Lattices that Admit Logarithmic Worst-Case to Average-Case Connection Factors

Chris Peikert* Alon Rosen†

April 13, 2007

Abstract

We demonstrate an *average-case* problem that is as hard as finding $\gamma(n)$-approximate shortest vectors in certain n-dimensional lattices in the *worst case*, where $\gamma(n) = O(\sqrt{\log n})$. The previously best known factor for any non-trivial class of lattices was $\gamma(n) = \tilde{O}(n)$.

Our results apply to families of lattices having special algebraic structure. Specifically, we consider lattices that correspond to *ideals* in the ring of integers of an algebraic number field. The worst-case problem we rely on is to find approximate shortest vectors in these lattices, under an appropriate form of preprocessing of the number field.

For the connection factors $\gamma(n)$ we achieve, the corresponding *decision* problems on ideal lattices are *not* known to be NP-hard; in fact, they are in P. However, the *search* approximation problems still appear to be very hard. Indeed, ideal lattices are well-studied objects in computational number theory, and the best known algorithms for them seem to perform *no better* than the best known algorithms for general lattices.

To obtain the best possible connection factor, we instantiate our constructions with infinite families of number fields having constant *root discriminant*. Such families are known to exist and are computable, though no efficient construction is yet known. Our work motivates the search for such constructions. Even constructions of number fields having root discriminant up to $O(n^{2/3-\epsilon})$ would yield connection factors better than $\tilde{O}(n)$.

As an additional contribution, we give reductions between various worst-case problems on ideal lattices, showing for example that the shortest vector problem is no harder than the closest vector problem. These results are analogous to previously-known reductions for general lattices.

1 Introduction

In 1996, Ajtai established a remarkable connection between the worst-case and average-case complexity of certain computational problems on lattices [4]. He showed that approximating the length of the shortest nonzero vector in n-dimensional lattices to within a certain *connection factor* $\gamma(n) = \text{poly}(n)$ in the worst case reduces to solving a related problem on the average. This result opened the door to basing cryptography on a worst-case assumption, namely, that no efficient algorithm can approximately solve the shortest vector problem (SVP) to within a $\gamma(n)$ factor on *every* lattice of dimension n.

It seems reasonable to assume that SVP is indeed hard. The problem dates back over 150 years, and it has been heavily scrutinized ever since. A major breakthrough occurred in 1982, when

*SRI International, cpeikert@alum.mit.edu

†Harvard CRCS, SEAS, alon@eecs.harvard.edu

Lenstra, Lenstra, and Lovász designed an efficient algorithm that approximates SVP to within a $2^{O(n)}$ factor [31] (which was later improved to $2^{O(n(\log\log n)^2/\log n)}$ by Schnorr [47]). While the so-called LLL algorithm has proved to be useful in many diverse applications, its approximation factors are much too large to undermine the hardness assumption associated with Ajtai's result. To date, the asymptotically best algorithm for SVP offers a trade-off between solution quality and running time, and for any poly(n) approximation factor still requires time exponential in n [7].

In Ajtai's original reduction, the connection factor $\gamma(n)$ was a rather large polynomial in n. Due to the known time/quality tradeoffs, it is desirable from both a theoretical and practical point of view to obtain reductions for as small of a $\gamma(n)$ as possible. This was a main goal of several follow-up works [16, 37], resulting in the currently best known connection factor of $\gamma(n) = \tilde{O}(n)$ by Micciancio and Regev [38].

The SVP also has the interesting property that it is NP-hard for small approximation factors (under randomized reductions). Ajtai first showed hardness for its exact version (in the Euclidean norm) [5]. This result was improved in a series of works to hardness for any constant approximation factor unless NP $\not\subseteq$ RP [17, 34, 30], and for almost-polynomial factors $2^{\log^{1-\epsilon} n}$ unless NP $\not\subseteq$ RTIME($2^{\mathrm{polylog}(n)}$) [30, 29]. The latter results already approach the perceived limits on the hardness of approximating SVP, as NP-hardness beyond $O(\sqrt{n})$ factors would imply that NP $\subseteq$ coNP [3] (or NP $\subseteq$ coAM, for factors beyond $O(\sqrt{n/\log n})$ [22]).

In light of the above, improving the worst-case/average-case connection factor to $\gamma(n) = n^{1/2-\epsilon}$ appears problematic. In particular, it would imply cryptographic functions based on problems not known to be in coNP or coAM; as shown by Akavia *et al*, this would require significantly new ideas [8]. Going even further to, say, $\gamma(n) = \mathrm{polylog}(n)$ would imply cryptography based on quasi-NP-hardness, a feat which appears far beyond our current capabilities. (Though see [26] for an interesting first step in this direction.)

Note that all the perceived barriers to improving the connection factor are based on the complexity of the *decision* version of SVP. However, Ajtai's reduction actually solves certain *search* problems on lattices. So one potential route to tighter connection factors would be to identify a suitable worst-case problem whose search version appears hard, *but whose decision version is easy*. Doing so would render the perceived barriers vacuous, while preserving or even improving the concrete hardness of the average-case problem.

1.1 Our Results

We put forward a new class of lattices that admit *very small* worst-case to average-case connection factors. These lattices correspond to certain algebraic structures, namely *ideals* in the ring of integers of a suitable algebraic number field. Our worst-case problem is to find approximate shortest vectors in such lattices, under an appropriate form of preprocessing of the number field.

For the connection factors we achieve, the corresponding *decision* problems on these lattices are not known to be NP-hard; in fact, they are in P. However, the *search* approximation problems still appear to be very hard. Indeed, the best known algorithms for these special lattices seem to perform no better than the best known algorithms for general lattices.

The high-level structure of our worst-case/average-case reduction inherits from a sequence of works starting with Ajtai's original paper [4] and the improvements proposed by Micciancio and Regev [38], as well as the works of Micciancio [36], Peikert and Rosen [42], and Lyubashevsky and Micciancio [33]. The latter works generalized the role of the integers $\mathbb{Z}$ in prior reductions, replacing them with some "larger" ring to obtain efficient cryptographic primitives. We show

that by substituting $\mathbb{Z}$ with a ring of *algebraic integers*, one can also obtain significantly better connection factors. Our analysis identifies the *root discriminant* of the number field as the main quantity governing this improvement.

Informal Theorem. *Let K be a number field of degree n having root discriminant $\mathcal{D}_K$. Then there exists an average-case problem which is as hard as finding a γ-approximate shortest nonzero vector in* any *ideal lattice over K, where $\gamma \sim \mathcal{D}_K^{1.5} \cdot \sqrt{\log n}$.*

It is a known fact of algebraic number theory that there exist *computable* infinite families of number fields (of increasing degree n) having *constant* root discriminant [46], though we do not know of any *efficient* construction. In lattices defined over these families, therefore, we obtain a connection factor of $O(\sqrt{\log n})$. More generally, any family of number fields whose root discriminants are as large as $O(n^{2/3-\epsilon})$ yield lattices admitting connection factors better than $\tilde{O}(n)$.

1.2 Explicitness

As far as we are aware, it is still unknown how to *efficiently* compute families of number fields having very small root discriminant. A review of the literature suggests that a fair amount of attention has been devoted to searching for number fields having highly-optimized root discriminants for small *fixed* degrees (see, e.g. [19]). To our knowledge, the problem of efficiently constructing good *asymptotic* families of number field has not received nearly as much attention. The best construction we know of is an infinite family of cyclotomic number fields having root discriminants as small as $O(n(\log\log n)/(\log n))$ [48]. As mentioned above, families having root discriminants even up to $O(n^{2/3-\epsilon})$ would yield improved connection factors for ideal lattices.

We remark that, while number fields having small root discriminant appear to be rare, one can efficiently compute the root discriminant of any given number field. Therefore it is easy to recognize a good number field once it is discovered.

1.3 Uniformity

Our reductions require a small amount of non-uniform advice, which is simply a form of preprocessing: the computational problems are parameterized by some fixed choice of number field, and the non-uniform advice depends only on this choice (not on the input instance). Preprocessing is a standard notion for computational problems over codes and lattices [14, 35, 21], and it seems to be the proper way of stating problems in our setting, given that in real applications the number fields will be chosen well in advance of any particular problem instance. We remark that preprocessing does not seem to help solve our worst-case problems.

A certain amount of advice about the number field also seems necessary for obtaining useful cryptographic hardness, e.g. collision-resistant hash functions. The reason is that we need a way to map inputs of the cryptographic function to "short" algebraic integers. On the face of it, computing such a mapping appears to require some particular information about the number field. See Section 10 for further discussion.

Note that explicit constructions of number fields may actually come with the required advice "by design," removing the non-uniformity from our reductions entirely and enabling cryptographic hardness. This is indeed the case for the cyclotomic number fields mentioned above.

