantly simpler and illustrates most of the new ideas necessary for the ring case. In particular, we prove the following result, which is identical to [Reg05, Theorem 3.1], except that we reduce to *decision* LWE, rather than search.

Theorem 5.1. *Let $\alpha = \alpha(n) \in (0,1)$, and let $q = q(n)$ be an integer such that $\alpha q \geq 2\sqrt{n}$. There is a polynomial-time quantum reduction from DGS_γ to (average-case, decision) $\mathsf{LWE}_{q,\alpha}$, where $\gamma := \sqrt{2n}\eta(\mathcal{L})/\alpha$.*

Using known reductions (see, e.g., [Reg05, Section 3.3]), we immediately derive the following corollary.

Corollary 5.2. *Let $\alpha = \alpha(n) \in (0,1)$, and let $q = q(n)$ be an integer such that $\alpha q \geq 2\sqrt{n}$. Then, there is a polynomial-time quantum reduction from $\mathsf{SIVP}_{\gamma'}$ and $\mathsf{GapSVP}_{\gamma'}$ to (average-case, decision) $\mathsf{LWE}_{q,\alpha}$ for some $\gamma' = \widetilde{O}(n/\alpha)$.*

The high-level structure of the proof of Theorem 5.1 is nearly identical to that of the corresponding statement in [Reg05]. In particular, [Reg05] shows (1) a quantum reduction from $\mathsf{DGS}_{\sqrt{n/2}/d}$ to BDD_d over $\mathcal{L}^*$ for $d < \lambda_1(\mathcal{L}^*)/2$; and (2) a classical reduction that allows us to solve BDD_d over $\mathcal{L}^*$ using discrete Gaussian samples over $\mathcal{L}$ with parameter $r \geq \sqrt{2}q\eta(\mathcal{L})$ together with a *search* $\mathsf{LWE}_{q,\alpha}$ oracle, with $d := \alpha q/(\sqrt{2}r)$. By combining these two steps, we can use an $\mathsf{LWE}_{q,\alpha}$ oracle to convert discrete Gaussian samples with parameter γ to samples with parameter γ', where

$$\gamma' := \max\{\sqrt{2n}\eta(\mathcal{L})/\alpha,\ \sqrt{n}\gamma/(\alpha q)\} \ .$$

The full reduction is obtained by starting with very large Gaussian samples (which we can sample efficiently) and then repeatedly running this procedure to lower the parameter.

The main difference between our reduction and that of [Reg05] is that we use a novel classical reduction from BDD to *decision* LWE, rather than search LWE. (See Lemma 5.4.) Like [Reg05], we also require access to discrete Gaussian samples to achieve this. But, while this part of Regev's reduction uses samples with a single parameter r, we require samples with many parameters $r' \geq r$. (Recall from Section 1.3 that we vary this parameter in order to obtain LWE samples with a variety of errors.) Regev's quantum reduction allows for this, so this is a nonissue.

In Lemma 5.4, we show our modified classical reduction. We will need the following lemma from [Reg05], which shows how to convert discrete Gaussian samples together with a BDD instance into LWE samples.

Lemma 5.3. *There is an efficient algorithm that takes as input an integer $q \geq 2$, a lattice $\mathcal{L} \subset \mathbb{R}^n$, a coset $\mathbf{x} + \mathcal{L}^*$, bound $d \geq \|\mathbf{x}\|$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{L})$, and samples from $D_{\mathcal{L},r}$. It outputs samples that are within negligible statistical distance of LWE samples with a uniformly random secret $\mathbf{s} \in \mathbb{Z}_q^n$ and error rate $(r')^2 := (r\|\mathbf{x}\|/q)^2 + (rd/q)^2$.*

With this, we can present our modified classical reduction.

Lemma 5.4. *There exists a probabilistic polynomial-time (classical) algorithm with access to an oracle that solves $\mathsf{LWE}_{q,\alpha}$ that takes as input a number $\alpha \in (0,1)$ and an integer $q \geq 2$, a lattice $\mathcal{L} \subset \mathbb{R}^n$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{L})$, and $\mathrm{poly}(n)$ samples from the discrete Gaussian distribution $D_{\mathcal{L},r_i}$ for $\mathrm{poly}(n)$ parameters $r_i \geq r$, solves $\mathsf{BDD}_{\mathcal{L}^*,d}$ for $d := \alpha q/(\sqrt{2}r)$.*

Proof. Let $\kappa = \mathrm{poly}(n)$ with $\kappa \geq 100n^2\ell$ be such that the advantage of our $\mathsf{LWE}_{q,\alpha}$ oracle is at least $1/\kappa$, where $\ell \geq 1$ is the number of samples required by the oracle.

The reduction takes as input an integer $q \geq 2$, a lattice $\mathcal{L} \subset \mathbb{R}^n$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{L})$, samples from $D_{\mathcal{L},r_i}$ for $1 \leq i \leq \mathrm{poly}(n)$, and a coset $\mathbf{x} + \mathcal{L}^*$ where $\|\mathbf{x}\| \leq d$. The goal is to recover $\mathbf{x}$. Without loss of generality, we may assume that $\|\mathbf{x}\| \geq \exp(-n)\lambda_1(\mathcal{L}^*) \geq (q/r) \cdot \exp(-n/2)$ (where the second inequality follows from Lemma 2.6), since otherwise we can find $\mathbf{x}$ efficiently using Lemma 2.12.