1.4 The Worst-Case Problem

Our results are based on the worst-case problem of *finding* a relatively short non-zero vector in any ideal lattice over certain families of number fields of increasing degree, allowing for arbitrary preprocessing of the number fields. The vector must be short relative to *every* ℓ_p length.

Due to the algebraic structure of ideal lattices, it is actually trivial to closely approximate the *length* λ_1 of a shortest nonzero vector. This is because in ideal lattices, λ_1 is always within a $\sqrt{\mathcal{D}_K}$ factor of Minkowski's upper bound (where $\mathcal{D}_K$ is the root discriminant of the number field), and this bound is easily computed.[1] However, it does not appear that this makes the search problem any easier. In particular, the search and decision problems are not known to be equivalent (and for general lattices, they are only known to be equivalent when the decision problem is NP-hard).

Finding short vectors in ideal lattices over number fields is a long-standing open problem in algebraic number theory, and is considered to be one of the motivations for the development of the celebrated LLL algorithm [31]. The problem also plays a role in the Number Field Sieve factoring algorithm and in "ideal reduction," which is, for example, an essential step in the computation of the unit group and class group of a number field (this is a reason why the recent quantum algorithm of Hallgren [28] is limited to *fixed* degree). Any efficient algorithm for finding short vectors in ideal lattices in the worst case would be considered a major breakthrough in computational number theory [48, 12].

It is hard to qualitatively compare our worst-case/average-case connections with those known for general lattices. On the one hand, our worst-case problem is restricted to a subclass of lattices and hence could be seen as potentially easier than the problems on general lattices. On the other hand, our connection factor is substantially smaller, hence our reduction could be seen as solving a potentially harder problem.

1.5 Additional Contributions

We additionally give relationships between various worst-case problems on ideal lattices. Specifically, we establish approximation-preserving reductions (in any ℓ_p length) from the shortest vector problem (SVP) to the closest vector problem (CVP), and from the exact search version of CVP to the corresponding decision version. Analogous results were already known for general lattices [24], but these reductions do not preserve the "ideal" structure of their inputs. (That is, the instances generated by the reduction are not necessarily ideal lattices, even if the input lattice is ideal.) Our new reductions rely crucially on the splitting behavior of integer primes over number fields.

We give bounds on many standard lattice quantities for ideal lattices, including the successive minima, basis minimum, and covering radius, in arbitrary ℓ_p lengths. We also give a new bound on the smoothing parameter which, for ideal lattices over number fields with small root discriminant, is significantly stronger than a prior bound [38].

Finally, we point out that number fields having constant root discriminant give rise to a collection of lattices that exemplify the tightness of known *transference theorems* [10, 11], *simultaneously* in all ℓ_p lengths, up to constant factors (or $O(\sqrt{\log n})$ factors for $p = 1, \infty$). This provides a more general alternative to a similar family of lattices by Conway and Thompson [39], which are tight for the ℓ_2 length.

[1] The structure of ideal lattices also makes it easy to estimate the value of several other lattice parameters, such as the successive minima λ_i and the covering radius μ.

2 Overview of Techniques

2.1 Ajtai's Framework

Like almost all prior works on worst-case/average-case reductions for lattice problems [23, 16, 37, 43], we follow the framework initiated by Ajtai [4]. This framework shows how to reduce worst-case lattice problems to finding "small" nonzero solutions to random linear equations over certain additive groups G. Specifically, the equations to be solved are of the form $\sum_j a_j \cdot z_j = 0 \in G$, for independent and uniformly-random elements $a_j \in G$ and unknown variables $z_j \in \mathbb{Z}$.

To find a short vector in an arbitrary input lattice Λ, the reduction samples random $a_j \in G$ in a way that is related to Λ. The core idea is a method of sampling $a_j \in G$ together with a short *"offset" vector* $\mathbf{d}_j \in \mathbb{R}^n$. The crucial property of the sampling procedure is this: for any z_j's satisfying the equation $\sum a_j \cdot z_j = 0$, the vector $\mathbf{v} = \sum \mathbf{d}_j \cdot z_j \in \mathbb{R}^n$ is a vector in the input lattice Λ (and is non-zero with significant probability). In particular, when the coefficients z_j are *small*, the resulting vector $\mathbf{v}$ is relatively *short*. The length of $\mathbf{v}$, and consequently the connection factor of the reduction, is therefore governed by two main quantities: (1) the "effective size" of the average-case solution (i.e., the amount by which it expands the offset vectors), and (2) the lengths of the offset vectors $\mathbf{d}_j$ themselves.

In Ajtai's original work [4], the group G is $\mathbb{Z}_q^n$, for an appropriate $q = \text{poly}(n)$. That is, the average-case problem is to find small integers z_j such that $\sum \mathbf{a}_j \cdot z_j = 0 \bmod q$, for $\mathbf{a}_j$ chosen uniformly from $\mathbb{Z}^n$ modulo q. In Ajtai's reduction, the effective size of the average-case solution is a small polynomial in n, and the length of the offset vectors is a $\text{poly}(n)$ factor larger than the nth successive minimum λ_n of Λ. This results in polynomial connection factors for several worst-case lattice problems.

In later work, Micciancio and Regev [38] proposed an elegant method of implementing the sampling procedure, which improved both the effective size of the solution and the lengths of the offset vectors. Their method relies on adding n-dimensional Gaussian noise to "blur" the lattice, destroying its discrete structure. The amount of noise required to completely blur the lattice is called the *smoothing parameter*, which was shown to be bounded by $\sim \lambda_n$. Based on this bound, the sampling procedure produces offset vectors of length $\sim \sqrt{n} \cdot \lambda_n$, and it guarantees an average-case solution of effective size $\sim \sqrt{n}$. The reduction therefore produces lattice points of length $\sim n \cdot \lambda_n$. Using several additional ideas, these techniques can also be used to approximate the *length* λ_1 of the shortest vector to within a factor $\sim n$.

2.2 Our Approach

We retain all the essential elements of Ajtai's framework and its subsequent improvements, but in a more general setting. The core of our approach is to replace the ring of integers $\mathbb{Z}$ with the ring of *algebraic integers* $\mathcal{O}_K$ contained in a *number field* K of degree n. (The idea of replacing $\mathbb{Z}$ with some "larger" ring is rooted in work of Micciancio [36], whose motivation was improved cryptographic efficiency. See Section 3 for details and a comparison to our work.)

A number field is a certain kind of field extension of the rationals $\mathbb{Q}$. Every number field K contains a discrete ring $\mathcal{O}_K$, called the algebraic integers, which acts as an analog of the integers $\mathbb{Z}$ in $\mathbb{Q}$. Strictly speaking, K (and in particular, $\mathcal{O}_K$) is a subset of the complex numbers $\mathbb{C}$. But K also corresponds very naturally to an n-dimensional geometric space via its *canonical embedding*. In this manner, elements in K can be viewed as vectors which can be added and multiplied, and

which have magnitudes in various ℓ_p lengths.

In Ajtai's framework, replacing $\mathbb{Z}$ with $\mathcal{O}_K$ alters the type of lattices underlying the worst-case problem. Instead of general lattices, which are made up of all the $\mathbb{Z}$-combinations of some set of basis vectors, we end up with all $\mathcal{O}_K$-*combinations* of some basis elements in $\mathcal{O}_K$. Specifically, we get a set $\{\sum c_i \cdot b_i \ : \ c_i \in \mathcal{O}_K\}$ for some collection of elements $b_i \in \mathcal{O}_K$, which is called an *ideal* in $\mathcal{O}_K$. When embedded in geometric space, an ideal yields an n-dimensional *ideal lattice*. This notion is quite natural and standard in algebraic number theory; see e.g. [20, Chapter 8], [13].

Our average-case problem is defined in a similar way. An instance is given by a vector $\mathbf{a} = (a_1, \ldots, a_m) \in \mathcal{O}_K^m$ where the a_j are uniform and independent representatives from $\mathcal{O}_K$ modulo q. The problem is to find a "small" solution $\mathbf{z} = (z_1, \ldots, z_m) \in \mathcal{O}_K^m$ (for an appropriate notion of size) satisfying the equation $\sum_{j=1}^m a_j z_j = 0 \bmod q$.

To obtain our improved connection factors, we will use number fields K having small *root discriminant* $\mathcal{D}_K$ (as a function of the degree n). The best number fields for our purposes will have $\mathcal{D}_K$ bounded by a *constant*, which is optimal. Using such number fields will yield a factor of $\sim \sqrt{n}$ improvement in each of the two main aspects of the worst-case to average-case reduction:

1. *Smaller average-case solutions.* We show that our average-case problem admits *very short* solutions $\mathbf{z} \in \mathcal{O}_K^m$, having an "effective size" of $\sim \sqrt{\log n}$. In all prior work, the solution size was at least $\sqrt{n \log n}$.

   The improvement stems from properties of the algebraic integers $\mathcal{O}_K$. Like $\mathbb{Z}$, the ring $\mathcal{O}_K$ has a discrete structure and a geometric notion of "absolute value" $|\cdot|$. In particular, for an element $x \in \mathcal{O}_K$ with $|x| \leq \beta$, multiplying an element by x increases its "size" by at most β.

   The crucial difference between $\mathbb{Z}$ and $\mathcal{O}_K$ is this: while $\mathbb{Z}$ contains only $\sim 2\beta$ elements of absolute value at most β, the ring $\mathcal{O}_K$ contains $\sim \beta^n$ such elements. Intuitively, and in a precise geometric sense, $\mathcal{O}_K$ is much more "densely-packed" than $\mathbb{Z}$.

   Now consider any $\mathbf{a} \in \mathcal{O}_K^m$ specifying an instance of the average-case problem. Because $\mathcal{O}_K \bmod q$ has exactly q^n residue classes, then intuitively we expect (at least) one out of every q^n vectors $\mathbf{z} \in \mathcal{O}_K^m$ to satisfy the equation $\sum a_j z_j = 0 \bmod q$. The density of $\mathcal{O}_K$ implies that there are $\sim \beta^{mn}$ (non-zero) vectors $\mathbf{z} \in \mathcal{O}_K^m$ with $|z_j| \leq \beta$ for every j. By choosing $\beta = O(1)$ and $m = O(\log n)$ appropriately so that the number of such $\mathbf{z}$ exceeds q^n, we can show that at least one such $\mathbf{z}$ satisfies the equation.

   In the reduction, the net effect of using such a $\mathbf{z}$ is the following: because $|z_j| \leq \beta = O(1)$, each offset vector $\mathbf{d}_j$ expands by only a constant factor when multiplied by z_j. Because the reduction outputs the sum of $m = O(\log n)$ such (expanded) offset vectors, the output is only a $O(\log n)$ factor longer than the individual offsets. In fact, a more sophisticated analysis actually reveals that the overall expansion is only $O(\sqrt{\log n})$.

2. *Smaller smoothing parameter.* In addition, we show that "smoothing" ideal lattices requires much less noise than general lattices. For ideal lattices over number fields with constant $\mathcal{D}_K$, the smoothing parameter is $\sim \lambda_1/\sqrt{n}$, whereas for general lattices it can be as large as $\sim \lambda_n \geq \lambda_1$ [38]. In the reduction, then, the length of the offset vectors $\mathbf{d}_j$ is $\sim \lambda_1$. Combined with the size of the average-case solutions, this fact accounts for our $O(\sqrt{\log n})$ connection factor for SVP.

   The improved smoothing parameter stems from the fact that ideal lattices and their duals simultaneously have large minimum distances λ_1 (which in turn is due to their algebraic

properties). In fact, the relationship is near optimal: for an ideal lattice Λ and its dual Λ^* (over some K having constant $\mathcal{D}_K$), we have $\lambda_1(\Lambda) \cdot \lambda_1(\Lambda^*) = \Omega(n)$, whereas the bound $\lambda_1(\Lambda) \cdot \lambda_1(\Lambda^*) \leq n$ applies for *any* lattice Λ [10]. We find it a remarkable coincidence that the same property guaranteeing small average-case solutions (namely, small root discriminant) is also exactly what accounts for this improvement in the smoothing parameter.

In addition to the improved connection factors, we also obtain a unified reduction where the worst-case problem can be stated in terms of *any* ℓ_p length, $p \in [1, \infty]$. The connection factor is (essentially) the same for all p. This result relies upon an analysis of Gaussian distributions over lattices from a concurrent work by Peikert [41].

Our treatment of general ℓ_p lengths is partly motivated by a recent result of Regev and Rosen [45], who showed that worst-case lattice problems are *easiest* in the ℓ_2 length (at least for general lattices). In light of this fact, obtaining reductions for arbitrary ℓ_p lengths under a unified connection factor is much more desirable. Furthermore, in algebraic number theory there are multiple notions of magnitude which correspond to different ℓ_p lengths, e.g. ℓ_∞ for the "height" of a number, and often ℓ_1 length for its "size."

2.3 Relationships Between Worst-Case Problems

We also provide some new reductions between worst-case problems on ideal lattices, showing, for example, that the shortest vector problem is no harder than the closest vector problem. For all of our results, analogous reductions are known for *general* lattices [24], but those reductions do not work for *ideal* lattice problems because they perform transformations that can destroy the ideal property. Nevertheless, our reductions use techniques similar to those for general lattice problems.

The essential technique from [24] for reducing, say, the shortest vector problem to the closest vector problem involves constructing a special sequence of sublattices of the input lattice. In this work, we will generate subideals of the input ideal $\mathcal{I}$ by *multiplying* $\mathcal{I}$ by a certain collection of appropriately-chosen (fixed) ideals. The crucial property of this collection is that all its ideals are individually "small," and that their product is the ideal $\langle q \rangle$ generated by some integer $q \in \mathbb{Z}$. Such a collection can be derived from any prime $q \in \mathbb{Z}$ that "splits well" over the number field K.

3 Comparison to Related Work

The idea of exploiting special families of lattices is not new. Some of the results in Ajtai's original paper [4] are based on lattices that a have "unique" shortest vector (in some formal sense), as are the cryptosystems of Ajtai and Dwork [6] and Regev [43].

Micciancio generalized Ajtai's framework by replacing the integers $\mathbb{Z}$ with an alternate ring R for the purpose of cryptographic efficiency [36]. He proposed the ring of n-dimensional integer vectors $\mathbb{Z}^n$, with *cyclic convolution* as the product operation. This yielded an efficient and "compact" *one-way* function assuming the worst-case hardness of problems on *cyclic* lattices.

In later independent works, Peikert and Rosen [42] and Lyubashevsky and Micciancio [33] observed that cyclic lattices actually correspond to ideals in the polynomial ring $R = \mathbb{Z}[x]/\langle x^n - 1 \rangle$, and obtained efficient *collision-resistant* hash functions by exploiting the algebraic structure of this ring. The latter work also suggested generalizing to other rings of the form $R = \mathbb{Z}[x]/\langle f(x) \rangle$ for other degree-n polynomials $f(x)$ satisfying certain special properties.

Our work is closely related to [36, 42, 33], but differs in a couple of crucial ways. First, we use a different kind of ring, namely the algebraic integers $\mathcal{O}_K$ in a number field K. In general, $\mathcal{O}_K$ is *not* isomorphic to any ring of the form $\mathbb{Z}[x]/\langle f(x)\rangle$. However, $\mathcal{O}_K$ always *contains* such a subring, and in certain special cases they can coincide. For example, let n be a power of 2. Then for the *cyclotomic* number field $\mathbb{Q}(\zeta)$ where $\zeta = \exp(\pi i/n)$ is a root of the irreducible polynomial $f(x) = x^n + 1$, the ring $\mathcal{O}_K$ is indeed isomorphic to $\mathbb{Z}[x]/\langle f(x)\rangle$. We caution that cyclotomic number fields are a very special case in this respect, and that they unfortunately do not have the other properties we need to obtain improved connection factors.

Just as importantly, we use a different correspondence between elements of our ring $\mathcal{O}_K$ and n-dimensional vectors. In [36, 42, 33], a polynomial residue $g(x) \in R = \mathbb{Z}[x]/\langle f(x)\rangle$ corresponds to a point in $\mathbb{Z}^n$ simply by reading g's coefficients as an n-dimensional vector. In this work, we use the *canonical embedding* of K into an n-dimensional geometric space. Even when $\mathcal{O}_K$ is isomorphic to a ring of the form $R = \mathbb{Z}[x]/\langle f(x)\rangle$, the embedding of $g \in R \equiv \mathcal{O}_K$ is in general quite different from the vector of g's coefficients. Our use of the canonical embedding will prove crucial in characterizing the geometric structure of ideal lattices.

Implicit in [38] is a bound on the smoothing parameter (similar to ours) for the special class of lattices having near-optimal (dual) minimum distances. This yields a connection factor of $\tilde{O}(\sqrt{n})$ for such lattices (because the solutions to the average-case problem are still of size $\sim\sqrt{n}$). Prior to our work, however, it was not clear whether there were any candidate families of such lattices that were amenable to a worst-case hardness assumption. (The sequence of lattices constructed by Conway and Thompson [39], for example, are optimal with respect to minimum distance, but it is not clear whether the sequence even contains more than one lattice per dimension n.)