The reduction will use its LWE oracle to simulate an oracle $\mathcal{O}$ such that the probability that $\mathcal{O}(\mathbf{z}, t)$ outputs 1 depends only on $\exp(t)\|\mathbf{z} - \mathbf{x}\|$. In other words, $\mathcal{O}$ is an oracle with a "hidden center" $\mathbf{x}$ as in our definition of OHCP (Definition 4.3). We will then use Proposition 4.4 to find good approximations to $\mathbf{x}$, which will allow us to recover $\mathbf{x}$.

On input $(\mathbf{x}, t)$, the oracle $\mathcal{O}$ uses fresh samples from $D_{\mathcal{L},\exp(t)r}$. It then runs the transformation from Lemma 5.3 on these samples, the coset $\mathbf{x} - \mathbf{z} + \mathcal{L}^*$, parameter r, and distance bound d . Let $A_{\mathbf{z},t}$ be the resulting samples. $\mathcal{O}$ then calls the LWE oracle on $A_{\mathbf{z},t}$ and outputs 1 if and only if it accepts.

Using the oracle $\mathcal{O}$, the reduction runs the procedure from Proposition 4.4 with confidence parameter κ and distance bound d, receiving as output some approximation $\mathbf{z}$ to the oracle's "center." Finally, the reduction runs Babai's algorithm (Lemma 2.12) on the coset $\mathbf{x} - \mathbf{z} + \mathcal{L}^*$, receiving as output $\widehat{\mathbf{z}} \in \mathbf{x} - \mathbf{z} + \mathcal{L}^*$, and returns $\widehat{\mathbf{z}} + \mathbf{z}$ as its own output.

The running time is clear. We first assume that $\mathbf{z}$ is a valid solution to $(\exp(-\kappa), \exp(-\kappa))$-OHCP with "hidden center" $\mathbf{x}$ and show that the reduction outputs the correct answer in this case. By definition, we have

$$\|\mathbf{z} - \mathbf{x}\| \leq \exp(-\kappa)d \leq \alpha \exp(-\kappa)/\eta(\mathcal{L}) \leq 2^{-n}\lambda_1(\mathcal{L}^*) \ ,$$

where the last inequality follows from Claim 2.7. So, the procedure from Lemma 2.12 will output exactly $\widehat{\mathbf{z}} = \mathbf{x} - \mathbf{z}$. Therefore, the reduction returns the correct answer, $\mathbf{x} = \widehat{\mathbf{z}} + \mathbf{z}$.

It remains to show that (1) the oracle $\mathcal{O}$ represents a valid instance of $(\exp(-\kappa), \exp(-\kappa))$-OHCP with "hidden center" $\mathbf{z}$; and (2) the oracle additionally satisfies the requirements needed to apply Proposition 4.4.

To prove validity, we simply observe that Lemma 5.3 implies that the distribution of $A_{\mathbf{z},t}$ depends only on $\exp(t)\|\mathbf{z} - \mathbf{x}\|$ (up to negligible statistical distance). Furthermore, by assumption $\|\mathbf{x}\| \geq \exp(-n)d \geq \exp(-\kappa)d$. Therefore, $\mathcal{O}$, κ, and d correspond to a valid instance of $(\exp(-\kappa), \exp(-\kappa))$-OHCP with "hidden center" $\mathbf{x}$, as needed.

It remains to show that $\mathcal{O}$ satisfies the requirements from Proposition 4.4. We write $p(\mathbf{z}, t)$ for the probability that $\mathcal{O}$ outputs 1 on input $(\mathbf{z}, t)$ and $p(\infty)$ for the probability that our LWE oracle outputs 1 on uniformly random input. Again by Lemma 5.3, we have that $A_{\mathbf{z},t}$ is statistically close to LWE samples with error rate $(r/q) \cdot \sqrt{\exp(2t)\|\mathbf{z} - \mathbf{x}\|^2 + d^2}$. In particular, the error rate is α when $\mathbf{z} = \mathbf{0}$ and $t = \log(d/\|\mathbf{x}\|)$,

and it follows that $p(\mathbf{0}, \log(d/\|\mathbf{x}\|)) - p(\infty) \geq 1/\kappa$. I.e., Item 1 holds. For Item 2, notice that by Lemma 2.5, Lemma 2.6, and the union bound, the distribution of $A_{\mathbf{0},t}$ is within statistical distance

$$\min\{1, 2\ell \exp(-\exp(t-n))\} \leq 2\exp(-t/\kappa)$$

of the uniform distribution. It follows that $|p(\mathbf{0}, t) - p(\infty)| \leq 2\exp(-t/\kappa)$, as needed. Finally, for Item 3, notice that by Lemma 2.1 the distributions of $A_{\mathbf{z},t_1}$ and $A_{\mathbf{z},t_2}$ are within statistical distance

$$\min\big\{1, 10\ell(\exp(|t_1 - t_2|) - 1)\big\} \leq \kappa|t_1 - t_2|$$

of each other, where we have applied the union bound over the ℓ samples in $A_{\mathbf{z},t_j}$. Therefore, $p(\mathbf{z}, t)$ is κ-Lipschitz in t, as needed. □

6 Main Theorem

Throughout this section, let K be a number field having s_1 real embeddings and s_2 pairs of complex embeddings, with $n = s_1 + 2s_2$, and denote its ring of integers by $R = \mathcal{O}_K$.