4 Preliminaries

4.1 Notation

The complex numbers are denoted by $\mathbb{C}$, the reals by $\mathbb{R}$, the rationals by $\mathbb{Q}$, and the integers by $\mathbb{Z}$. For a real r, $\lfloor r \rceil = \lfloor r + \frac{1}{2} \rfloor$ denotes a closest integer to r. For $z \in \mathbb{C}$, $\overline{z}$ denotes the complex conjugate of z. For a positive integer n, $[n]$ denotes $\{1, \ldots, n\}$. The function log denotes the natural logarithm. For simplicity, we adopt the convention $x^{1/\infty} = 1$ for any positive $x \in \mathbb{R}$. For a function f and countable set A, $f(A)$ denotes $\sum_{x\in A} f(x)$.

Vectors are represented as bold lower-case letters, e.g. $\mathbf{x}$. For a vector $\mathbf{x}$, the ith component of x is denoted by x_i. The Hermitian inner product between $\mathbf{x}, \mathbf{y} \in \mathbb{C}^d$ is $\langle \mathbf{x}, \mathbf{y}\rangle = \sum_{i\in[n]} x_i\overline{y_i}$. Note that when $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, the Hermitian inner product and the standard inner product coincide.

For $p \in [1, \infty)$, the ℓ_p length of a vector $\mathbf{x} \in \mathbb{C}^d$ is $\|\mathbf{x}\|_p = (\sum_{i\in[n]} |x_i|^p)^{1/p}$.[2] For $p = \infty$, the ℓ_∞ length is $\|\mathbf{x}\|_\infty = \max_{i\in[n]} |x_i|$. Let D be any domain supporting an ℓ_p length. For any set $V = \{v_i\} \subseteq D$ define $\|V\|_p = \sup_i \|v_i\|_p$, and for $t \in D$, define $\mathrm{dist}^p(t, V) = \inf_i \|t - v_i\|_p$. We always take $p = 2$ whenever it is omitted. Note that $\|\mathbf{x}\|_2^2 = \langle \mathbf{x}, \mathbf{x}\rangle$ for any $\mathbf{x} \in \mathbb{C}^d$.

We write $\mathrm{poly}(\cdot)$ for some unspecified polynomial function in its parameter. We say that a function $f(n)$ is *negligible* in n if it decreases faster than the inverse of any polynomial in n, and write $\nu(n)$ for some unspecified negligible function in n.

The statistical distance between two probability distributions A and B is denoted $\Delta(A, B)$. The uniform distribution over a set S is denoted $\mathsf{U}(S)$.

[2] Usually this is called the ℓ_p *norm*, but the term "norm" will claimed by another notion from number theory.

4.2 Lattices

Let $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ be a set of n linearly independent vectors in $\mathbb{C}^d \equiv \mathbb{R}^{2d}$. The *lattice* generated by $\mathbf{B}$ is the set of all integer combinations

$$\mathcal{L}(\mathbf{B}) = \left\{ \sum_{i \in [n]} c_i \mathbf{b}_i \ : \ \forall i \in [n],\ c_i \in \mathbb{Z} \right\}.$$

The set $\mathbf{B}$ is called a *basis* of the lattice, and its size $n = |\mathbf{B}|$ the *rank*. Every lattice has an infinite number of bases generating it. For any basis $\mathbf{B}$, its *fundamental region* $\mathcal{P}(\mathbf{B}) = \{\sum_{i \in [n]} c_i \mathbf{b}_i \ : \ c_i \in [0, 1)\}$. The n-dimensional volume $\mathrm{vol}(\mathcal{P}(\mathbf{B}))$, denoted $\det \mathcal{L}(\mathbf{B})$, is called the *fundamental volume* and is invariant over any basis $\mathbf{B}$ of the lattice.

The *minimum distance* in ℓ_p length of a lattice Λ, denoted $\lambda_1^p(\Lambda)$, is the length of its shortest nonzero element (in ℓ_p length): $\lambda_1^p(\Lambda) = \min_{0 \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\|_p$. Minkowski's Theorem relates the minimum distance to the fundamental volume of a lattice:

Proposition 4.1 (Minkowski's Theorem)**.** *Let Λ be any lattice of rank n and $\mathcal{B} \subseteq \mathrm{span}(\Lambda)$ be any convex body symmetric about the origin having n-dimensional volume $\mathrm{vol}(\mathcal{B}) > 2^n \cdot \det \Lambda$. Then $\mathcal{B}$ contains some nonzero $\mathbf{x} \in \Lambda$.*

Generalizing the minimum distance, the *ith successive minimum* in ℓ_p length $\lambda_i^p(\Lambda)$ is the smallest radius r such that the ball $r\mathcal{B}_n^p$ contains i linearly independent lattice points, where $\mathcal{B}_n^p$ is the closed unit ball under the ℓ_p length. A set of n linearly independent lattice points is *not necessarily* a basis for the lattice. Let $g^p(\Lambda)$, which we call the *basis minimum* (in ℓ_p length), be the minimum r such that the ball $r\mathcal{B}_n^p$ contains a set of lattice vectors that are a *basis* of Λ. The covering radius (in ℓ_p length) $\mu^p(\Lambda)$ is the shortest radius r such ℓ_p balls $r\mathcal{B}_n^p$ centered at the points of Λ cover the subspace it spans: $\mu^p(\Lambda) = \max_{\mathbf{x} \in \mathrm{span}(\Lambda)} \mathrm{dist}^p(\mathbf{x}, \Lambda)$.

The *dual lattice* of Λ, denoted Λ^*, is defined to be $\Lambda^* = \{\mathbf{x} \in \mathrm{span}(\Lambda) \ : \ \forall\, \mathbf{v} \in \Lambda, \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbb{Z}\}$. The *transference theorems* of Banaszczyk give relations between lattices and their duals, in both the standard ℓ_2 length [10] and in general ℓ_p lengths [11]. Following Cai [15] in a straightforward manner, we can generalize Banaszczyk's results to relate the basis minimum of Λ (under any ℓ_p length) to the minimum distance of Λ^* (under the *dual length* of ℓ_p):

Lemma 4.2 (Synthesis of [11] and [15])**.** *There is a constant C such that for any lattice Λ of rank n and any $p, r \in [1, \infty]$ with $1/p + 1/r = 1$,*

$$g^p(\Lambda) \cdot \lambda_1^r(\Lambda^*) \leq C \cdot n\sqrt{\log n}.$$

4.3 Gaussian Measures

Our review of Gaussian measures over lattices follows the development of prior works [3, 43, 38]. For any $s > 0$ define the Gaussian function over $\mathbb{C}^d$ centered at $\mathbf{c} \in \mathbb{C}^d$ with parameter s as:

$$\forall \mathbf{x} \in \mathbb{C}^d,\ \rho_{s,\mathbf{c}}(\mathbf{x}) = e^{-\pi \|\mathbf{x} - \mathbf{c}\|^2 / s^2}.$$

The subscripts s and $\mathbf{c}$ are taken to be 1 and $\mathbf{0}$ (respectively) when omitted. The total measure of $\rho_{s,\mathbf{c}}(\mathbf{x})$ over any subspace $H \subseteq \mathbb{C}^d$ of dimension n is s^n, therefore we can define a continuous Gaussian probability distribution over H as $D_{s,\mathbf{c}}^H(\mathbf{x}) = s^{-n} \cdot \rho_{s,\mathbf{c}}(\mathbf{x})$. We often take the subspace H to be implicit, omitting the superscript.

$D^H_{s,\mathbf{c}}$ is the sum of n one-dimensional Gaussian distributions (aligned in orthogonal directions within H), which can each be approximated and sampled arbitrarily well using standard algorithms. For simplicity, we will assume that algorithms can efficiently sample from $D^H_{s,\mathbf{c}}$ exactly; their analysis can be made rigorous by using a sufficiently precise approximation.

For $\mathbf{c} \in \mathbb{C}^d$, $s > 0$, and lattice Λ, define the *discrete Gaussian distribution over* Λ:

$$\forall \mathbf{x} \in \Lambda,\ D_{\Lambda,s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\Lambda)}.$$

(As above, we may omit the parameters s or $\mathbf{c}$.) Intuitively, $D_{\Lambda,s,\mathbf{c}}$ can be viewed as a "conditional" distribution, resulting from sampling an $\mathbf{x}$ from $D^{\mathrm{span}(\Lambda)}_{s,\mathbf{c}}$ and conditioning on $\mathbf{x} \in \Lambda$.

The smoothing parameter. Micciancio and Regev [38] proposed a new lattice quantity which they called the *smoothing parameter*:

Definition 4.3 ([38])**.** For a lattice Λ and positive real $\epsilon > 0$, the *smoothing parameter* $\eta_\epsilon(\Lambda)$ is defined to be the smallest s such that $\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$.

The name "smoothing parameter" is motivated by the following (informal) fact: if a lattice Λ is "blurred" by adding Gaussian noise with parameter $s \geq \eta_\epsilon(\Lambda)$, the resulting distribution is within ϵ of uniform. The following lemma makes this formal:

Lemma 4.4 ([38])**.** *Let* $\mathbf{B}$ *be a lattice basis. For any* $\epsilon > 0$, $\mathbf{c} \in \mathbb{C}^d$, *and* $s \geq \eta_\epsilon(\mathcal{L}(\mathbf{B}))$,

$$\Delta(D^{\mathrm{span}(\mathcal{L}(\mathbf{B}))}_{s,\mathbf{c}} \bmod \mathcal{P}(\mathbf{B}), \mathsf{U}(\mathcal{P}(\mathbf{B}))) \leq \epsilon/2.$$

We will need the following simple bound on the smoothing parameter:

Lemma 4.5 ([38])**.** *For any lattice* Λ *of rank* n *and* $\epsilon = 2^{-n}$, $\eta_\epsilon(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^*)$.

The behavior of the discrete Gaussian distribution $D_{\Lambda,s,\mathbf{c}}$ also depends on the smoothing parameter. We will need a bound, shown by Peikert and Rosen [42], on *maximum* value of $D_{\Lambda,s,\mathbf{c}}$ (i.e., the probability of the mode):

Lemma 4.6 ([42])**.** *Let* Λ *be a lattice of rank* n. *For any* $\epsilon > 0$, $s \geq 2 \cdot \eta_\epsilon(\Lambda)$, $\mathbf{c} \in \mathbb{C}^d$, *and* $\mathbf{x} \in \Lambda$,

$$D_{\Lambda,s,\mathbf{c}}(\mathbf{x}) \leq 2^{-n} \cdot \tfrac{1+\epsilon}{1-\epsilon}.$$

5 Basic Algebraic Number Theory

In this section we review the necessary background in algebraic number theory. Due to lack of space, we will present most facts without proof (which can be found in any number of introductory books on the topic, e.g. [9, 40].) As new concepts are introduced, the reader may wish to follow along with an extended example which appears at the end of this section.

An *algebraic number* $\zeta \in \mathbb{C}$ is any root of some polynomial in $\mathbb{Q}[x]$. The *minimal polynomial* of ζ is the unique monic, irreducible polynomial $f(x) \in \mathbb{Q}[x]$ of minimal degree having ζ as a root. An *algebraic integer* is an algebraic number whose minimal polynomial has integer coefficients.

A *number field* is a field extension $K = \mathbb{Q}(\zeta)$ that is constructed by adjoining an algebraic integer ζ to the rationals $\mathbb{Q}$. The minimal polynomial $f(x)$ of ζ is called the *generating polynomial*

of K, and the *degree* of K is the degree of f. Because $f(\zeta) = 0$, there is a natural isomorphism between $\mathbb{Q}[x]/\langle f(x)\rangle$ and K, given by $x \mapsto \zeta$. Due to this isomorphism, it actually does not matter which root ζ of f we adjoin to $\mathbb{Q}$, as they all yield isomorphic number fields. Therefore, it is often more convenient to view K as the field of polynomials (having rational coefficients) modulo f, rather than as a subfield of $\mathbb{C}$.

5.1 Embeddings and Geometry

Here we describe how a number field corresponds naturally to an n-dimensional geometric space.

An *embedding* is a ring homomorphism (i.e., a function that preserves multiplication and addition, and their respective identity elements). A number field $K = \mathbb{Q}(\zeta)$ of degree n has exactly n embeddings $\{\sigma_j\}_{j\in[n]}$ into $\mathbb{C}$. Concretely, these embeddings are given by mapping ζ to each root of the generating polynomial $f(x)$. An embedding whose image lies in $\mathbb{R}$ (corresponding to a real root of f) is called a real embedding; otherwise (for a complex root of f) it is called a complex embedding. Just as f has pairs of complex conjugate roots, the complex embeddings of K are also paired into complex conjugates $\tau, \overline{\tau}$ where $\overline{\tau}(x) = \overline{\tau(x)}$ for all $x \in K$. The number of real embeddings is denoted r_1, and the number of pairs of conjugate complex embeddings is denoted r_2, so we have $n = r_1 + 2r_2$. The pair (r_1, r_2) is called the *signature* of K. By convention, we let $\{\sigma_j\}_{j\in[r_1]}$ be the real embeddings, and we number the remaining complex embeddings so that $\sigma_{j+r_1+r_2} = \overline{\sigma_{j+r_1}}$ for all $j \in [r_2]$.

The *canonical embedding* $\sigma : K \to \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ is defined by

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

One can see that σ is an embedding (i.e., a ring homomorphism) from K to $\mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$, where multiplication and addition in $\mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ are component-wise. Due to the r_2 pairs of conjugate embeddings, $\sigma(K)$ spans the n-dimensional subspace

$$H = \left\{(x_1, \ldots, x_n) \in \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2} \ : \ x_{j+r_1+r_2} = \overline{x_{j+r_1}},\ \forall\, j \in [r_2]\right\} \subseteq \mathbb{C}^n.$$

With the canonical embedding in hand, we can define geometric norms ("lengths") on K. For any $x \in K$ and any $p \in [1, \infty]$, define the ℓ_p length of x to be $\|x\|_p = \|\sigma(x)\|_p = (\sum_{i\in[n]} |\sigma_i(x)|^p)^{1/p}$ for $p < \infty$, and $\max_{i\in[n]} |\sigma_i(x)|$ for $p = \infty$. As always, we assume the ℓ_2 length when p is omitted. From these definitions and because σ is a ring homomorphism, one can see that for any $x, y \in K$ and any $p \in [1, \infty]$:

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

Therefore the ℓ_∞ length acts as an "absolute value" for elements in K (as alluded to in the discussion from Section 2.2).

5.2 The Ring of Integers and Its Ideals

Here we describe how K contains a discrete n-dimensional analog of the integers $\mathbb{Z}$, called the *ring of integers* $\mathcal{O}_K$.

Let $\mathcal{O}_K \subset K$ be the set of all algebraic integers contained in K. This set forms a ring under standard addition and multiplication of complex numbers. Additionally, $\mathcal{O}_K$ is a free $\mathbb{Z}$-module of rank n, i.e. it is the set of all $\mathbb{Z}$-combinations of some basis $B = \{b_1, \ldots, b_n\} \subset \mathcal{O}_K$. Such a basis is called an *integral basis* for K.

An *ideal* $\mathcal{I} \subseteq \mathcal{O}_K$ in the ring of integers is a nontrivial (i.e., $\mathcal{I} \neq \{0\}$) set which is a group under addition and which is closed under multiplication by $\mathcal{O}_K$, i.e. $xr \in \mathcal{I}$ for all $x \in \mathcal{I}$ and all $r \in \mathcal{O}_K$.[3] An ideal in $\mathcal{O}_K$ is the set of all $\mathcal{O}_K$-combinations of some number of elements $g_1, g_2, \ldots \in \mathcal{O}_K$, and is denoted $\langle g_1, g_2, \ldots \rangle$. Similarly to $\mathcal{O}_K$, an ideal is also a free $\mathbb{Z}$-module of rank n with some $\mathbb{Z}$-basis $\{u_1, \ldots, u_n\}$. The product of two ideals $\mathcal{I}$ and $\mathcal{J}$ is another ideal that is the ideal generated by all products xy, where $x \in \mathcal{I}$ and $y \in \mathcal{J}$.

An ideal $\mathfrak{q} \subsetneq \mathcal{O}_K$ is *prime* if whenever $a, b \in \mathcal{O}_K$ and $ab \in \mathfrak{q}$, then $a \in \mathfrak{q}$ or $b \in \mathfrak{q}$ (or both). The ring $\mathcal{O}_K$ has *unique factorization of ideals*, that is, every ideal $\mathcal{I} \subseteq \mathcal{O}_K$ can be uniquely expressed as a product of *prime* ideals. For any prime $q \in \mathbb{Z}$, the ideal $\langle q \rangle$ factors into prime ideals as $\langle q \rangle = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_L^{e_L}$ where the $\mathfrak{q}_i$ are distinct prime ideals and $1 \leq e_i \leq n$. The prime q is said to *split completely* if $L = n$ and every $e_i = 1$, and is said to be *fully ramified* if $L = 1$ and $e_1 = n$.

A *fractional ideal* $\mathcal{I}$ is a generalization of an ideal, all of whose elements can be written as fractions with some fixed denominator. Formally, there is some $d \in \mathcal{O}_K$ such that $d\mathcal{I} = \{dx : x \in \mathcal{I}\}$ is an ideal in $\mathcal{O}_K$.