We now define the distribution Υ_α over error distributions that we will use in our reduction. We note that this distribution is slightly different than the one used in [LPR10], and is in particular narrower by a factor of essentially $n^{1/4}$.

Definition 6.1. Fix an arbitrary $f(n) = \omega(\sqrt{\log n})$. For $\alpha > 0$, a distribution sampled from Υ_α is an elliptical Gaussian $D_{\mathbf{r}}$, where $\mathbf{r} \in G$ is sampled as follows: for $i = 1, \ldots, s_1$, sample $x_i \leftarrow D_1$ and set $r_i^2 = \alpha^2(x_i^2 + f^2(n))/2$. For $i = s_1 + 1, \ldots, s_1 + s_2$, sample $x_i, y_i \leftarrow D_{1/\sqrt{2}}$ and set $r_i^2 = r_{i+s_2}^2 = \alpha^2(x_i^2 + y_i^2 + f^2(n))/2$.

Equivalently, r_i^2 is distributed as a shifted chi-squared distribution for real embeddings, and a shifted chi-squared distribution with two degrees of freedom for complex embeddings. The following is the main result of the paper.

Theorem 6.2. *Let K be an arbitrary number field of degree n and $R = \mathcal{O}_K$. Let $\alpha = \alpha(n) \in (0, 1)$, and let $q = q(n) \geq 2$ be an integer such that $\alpha q \geq 2 \cdot \omega(1)$. There is a polynomial-time quantum reduction from K-DGS_γ to (average-case, decision) R-$\mathsf{LWE}_{q,\Upsilon_\alpha}$ for any*

$$\gamma = \max\Big\{\eta(\mathcal{I}) \cdot \sqrt{2}/\alpha \cdot \omega(1),\ \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)\Big\}. \tag{6.1}$$

Using the easy inequality $\eta(\mathcal{I}) > \omega(\sqrt{\log n})/\lambda_1(\mathcal{I}^\vee)$ (Claim 2.7) we get that as long as $\alpha < \sqrt{\log n/n}$ (which is typically the case in applications), the first term in the maximum in Equation (6.1) dominates. Using known reductions (see, e.g., [Reg05, Section 3.3]), we immediately derive the following corollary.

Corollary 6.3. *Let K be an arbitrary number field of degree n and $R = \mathcal{O}_K$. Let $\alpha = \alpha(n) \in (0, 1)$, and let $q = q(n) \geq 2$ be an integer such that $\alpha q \geq \omega(1)$. There is a polynomial-time quantum reduction from K-SIVP_γ to (average-case, decision) R-$\mathsf{LWE}_{q,\Upsilon_\alpha}$ for any*

$$\begin{aligned}\gamma &= \max\Big\{\omega(\sqrt{n}/\alpha) \cdot \eta(\mathcal{I})/\lambda_n(\mathcal{I}),\ \sqrt{2}n/(\lambda_1(\mathcal{I}^\vee)\lambda_n(\mathcal{I}))\Big\} \\ &\leq \max\Big\{\omega(\sqrt{n\log n}/\alpha),\ \sqrt{2}n\Big\}.\end{aligned}$$

We remark that for the families of number fields in [Roq67, PR07], $\lambda_1(\mathcal{I}^\vee)\lambda_n(\mathcal{I}) = \Theta(n)$ and $\eta(\mathcal{I})/\lambda_n(\mathcal{I}) = \Theta(1/\sqrt{n})$ for any ideal $\mathcal{I}$. So, in this special case, we obtain a worst-case approximation factor of $\gamma = \omega(1/\alpha)$. Furthermore, these number fields can potentially allow for applications with constant α, in which case the worst-case approximation factor is essentially constant.

We now proceed to describe the the high-level structure of the proof of Theorem 6.2, which is nearly identical to that of the corresponding statement in [LPR10, Theorem 4.1] (which in turn is very similar to the proof in [Reg05]). As in Section 5, the most significant difference is that we use a direct reduction to the *average-case* Ring-LWE *decision* problem in our classical step, Lemma 6.6. Furthermore, we again require discrete Gaussian samples with various different parameters, but unlike in Section 5, we now need *non-spherical* parameters, as follows.

Definition 6.4. For $r > 0$, $\zeta > 0$, and $T \geq 1$, define $W_{r,\zeta,T}$ as the set of cardinality $(s_1 + s_2) \cdot (T + 1)$ containing for each $i = 1, \ldots, s_1 + s_2$ and $j = 0, \ldots, T$ the vector $\mathbf{r}_{i,j} \in G$ which is equal to r in all coordinates except in the ith, and the $(i + s_2)$th if $i > s_1$, where it is equal to $r \cdot (1 + \zeta)^j$.

The reduction works by repeated applications of the following *iterative step*.

Lemma 6.5. *There exists an efficient quantum algorithm that given an oracle that solves R-$\mathsf{LWE}_{q,\Upsilon_\alpha}$ on input a number $\alpha \in (0, 1)$ and an integer $q \geq 2$, a fractional ideal $\mathcal{I} \subset K$, a number $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$ such that $r' := r \cdot \omega(1)/(\alpha q) > \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)$, polynomially many samples from the discrete Gaussian distribution $D_{\mathcal{I},\mathbf{r}}$ for each $\mathbf{r} \in W_{r,\zeta,T}$ (for some $\zeta = 1/\operatorname{poly}(n)$ and $T = \operatorname{poly}(n)$), and a vector $\mathbf{r}' \in G$ where $\mathbf{r}' \geq r'$, outputs an independent sample from $D_{\mathcal{I},\mathbf{r}'}$.*

Theorem 6.2 follows easily from this iterative step, as we now sketch. We start with a very large value of r, say $r \geq 2^{2n}\lambda_n(\mathcal{I})$, so that samples from $D_{\mathcal{I},\mathbf{r}}$ for each $\mathbf{r} \in W_{r,\zeta,T}$ can be generated classically (see [Reg05, Lemma 3.2]). Then, given those samples, we apply the algorithm of Lemma 6.5 a polynomial number of times with various values of $\mathbf{r}'$ (using the same samples) to obtain a polynomial number of samples from $D_{\mathcal{I},\mathbf{r}}$ for each $\mathbf{r} \in W_{r',\zeta,T}$ where $r' := r \cdot \omega(1)/(\alpha q) \leq r/2$. Repeating this, we obtain samples from progressively narrower and narrower distributions, until we get samples with the desired Gaussian parameter $s \geq \gamma$. Note that the γ given in (6.1) corresponds to values of r, r' satisfying the hypotheses of Lemma 6.5.

The iterative step given by Lemma 6.5 is obtained by combining two reductions, described next in Lemmas 6.6 and 6.7. The first, whose proof is given in Section 6.1, is a reduction from GDP to (average-case, decision) R-LWE, which uses Gaussian samples. This is the main novel component in our paper.

Lemma 6.6. *There exists a probabilistic polynomial-time (classical) algorithm that given an oracle that solves R-$\mathsf{LWE}_{q,\Upsilon_\alpha}$ and input a number $\alpha \in (0, 1)$ and an integer $q \geq 2$ together with its factorization, a fractional ideal $\mathcal{I}$ in K, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$, and polynomially many samples from the discrete Gaussian distribution $D_{\mathcal{I},\mathbf{r}}$ for each $\mathbf{r} \in W_{r,\zeta,T}$ (for some $\zeta = 1/\operatorname{poly}(n)$ and $T = \operatorname{poly}(n)$), solves $\mathsf{GDP}_{\mathcal{I}^\vee,g}$ for any $g = o(1) \cdot \alpha q/(\sqrt{2}r)$.*

The second part is quantum, and is a very slight extension of [Reg05, Lemma 3.14].

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

Proof. The case $\mathbf{r} = \mathbf{1}$ is precisely [LPR10, Lemma 4.4] (which is essentially [Reg05, Lemma 3.14]). We solve the general case by reducing to this special case by rescaling coordinates. Namely, we apply the special

case to the lattice $\mathbf{R}^{-1}\mathcal{L}$ and g where $\mathbf{R} = \mathrm{diag}(\mathbf{r})$, and obtain as output a sample from $D_{\mathbf{R}^{-1}\mathcal{L},1/(2g)}$, which by left-multiplying by $\mathbf{R}$ gives us the desired sample from $D_{\mathcal{L},\mathbf{r}/(2g)}$. The special case also requires an oracle that solves $\mathsf{GDP}_{\mathbf{R}\mathcal{L}^*,g}$. We construct such an oracle from our given $\mathsf{GDP}_{\mathcal{L}^*,g}$ oracle $\mathcal{O}$ in the following way: on input a coset $\mathbf{z} + \mathbf{R}\mathcal{L}^*$ where $\mathbf{z} \leftarrow D_g$, sample a point $\mathbf{y}$ from $D_{\mathbf{r}'g}$ where $\mathbf{r}' := (\mathbf{1} - \mathbf{r}^{-2})^{1/2}$, then call $\mathcal{O}$ with the coset $\mathbf{y} + \mathbf{R}^{-1}\mathbf{z} + \mathcal{L}^*$ and return its output multiplied by $\mathbf{R}$. It is easy to check that the input to $\mathcal{O}$ is properly distributed, and that our oracle works as needed. □

6.1 Proof of Lemma 6.6

In this section we prove Lemma 6.6, providing a reduction from GDP on any ideal lattice in any number field to a corresponding Ring-LWE decision problem. (The reduction uses discrete Gaussian samples over the dual ideal.) We adopt the notation from Section 2.3 for the ring embeddings $\sigma_i \colon K \to \mathbb{C}$ for $i = 1, \ldots, n = s_1 + 2s_2$, the (fractional) codifferent ideal $R^\vee \subset K$, and the discriminant Δ_K of the number field K.