5.3 Field Norm

Here we describe a quantity which measures the "size" of a field element or a (fractional) ideal, relative to the ring of integers.

The *(field) norm* of an element $x \in K$ is the product of all its embeddings $N(x) = \prod_{i \in [n]} \sigma_i(x)$; it is nonzero if $x \neq 0$, is always in $\mathbb{Q}$, and is in $\mathbb{Z}$ if $x \in \mathcal{O}_K$. Because the σ_i are ring homomorphisms, the norm is multiplicative: $N(xy) = N(x)N(y)$.

The above definition generalizes to (fractional) ideals. For any *integral* ideal $\mathcal{I} \subseteq \mathcal{O}_K$, the norm is defined to be $N(\mathcal{I}) = |\mathcal{O}_K / \mathcal{I}|$ (i.e., the number of distinct residues in $\mathcal{O}_K \bmod \mathcal{I}$). The norm of the element $x \in \mathcal{O}_K$ and the norm of its principal ideal coincide (up to sign): $N(\langle x \rangle) = |N(x)|$. For any $x \in \mathcal{I}$, $N(x)$ is divisible by $N(\mathcal{I})$ (because $\langle x \rangle$ is a subideal of $\mathcal{I}$).

For a *fractional* ideal $\mathcal{I}$ over K with denominator $d \in \mathcal{O}_K$ (i.e., $d\mathcal{I}$ is an integral ideal), the norm is defined to be $N(\mathcal{I}) = N(d\mathcal{I})/N(d)$, so it is multiplicative for (fractional) ideals as well.

5.4 Ideal Lattices

Here we describe how the ring of integers and its ideals yield lattices under the canonical embedding.

Recall that $\mathcal{O}_K$ is a $\mathbb{Z}$-module having some basis $\{b_1, \ldots, b_n\}$. Therefore, it embeds as an n-dimensional lattice $\sigma(\mathcal{O}_K)$ spanning $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ and having basis $\{\sigma(b_1), \ldots, \sigma(b_n)\}$. The *(absolute) discriminant* of K, denoted Δ_K, is the squared fundamental volume $(\det \sigma(\mathcal{O}_K))^2$ of this lattice.[4] The *root discriminant* of K, denoted $\mathcal{D}_K$, is defined to be $\Delta_K^{1/n}$. Intuitively, this is a measure of the "density" of the the algebraic integers (where smaller $\mathcal{D}_K$ means more dense).

The same ideas apply to any (fractional) ideal $\mathcal{I}$, which has some basis $\{u_1, \ldots, u_n\}$. Then $\sigma(\mathcal{I})$ is a lattice spanning the subspace H and having basis $\{\sigma(u_1), \ldots, \sigma(u_n)\}$. We call such a lattice an *ideal lattice* (over K). The fundamental volume of an ideal lattice $\sigma(\mathcal{I})$ is $N(\mathcal{I})\sqrt{\Delta_K}$.

The dual of an ideal lattice $\sigma(\mathcal{I})$ is another ideal lattice corresponding to a (possibly fractional) ideal $\mathcal{I}^*$ over an isomorphic number field $\overline{K} \equiv K$. The precise form of $\mathcal{I}^*$ is somewhat complicated and involves an ideal called the *different* [13]; we will only need the fact that $N(\mathcal{I}^*) = (N(\mathcal{I})\Delta_K)^{-1}$.

[3] Nontriviality is a somewhat non-standard condition; we adopt it for ease of exposition.

[4] Some texts define the discriminant as a signed quantity; because we will only be concerned with its magnitude, we adopt the simpler absolute definition.

For ease of notation, when referring to an ideal lattice we will often omit the embedding σ. For example, we will write $\lambda_1(\mathcal{I})$ instead of $\lambda_1(\sigma(\mathcal{I}))$, $\det \mathcal{I}$ instead of $\det \sigma(\mathcal{I})$, etc.

5.5 Distributions over Number Fields

By treating K as a Euclidean space (i.e., one with an ℓ_2 length), we can also define Gaussian distributions over K, and discrete Gaussians on ideals over K. For $c \in K$, real $s > 0$, and $x \in K$, define the Gaussian function $\rho_{s,c}(x) = \exp(-\pi \|x - c\|^2 / s^2)$. We define the continuous Gaussian probability distribution over K as $D^K_{s,c}(x) = s^{-n} \cdot \rho_{s,c}(x)$.[5] The discrete Gaussian distribution on a fractional ideal $\mathcal{I}$ is $D_{\mathcal{I},s,c}(x) = \rho_{s,c}(x)/\rho_{s,c}(\mathcal{I})$ for every $x \in \mathcal{I}$ (and zero elsewhere on K). We may efficiently sample from (a close approximation of) $D^K_{s,c}$ by sampling from $D^H_{s,\mathbf{c}}$ where $\mathbf{c} = \sigma(c)$, and applying σ^{-1} to the result.

5.6 Computational Issues

We next describe how to represent a number field K with its ring of integers $\mathcal{O}_K$, and how to perform basic operations in polynomial time. Our exposition mainly follows [28]; a detailed treatment of these issues can be found in [18, Sections 4.2–4.7] and [27, Section 2].

For measuring the computational complexity of working with a number field K, "polynomial" is taken to mean some polynomial in both n and $\log \Delta_K$. Elements in K will be represented relative to some *integral* basis $B = \{b_1, \ldots, b_n\}$: an element $x \in K$ is represented by a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Q}^n$, where $x = \sum_{i \in [n]} x_i b_i$ is the unique representation of x relative to B. Membership in $\mathcal{O}_K$ can be tested simply by checking if $\mathbf{x} \in \mathbb{Z}^n$, and addition is implemented simply by adding representations component-wise. For multiplication, it suffices by linearity to have the representation (in $\mathbb{Z}^n$) of each product $b_i b_j \in \mathcal{O}_K$ for $i, j \in [n]$, which are of polynomial size. Given these products $b_i b_j$, we can compute all of the following in polynomial time: multiplicative inverses, the discriminant Δ_K, and the embeddings from K into $\mathbb{C}$ and their inverses.

An *integral* ideal $\mathcal{I} \subseteq \mathcal{O}_K$ is represented by a $\mathbb{Z}$-basis $U = \{u_1, \ldots, u_n\} \subset \mathcal{O}_K$, where the u_i are given relative to B. A *fractional* ideal $\mathcal{I}$ is additionally represented by a denominator $d \in \mathcal{O}_K$ for which $d\mathcal{I} \subseteq \mathcal{O}_K$. Given a set U, it is possible in polynomial time to confirm that U is a $\mathbb{Z}$-basis for an ideal in $\mathcal{O}_K$ and to compute its norm. The basis U can also be kept efficiently in *Hermite Normal Form* (HNF), which makes the representation of the ideal unique. It is possible to multiply ideals and to reduce an element modulo an ideal in polynomial time. Given two ideals $\mathcal{I}' \subseteq \mathcal{I}$, it is possible to sample uniformly from the quotient group $\mathcal{I}/\mathcal{I}'$ in polynomial time and enumerate $\mathcal{I}/\mathcal{I}'$ in time polynomial in $|\mathcal{I}/\mathcal{I}'|$, $\log \Delta_K$, and n.

5.7 Example

Let $\zeta = \sqrt{13}$, an algebraic integer whose minimal polynomial is $f(x) = x^2 - 13$. The number field $K = \mathbb{Q}(\zeta)$ is of degree 2, and consists of all numbers of the form $a + b\zeta$ for $a, b \in \mathbb{Q}$.

The embeddings from K into $\mathbb{C}$ are both real, and correspond to the roots ζ and $-\zeta$ of f, so the canonical embedding σ is given by $\sigma(a + b\zeta) = (a + b\zeta, a - b\zeta)$. The subspace H spanned by $\sigma(K)$ is simply $\mathbb{R}^2$. The ℓ_2 length (for example) on K is $\|a + b\zeta\|_2 = \sqrt{2a^2 + 26b^2}$.

[5] Strictly speaking, K is not a continuous space, and therefore cannot support a continuous probability distribution. This problem can be circumvented by formally defining the Gaussian distribution $D^K_{s,c}$ over the continuous vector space $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R} \equiv H$, where $\otimes$ denotes the field tensor product.

An integral basis $B = \{b_1, b_2\}$ for K is given by $b_1 = 1$, $b_2 = \frac{1+\zeta}{2}$ (we will not prove this). Note that $\{1, \zeta\}$ is *not* an integral basis, because it does not generate the algebraic integer $\frac{1+\zeta}{2} \in K$ (which is a root of $x^2 - x - 3$). The ideal $\langle 2 \rangle$ is itself prime and has norm $N(2) = 4$, while the ideal $\langle 3 \rangle$ factors as the product of prime ideals $\langle 3, 1+\zeta \rangle$ and $\langle 3, 1-\zeta \rangle$, each having norm 3. The prime ideal $\langle 3, 1+\zeta \rangle$ has a $\mathbb{Z}$-basis $\{3, \frac{1+\zeta}{2}\}$.