As our starting point, the following restatement and slight generalization of [LPR10, Lemma 4.7] describes a transformation from an instance of BDD on an ideal lattice to Ring-LWE samples, which uses discrete Gaussian samples over the dual ideal. The generalization considers samples from an *elliptical* discrete Gaussian, rather than just a spherical one. The proof is nearly identical, except that in [LPR10, Lemma 4.8], one needs to replace the use of [Reg05, Claim 3.9] with the more general [Pei10, Theorem 3.1] in order to analyze the sum of an elliptical discrete Gaussian and a continuous Gaussian. (In addition, for convenience the lemma says that the Ring-LWE secret is uniformly random; this is achieved using the standard technique of randomizing the secret. See, e.g., [Reg10, Lemma 3.2].)

Lemma 6.8. *There is an efficient algorithm that takes as input an integer $q \geq 2$ with known factorization, a fractional ideal $\mathcal{I}^\vee \subset K$, a coset $e + \mathcal{I}^\vee$ and bound $d \geq \|e\|_\infty = \max_i |\sigma_i(e)|$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$, and samples from $D_{\mathcal{I},\mathbf{r}}$ for some $\mathbf{r} \geq r$. It outputs samples that are within negligible statistical distance of the Ring-LWE distribution $A_{s,\mathbf{r}'}$ for a uniformly random $s \in R_q^\vee$, where the coordinates of $\mathbf{r}'$ are given by $(r_i')^2 := (r_i|\sigma_i(e)|/q)^2 + (rd/q)^2$.*

In particular, notice that if we sample a coset $e + \mathcal{I}^\vee$ as in $\mathsf{GDP}_{\mathcal{I}^\vee,g}$ with $g := \alpha q/(\sqrt{2}r)$ and then apply Lemma 6.8 with $d := \alpha q f(n)/(\sqrt{2}r)$, then the distribution of the resulting error rate $\mathbf{r}'$ is distributed exactly Υ_α, as defined in Definition 6.1. This is the reason that we work with Υ_α.

We will also use the following lemma, which says that any elliptical Gaussian whose parameters' product is sufficiently large is "smooth" modulo an ideal.

Lemma 6.9. *For any fractional ideal $\mathcal{I} \subset K$ and $\mathbf{r} \in G$ where*

$$c := \Big(\prod_{i=1}^{n} r_i\Big)^{1/n} \cdot (\mathrm{N}(\mathcal{I}) \cdot \Delta_K)^{-1/n} \geq 1, \tag{6.2}$$

we have $\mathbf{r} \geq \eta_\varepsilon(\mathcal{I})$ for $\varepsilon = \exp(-c^2 n)$.

Our particular case of interest is $\mathcal{I} = R^\vee$, where $\mathrm{N}(\mathcal{I}) \cdot \Delta_K = 1$ and so $c = (\prod_i r_i)^{1/n}$.

Proof. Let $\mathbf{R} = \mathrm{diag}(\mathbf{r})$ and $\mathcal{L} = \mathbf{R}^{-1} \cdot \sigma(\mathcal{I})$, so $\mathcal{L}^* = \mathbf{R} \cdot \sigma(\mathcal{I})^*$ and any nonzero $\mathbf{w} \in \mathcal{L}^*$ has the form $\mathbf{R} \cdot \overline{\sigma(w)}$ for some nonzero $w \in \mathcal{I}^\vee$. By the inequality of arithmetic and geometric means and the fact that

$\prod_i |\sigma_i(w)| = |\mathrm{N}(w)| \geq \mathrm{N}(\mathcal{I}^\vee) = (\mathrm{N}(\mathcal{I}) \cdot \Delta_K)^{-1}$, we have

$$\|\mathbf{w}\|^2 = \sum_i r_i^2 \cdot |\sigma_i(w)|^2 \geq n \left(\prod_i r_i^2 \cdot |\sigma_i(w)|^2 \right)^{1/n} \geq c^2 n \ ,$$

so $\lambda_1(\mathcal{L}^*) \geq c\sqrt{n}$. Lemma 2.6 then implies that $1 \geq \eta_\varepsilon(\mathcal{L})$, or equivalently, $\mathbf{r} \geq \eta_\varepsilon(\mathcal{I})$. □

Proof of Lemma 6.6. If $\alpha \leq \exp(-n)$, then with high probability the coset representative e from the GDP instance will satisfy

$$\|\sigma(e)\| \leq \sqrt{n} g \leq \alpha\sqrt{n}/\eta(\mathcal{I}) \leq 2^{-n}\lambda_1(\mathcal{I}^\vee) \ ,$$

where the last inequality follows from Claim 2.7. In this case, the problem can be solved efficiently using Babai's algorithm (Lemma 2.12). So, we may assume that $\alpha > \exp(-n)$. Furthermore, by Proposition 3.3, it suffices to solve $\mathsf{GDP}_{\mathcal{I}^\vee, g'}$ with non-negligible probability, where $g' := \alpha q/(\sqrt{2}r)$. Finally, we let $\kappa = \mathrm{poly}(n)$ with $\kappa \geq 100n^2\ell$ be such that the advantage of our R-$\mathsf{LWE}_{q,\Upsilon_\alpha}$ oracle is at least $2/\kappa$, where $\ell \geq 1$ is the number of samples required by the oracle.

The reduction takes as input an integer $q \geq 2$ with known factorization, a fractional ideal $\mathcal{I} \subset K$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$, samples from the distributions $D_{\mathcal{I},\mathbf{r}}$ for each $\mathbf{r} \in W_{r,\zeta,T} = \{\mathbf{r}_{i,j}\}$ (Definition 6.4), and a coset $e + \mathcal{I}^\vee$ where $e \leftarrow D_{g'}$. The goal is to recover e. The reduction will use its R-LWE oracle to simulate oracles $\mathcal{O}_i \colon \mathbb{R} \times \mathbb{R}^{\geq 0} \to \{0,1\}$ for $1 \leq i \leq s_1$ and $\mathcal{O}_i \colon \mathbb{C} \times \mathbb{R}^{\geq 0}$ for $s_1 < i \leq s_1 + s_2$ such that the probability that $\mathcal{O}_i(z,t)$ outputs 1 depends only on $\exp(t)|z - \sigma_i(e)|$. In other words, $\mathcal{O}_i$ is an oracle with a "hidden center" $\sigma_i(e)$ as in our definition of OHCP (Definition 4.3), with $k = 1$ for $1 \leq i \leq s_1$ and $k = 2$ for $s_1 < i \leq s_1 + s_2$. (Here, we assume that the oracles corresponding to complex embeddings implicitly identify $\mathbf{z} \in \mathbb{R}^2$ with $z \in \mathbb{C}$ in the natural way.) We will then use Proposition 4.4 to find good approximations to $\sigma_i(e)$ for each i, which will allow us to recover e.

For this purpose, for $1 \leq i \leq s_1$ define $k_i \colon \mathbb{R} \to K_\mathbb{R}$ as $k_i(z) = \sigma^{-1}(z \cdot \mathbf{e}_i)$, and for $s_1 < i \leq s_1 + s_2$ define $k_i \colon \mathbb{C} \to K_\mathbb{R}$ as $k_i(z) = \sigma^{-1}(z \cdot \mathbf{e}_i + \overline{z} \cdot \mathbf{e}_{i+s_2})$, where $\mathbf{e}_i \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ has 1 in the ith coordinate and 0 elsewhere. On input (z,t), the oracle $\mathcal{O}_i$ uses fresh samples from $D_{\mathcal{I},\mathbf{r}_{i,j}}$, where $(1+\zeta)^j = \exp(t)$. It then runs the transformation from Lemma 6.8 on these samples, the coset $e - k_i(z) + \mathcal{I}^\vee$, parameter r, and distance bound $d := g' f(n)/\sqrt{2}$ for some $f(n) = \omega(\sqrt{\log n})$ with $f(n) \leq n$. Let $A_{i,z,t}$ be the resulting samples. $\mathcal{O}_i$ then calls the LWE oracle on $A_{i,z,t}$ and outputs the 1 if and only if it accepts.

Using the oracles $\mathcal{O}_i$, the reduction runs the procedure from Proposition 4.4 for each $i = 1, \ldots, s_1 + s_2$, with oracle $\mathcal{O}_i$, confidence parameter κ, and distance bound d, receiving as output some approximation z_i to the oracle's "center." Finally, the reduction runs Babai's algorithm (Lemma 2.12) on the coset $e - \sum k_i(z_i) + \mathcal{I}^\vee$, receiving as output $\widehat{e} \in K$, and returns $\widehat{e} + \sum k_i(z_i)$ as its own output.

The running time is clear. We first assume that the z_i are valid solutions to $(\exp(-\kappa), \exp(-\kappa))$-OHCP with "hidden center" $\sigma_i(e)$ and show that the reduction outputs the correct answer in this case. Since the z_i are valid solutions, by definition we have

$$|z_i - \sigma_i(e)| \leq \exp(-\kappa)d = \exp(-\kappa)\alpha q f(n)/(\sqrt{2}r) \leq \exp(-\kappa) f(n)/\eta(\mathcal{I}) \leq 2^{-n}\lambda_1(\mathcal{I}^\vee)/\sqrt{n} \ ,$$

where the last inequality follows from Claim 2.7. So, $\|e - \sum k_i(z_i)\| \leq 2^{-n}\lambda_1(\mathcal{I}^\vee)$, and the procedure from Lemma 2.12 will therefore output exactly $\widehat{e} = e - \sum k_i(z_i)$. Therefore, the reduction returns the correct answer, $e = \widehat{e} + \sum k_i(z_i)$.