The lattices for $\mathcal{O}_K$ and the prime ideal $\langle 3, 1+\zeta \rangle$ (constructed using the embeddings of their $\mathbb{Z}$-bases) are depicted in Figure 1. The discriminant Δ_K of K is 13.

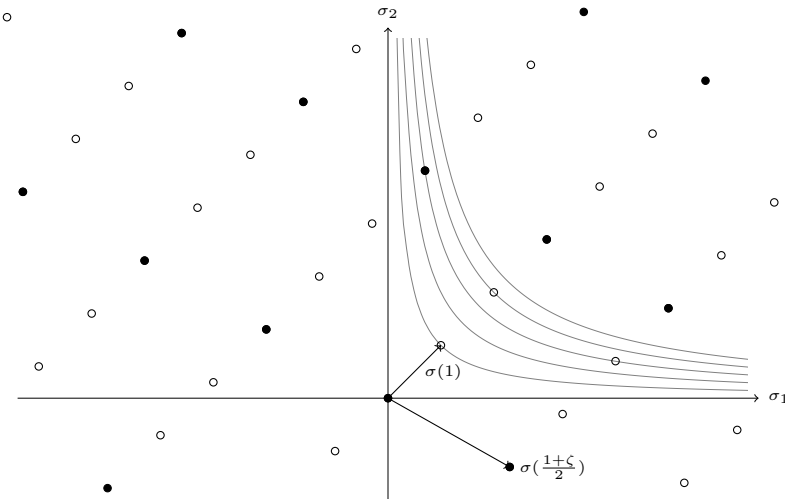

Figure 1: The lattices for the ideals $\mathcal{O}_K$ and $\mathcal{I} = \langle 3, 1+\zeta \rangle$ over the number field $K = \mathbb{Q}(\zeta)$, where $\zeta = \sqrt{13}$. All the points are elements of $\mathcal{O}_K$. The filled points are the elements of $\mathcal{I}$. All nonzero $x \in \mathcal{O}_K$ have nonzero integer field norm, hence lie on some hyperbola of the form $N(x) = \sigma_1(x) \cdot \sigma_2(x) \in \mathbb{Z} \setminus \{0\}$. All nonzero $x \in \mathcal{I}$ have norm divisible by $N(\mathcal{I}) = 3$, hence lie on some hyperbola $N(x) \in 3\mathbb{Z} \setminus \{0\}$.

6 Properties of Ideal Lattices

In this section we develop several useful facts about ideal lattices. It is likely that some of these facts are already known, however they play such an important role in our work that we prefer to present their proofs in full. Throughout this section, K denotes any number field of degree n.

6.1 Minima

Here we develop several useful facts about, and connections among, the various minima (minimum distance, basis minimum) of ideal lattices. We start with a straightforward upper bound on the minimum distance λ_1 of an ideal lattice.

Lemma 6.1. *Let K have signature (r_1, r_2), and let $\mathcal{I}$ be a fractional ideal over K. Then for any $p \in [1, \infty]$,*

$$\lambda_1^p(\mathcal{I}) \leq n^{1/p} \cdot N^{1/n}(\mathcal{I}) \cdot \sqrt{\mathcal{D}_K} \cdot \left(\tfrac{2}{\pi}\right)^{r_2/n}.$$

Proof. It suffices to prove the claim for $p = \infty$, because $\|x\|_p \le n^{1/p} \cdot \|x\|_\infty$ for any $x \in K$.

We will use Minkowski's Theorem (Proposition 4.1) to bound the minimum distance λ_1^∞. This requires some care because the ℓ_∞ length over H can involve complex coordinates.

Consider the n-dimensional closed "unit cube" $\mathcal{C} = \{\mathbf{x} \in H : \|\mathbf{x}\|_\infty \le 1\}$ in $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$. The cube $\mathcal{C}$ is the (Cartesian) product of r_1 one-dimensional cubes $\mathcal{C}_1 = \{x \in \mathbb{R} : |x| \le 1\}$, and r_2 two-dimensional "cubes" $\mathcal{C}_2 = \{(z, \overline{z}) \in \mathbb{C}^2 : |z| \le 1\}$. The first cube $\mathcal{C}_1$ has volume 2. The second cube $\mathcal{C}_2$ is a two-dimensional closed disc of radius $\sqrt{2}$, so its volume is 2π. All together, the volume of $\mathcal{C}$ is $\mathrm{vol}(\mathcal{C}) = 2^{r_1} \cdot (2\pi)^{r_2} = 2^n \cdot (\pi/2)^{r_2}$, where we have used $n = r_1 + 2r_2$.

Now for any $\beta > N^{1/n}(\mathcal{I}) \cdot \sqrt{\mathcal{D}_K} \cdot (2/\pi)^{r_2/n}$, we have

$$\mathrm{vol}(\beta\mathcal{C}) = \beta^n \,\mathrm{vol}(\mathcal{C}) > 2^n \cdot N(\mathcal{I}) \cdot \sqrt{\Delta_K} = 2^n \cdot \det \sigma(\mathcal{I}).$$

By Minkowski's Theorem (Proposition 4.1), $\beta\mathcal{C}$ contains a nonzero point of $\sigma(\mathcal{I})$, so $\lambda_1^\infty(\mathcal{I}) \le \beta$. □

Our next lemma provides a matching *lower bound* on λ_1 for ideal lattices, which differs by a factor of at most $\sqrt{\mathcal{D}_K}$ from the upper bound of Lemma 6.1. This trivially implies that the value λ_1 is easy to approximate to within a $\sqrt{\mathcal{D}_K}$ factor, because the value of the lower bound is easy to compute. This distinguishes ideal lattices in a crucial way from general lattices, in which λ_1 can be much shorter than the Minkowski bound, and is even hard to approximate [5, 34, 30, 29].

The lower bound on λ_1 is one of two *fundamental properties* of ideal lattices that yield our improved connection factors. It lies at the heart of the improved smoothing parameter (Lemma 6.5), and is also crucial for bounding the basis minimum of $\mathcal{O}_K$ (Lemma 6.3), which is needed for technical reasons in the reduction.

Lemma 6.2 (First Foundation)**.** *For any $x \in K$ and $p \in [1, \infty]$, we have $\|x\|_p \ge n^{1/p} \cdot |N(x)|^{1/n}$. Then for any fractional ideal $\mathcal{I}$ over K, we have $\lambda_1^p(\mathcal{I}) \ge n^{1/p} \cdot N^{1/n}(\mathcal{I})$.*

Proof. For $1 \le p < \infty$, by the inequality of arithmetic and geometric means we get:

$$\|x\|_p^p = \sum_{i \in [n]} |\sigma_i(x)|^p \ge n \cdot \left(\prod_{i \in [n]} |\sigma_i(x)|^p \right)^{1/n} = n \cdot |N(x)|^{p/n}.$$

Taking pth roots of both sides, we get the claimed bound.

For $p = \infty$, we see that $\|x\|_\infty = \max_{i \in [n]} |\sigma_i(x)| \ge \left(\prod_{i \in [n]} |\sigma_i(x)| \right)^{1/n} = |N(x)|^{1/n}$.

Now consider any fractional ideal $\mathcal{I}$ with denominator d, i.e. $d\mathcal{I}$ is an ideal of $\mathcal{O}_K$. For any $x \in \mathcal{I}$, $dx \in \mathcal{O}_K$ is an element of $d\mathcal{I}$. Therefore $N(dx) = N(d)N(x) \in \mathbb{Z}$ is divisible by $N(d\mathcal{I}) = |N(d)|\, N(\mathcal{I})$. For $x \ne 0$, we then have $|N(x)| \ge N(\mathcal{I})$, and the claim follows. □

Recall that the basis minimum $g^p(\mathcal{I})$ is the minimal length of a basis (in ℓ_p length) for $\mathcal{I}$.

Lemma 6.3. *There is a constant C such that for any fractional ideal $\mathcal{I}$ over K and any $p \in [1, \infty]$,*

$$g^p(\mathcal{I}) \le C \cdot n^{1/p} \cdot N^{1/n}(\mathcal{I}) \cdot \mathcal{D}_K \cdot \sqrt{\log n}.$$

In particular, $g^\infty(\mathcal{O}_K) \le C \cdot \mathcal{D}_K \cdot \sqrt{\log n}$.