It remains to show that, with non-negligible probability over the choice of e, for all i (1) $\mathcal{O}_i$ represents valid instances of $(\exp(-\kappa), \exp(-\kappa))$-OHCP with "hidden center" $\sigma_i(e)$; and (2) $\mathcal{O}_i$ satisfies the requirements needed to apply Proposition 4.4.

To prove validity, we first observe that Lemma 6.8 implies that the distribution of $A_{i,z,t}$ depends only on $\exp(t)|z - \sigma_i(e)|$ (up to negligible statistical distance). And, since e was drawn from $D_{g'}$ and $d = \omega(\sqrt{\log n}) \cdot g'$, we have

$$\exp(-\kappa)d \le |\sigma_i(e)| \le d \tag{6.3}$$

for all i except with negligible probability. Therefore, $\mathcal{O}_i$, κ, and d correspond to a valid instance of $(\exp(-\kappa), \exp(-\kappa))$-OHCP with "hidden center" $\sigma_i(e)$, except with negligible probability.

We now analyze the oracles $\mathcal{O}_i$. We write $p_i(z,t)$ for the probability that $\mathcal{O}_i$ outputs 1 on input (z,t) and $p(\infty)$ for the probability that our R-LWE oracle outputs 1 on uniformly random input. Notice that $p_i(0,0) = p_j(0,0)$ for all i, j, and $p_i(0,0) - p(\infty)$ is exactly the advantage that our R-LWE oracle has against the error rate that we get from applying Lemma 6.8 to $e + \mathcal{I}^\vee$ with the parameters chosen above. Furthermore, recall that when e is sampled as in $\mathsf{GDP}_{\mathcal{I}^\vee, g'}$, the resulting error rate is distributed exactly as Υ_α. Since our oracle has advantage $2/\kappa$ against this distribution of error rate, by a standard Markov argument, we may assume that $p_i(0,0) - p(\infty) \ge 1/\kappa$, which must hold with non-negligible probability.

We can now show that all of the items necessary to apply Proposition 4.4 hold, which will complete the proof. Indeed, we have already shown Item 1, which asks that $p_i(0,0) - p(\infty) \ge 1/\kappa$. For Item 2, notice that by Lemma 6.9, Lemma 2.5, and the union bound, for $t \ge \kappa/10$, the distribution of $A_{i,0,t}$ is within statistical distance

$$\ell \exp\Big(-n \exp(2t/n) r^2 \prod_j |\sigma_j(e)|^{2/n}\Big) \le \ell \exp\big(-n \exp(2t/n - 4n - 1) q^2\big) \le 2\exp(-t/\kappa)$$

of the uniform distribution, where we have used Equation (6.3) to get $|\sigma_j(e)| \ge \exp(-n)g' = \exp(-n)\alpha q/(\sqrt{2}r) > \exp(-2n-1)q/r$ except with negligible probability. It follows that $|p_i(0,t) - p(\infty)| \le 2\exp(-t/\kappa)$, as needed. Finally, for Item 3, notice that by Lemma 2.1 the distributions of A_{i,z,t_1} and A_{i,z,t_2} are within statistical distance

$$\min\big\{1, 10\ell(\exp(|t_1 - t_2|) - 1)\big\} \le \kappa |t_1 - t_2|$$

of each other, where we have applied the union bound over the ℓ samples in A_{i,z,t_j}. Therefore, $p_i(z,t)$ must be κ-Lipschitz in t, as needed. □

7 Spherical Error

The goal of this section is to extend the main theorem, Theorem 6.2, to the case of spherical Gaussian error (see Corollary 7.3). This would follow easily from Lemma 7.2, showing how to reduce LWE with non-spherical error into LWE with spherical error, with a loss in approximation factor depending on the number of samples used by the latter. The lemma is very similar to [LPR10, Lemma 5.16], and is included here for completeness. We start with the following claim from [LPR10].

Claim 7.1. *Let $r_1, \ldots, r_n \in \mathbb{R}^+$ and $s_1, \ldots, s_n \in \mathbb{R}^+$ be such that for all i, $|s_i/r_i - 1| < \sqrt{\log n/n}$. Then any set $A \subseteq \mathbb{R}^n$ whose measure under the Gaussian distribution $D_{r_1} \times \cdots \times D_{r_n}$ is non-negligible, also has non-negligible measure under $D_{s_1} \times \cdots \times D_{s_n}$.*

In the lemma below, we use R-LWE with a fixed error distribution, as opposed to a distribution over error distributions as in Lemma 2.15.

Lemma 7.2 (Worst-case to average-case with spherical error). *There is a randomized polynomial-time algorithm that given any $\alpha > 0$ and $\ell \ge 1$, as well as an oracle that solves R-LWE_{q,D_ξ} given only ℓ samples,*

where $\xi = \alpha(n\ell/\log(n\ell))^{1/4}$*, solves* R-$\mathsf{LWE}_{q,D_{\mathbf{r}}}$ *for any (possibly unknown)* $\mathbf{r}$ *satisfying that all* r_i *are in* $[0, \alpha]$.

Proof. For $e_1, \ldots, e_\ell \in \mathbb{T}$, consider the transformation mapping ℓ samples $(a_i, b_i)_{i=1}^{\ell}$ to $(a_i, b_i + e_i)_{i=1}^{\ell}$. Then it is easy to see that for all $s \in R_q^\vee$, ψ, $\mathbf{r}'$, if we sample from $(A_{s,\psi})^\ell$ (i.e., ℓ independent samples from $A_{s,\psi}$) and apply this transformation with $e_1, \ldots, e_\ell$ chosen independently from $D_{\mathbf{r}'}$, then the output distribution (averaged over the choice of $e_1, \ldots, e_\ell$) is $(A_{s,\psi+D_{\mathbf{r}'}})^\ell$.

The reduction repeats the following a polynomial number of times. Choose $e_1, \ldots, e_\ell$ independently from D_ξ. Then estimate the acceptance probability of the oracle on the following two input distributions: the first is obtained by applying the above transformation with $e_1, \ldots, e_\ell$ to our input samples (each time with fresh inputs); the second is the uniform distribution $(R_q \times \mathbb{T})^\ell$. If in any of these polynomial number of attempts a non-negligible difference is observed between the two acceptance probabilities, output "non-uniform" (accept); otherwise output "uniform" (reject).

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

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

it follows from Claim 7.1 that the measure of S under $(D_\xi)^\ell$ is also non-negligible, and we are done. □

Recalling the definition of Υ_α from Definition 6.1 and noting that with high probability all the coordinates of its error rate $\mathbf{r}$ are at most $\alpha \cdot \omega(\sqrt{\log n})$, we immediately obtain the following hardness result for average-case, decision Ring-LWE with spherical error.

Corollary 7.3. *With notation as in Theorem 6.2, there is a polynomial-time quantum reduction from* K-DGS_γ *to the (average-case, decision) problem of solving* R-LWE_{q,D_ξ} *using* ℓ *samples, where*

$$\gamma = \max\Big\{\eta(\mathcal{I}) \cdot (\sqrt{2}/\xi) \cdot (n\ell/\log(n\ell))^{1/4} \cdot \omega(\sqrt{\log n}),\ \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)\Big\}.$$

References

[ACPS09] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618. 2009.

[ADPS16] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange - a new hope. In *USENIX Security Symposium*, pages 327–343. 2016.

[Bab85] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.

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

[BCD+16] J. W. Bos, C. Costello, L. Ducas, I. Mironov, M. Naehrig, V. Nikolaenko, A. Raghunathan, and D. Stebila. Frodo: Take off the ring! practical, quantum-secure key exchange from LWE. In *CCS*, pages 1006–1018. 2016.

[BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. *TOCT*, 6(3):13, 2014. Preliminary version in ITCS 2012.

[BLP+13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. 2013.

[Bor85] C. Borell. Geometric bounds on the Ornstein-Uhlenbeck velocity process. *Z. Wahrsch. Verw. Gebiete*, 70(1):1–13, 1985.

[BS16] J. Biasse and F. Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In *SODA*, pages 893–902. 2016.

[BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput.*, 43(2):831–871, 2014. Preliminary version in FOCS 2011.

[CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT*, pages 559–585. 2016.

[CDW17] R. Cramer, L. Ducas, and B. Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In *EUROCRYPT*. 2017. To appear.

[CGS14] P. Campbell, M. Groves, and D. Shepherd. Soliloquy: a cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. Available at `http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf`.

[EHL14] K. Eisenträger, S. Hallgren, and K. E. Lauter. Weak instances of PLWE. In *SAC*, pages 183–194. 2014.

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

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

[GSW13] C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, pages 75–92. 2013.

[GVW13] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554. 2013.

[Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.

[LLM06] Y.-K. Liu, V. Lyubashevsky, and D. Micciancio. On bounded distance decoding for general lattices. In *APPROX-RANDOM*, pages 450–461. 2006.

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

[LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in Eurocrypt 2010.

[LS15] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015.

[MM11] D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *CRYPTO*, pages 465–484. 2011.

[MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718. 2012.

[MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.

[O'D14a] R. O'Donnell. *Analysis of Boolean functions*. Cambridge University Press, New York, 2014. ISBN 978-1-107-03832-5. doi:10.1017/CBO9781139814782.

[O'D14b] R. O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

[Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009.

[Pei10] C. Peikert. An efficient and parallel Gaussian sampler for lattices. In *CRYPTO*, pages 80–97. 2010.

[Pei14] C. Peikert. Lattice cryptography for the Internet. In *PQCrypto*, pages 197–219. 2014.

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

[PW08] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. *SIAM J. Comput.*, 40(6):1803–1844, 2011. Preliminary version in STOC 2008.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.

[Reg10] O. Regev. The learning with errors problem (invited survey). In *IEEE Conference on Computational Complexity*, pages 191–204. 2010.

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

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

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