Proof. Let r be such that $1/p + 1/r = 1$. By Lemma 4.2, there is some C such that

$$g^p(\mathcal{I}) \cdot \lambda_1^r(\mathcal{I}^*) \leq C \cdot n\sqrt{\log n}.$$

Because $\mathcal{I}^*$ is a fractional ideal over a number field of degree n having root discriminant $\mathcal{D}_K$, Lemma 6.2 applies, yielding

$$\lambda_1^r(\mathcal{I}^*) \geq n^{1/r} \cdot N^{1/n}(\mathcal{I}^*) = n^{1/r} \cdot N^{-1/n}(\mathcal{I}) \cdot \mathcal{D}_K^{-1}.$$

Division yields the claim. The particular case of $\mathcal{I} = \mathcal{O}_K$ follows from $N(\mathcal{O}_K) = 1$. □

The basis minimum of $\mathcal{O}_K$ actually yields a tight connection between the first and nth successive minima of any ideal lattice. This is because a single element can be multiplied by the n elements of a short integral basis, yielding n independent elements (though possibly not a basis) of similar length in the ideal. We remark that this technique is *constructive* (given a short integral basis), and indeed we will also use it in our worst-case to average-case reduction.

Lemma 6.4. *For any fractional ideal $\mathcal{I}$ over K and any $p \in [1, \infty]$,*

$$\lambda_n^p(\mathcal{I}) \leq g^\infty(\mathcal{O}_K) \cdot \lambda_1^p(\mathcal{I}).$$

Proof. Let $B = \{b_1, \ldots, b_n\}$ be an integral basis of K with $\|B\|_\infty = g^\infty(\mathcal{O}_K)$. Take $x \in \mathcal{I}$ such that $\|x\|_p = \lambda_1^p(\mathcal{I})$, and consider the set $X = \{b_1 x, \ldots, b_n x\}$. First, $X \subseteq \mathcal{I}$ because $b_i \in \mathcal{O}_K$ for all $i \in [n]$. Also, the elements in X are nonzero (because $\mathcal{O}_K$ is an integral domain) and independent (because $b_1, \ldots, b_n$ are independent), so $\lambda_n^p(\mathcal{I}) \leq \|X\|_p \leq \|B\|_\infty \cdot \|x\|_p = g^\infty(\mathcal{O}_K) \cdot \lambda_1^p(\mathcal{I})$. □

6.2 Smoothing Parameter

Here we present a bound on the smoothing parameter for ideal lattices which is especially strong for number fields of small root discriminant (see discussion below). The proof is quite straightforward given our tools from above; it mainly relies on the fact that (the duals of) ideal lattices have large minimum distance λ_1.

Lemma 6.5. *For any fractional ideal $\mathcal{I}$ over K, $\eta_\epsilon(\mathcal{I}) \leq N^{1/n}(\mathcal{I}) \cdot \mathcal{D}_K$, where $\epsilon = 2^{-n}$.*

Proof. We have

$$\eta_\epsilon(\mathcal{I}) \leq \frac{\sqrt{n}}{\lambda_1(\mathcal{I}^*)} \leq \frac{\sqrt{n}}{\sqrt{n} \cdot N^{1/n}(\mathcal{I}^*)} = \frac{1}{(N^{1/n}(\mathcal{I}) \cdot \mathcal{D}_K)^{-1}},$$

where the first inequality follows from Lemma 4.5, the second from Lemma 6.2, and the last equality by the norms of dual ideals. □

As mentioned in Section 2.2, Lemma 6.5 provides up to a $\Theta(\sqrt{n})$ factor improvement over a similar bound for general lattices [38]. Consider a number field K with *constant* root discriminant $\mathcal{D}_K$, and a fractional ideal $\mathcal{I}$ over K with $N(\mathcal{I}) = 1$ (without loss of generality). Then by Lemma 6.2 we have $\lambda_1(\mathcal{I}) \geq \sqrt{n}$, and by Lemma 6.5, $\eta_\epsilon(\mathcal{I}) \leq \mathcal{D}_K \sim \lambda_1/\sqrt{n}$. In contrast, the bound for general lattices is $\eta_\epsilon(\mathcal{I}) \sim \lambda_n$ (for a larger, but still negligible, function $\epsilon(n)$).

We also remark that the proof of Lemma 6.5 is oblivious to the particular ideal $\mathcal{I}$. The proof only depends on the discriminant of the number field (and $N(\mathcal{I})$, but that can be normalized away). We do not know if there is a stronger bound that uses more information about the lattice or has a better dependence on the discriminant, even for a negligible $\epsilon(n)$ larger than 2^{-n}.

6.3 Module Lattices

As described in Section 2.2, our average-case problem involves an equation over $\mathcal{O}_K^m$ for some positive integer m. As we will see in this subsection, the set of solutions to the equation forms a structure called a $\mathcal{O}_K$-*module*, which can be viewed as a straightforward generalization of an ideal. Here we develop some of the essential concepts about modules: we define ℓ_r lengths for $\mathcal{O}_K^m$, describe how modules embed as lattices, and give a crucial bound on the lengths of solutions to our average-case problem.

An $\mathcal{O}_K$-module in $\mathcal{O}_K^m$ is a set $\Psi \subseteq \mathcal{O}_K^m$ that is closed under (coordinate-wise) addition and scalar multiplication by any element in $\mathcal{O}_K$. Formally, for any $\mathbf{x} = (x_1, \ldots, x_m), \mathbf{y} = (y_1, \ldots, y_m) \in \mathcal{O}_K^m$ and any $c \in \mathcal{O}_K$, we have $\mathbf{x} + \mathbf{y} = (x_1 + y_1, \ldots, x_m + y_m) \in \Psi$ and $c\mathbf{x} = (cx_1, \ldots, cx_m) \in \Psi$. For $m = 1$, then, an $\mathcal{O}_K$-module is simply an ideal in $\mathcal{O}_K$.

Let σ be the canonical embedding of K into $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$. Overloading notation slightly, also define $\sigma : K^m \to H^m$ as $\sigma(\mathbf{z}) = (\sigma(z_1), \ldots, \sigma(z_m))$. For any $r \in [1, \infty)$, define the ℓ_r length

$$\|\mathbf{z}\|_r = \|\sigma(\mathbf{z})\|_r = \left(\sum\nolimits_{i\in[n], j\in[m]} |\sigma_i(z_j)|^r\right)^{1/r}.$$

The ℓ_∞ length is defined similarly, as $\|\mathbf{z}\|_\infty = \|\sigma(\mathbf{z})\|_\infty = \max_{i,j} |\sigma_i(z_j)|$.

For a positive $q \in \mathbb{Z}$ and $\mathbf{a} = (a_1, \ldots, a_m) \in \mathcal{O}_K^m$, define the set

$$\Psi_q^K(\mathbf{a}) = \left\{\mathbf{z} = (z_1, \ldots, z_m) \in \mathcal{O}_K^m \ : \ \sum\nolimits_{j\in[m]} a_j z_j \in \langle q \rangle\right\}.$$

(We omit K when it is clear from context.) One can verify that $\Psi = \Psi_q(\mathbf{a})$ is an $\mathcal{O}_K$-module. In addition, it is a free module of rank m, that is, it is generated by an $\mathcal{O}_K$-basis $\{\mathbf{b}_1, \ldots, \mathbf{b}_m\} \subset \mathcal{O}_K^m$. Therefore, the embedding $\sigma(\Psi) \subset H^m$ is a lattice of rank mn; we call it a *module lattice* (often omitting σ for clarity).

Our next lemma provides the second *fundamental fact* needed for obtaining our improved connection factors: the module lattices $\Psi_q(\mathbf{a})$ contain *very short* solution vectors (in any ℓ_r length). For simplicity, consider the ℓ_∞ length. The lemma states that $\Psi_q(\mathbf{a})$ contains a nonzero element of length at most $\sqrt{\mathcal{D}_K} \cdot q^{1/m}$. Looking ahead, this bound will be made as small as a *constant* for constant $\mathcal{D}_K$ and appropriately-chosen $q = \text{poly}(n)$, $m = O(\log n)$.

Lemma 6.6 (Second Foundation)**.** *Let K be a number field of degree n with signature (r_1, r_2), let $r \in [1, \infty]$, and let m, q be positive integers. Then for any $\mathbf{a} \in \mathcal{O}_K^m$,*

$$\lambda_1^r(\Psi_q^K(\mathbf{a})) \quad \leq \quad (mn)^{1/r} \cdot q^{1/m} \cdot \sqrt{\mathcal{D}_K} \cdot \left(\tfrac{2}{\pi}\right)^{r_2/n}.$$

Proof. It suffices to prove the claim for $r = \infty$, because $\|\mathbf{z}\|_r \leq (mn)^{1/r} \cdot \|\mathbf{z}\|_\infty$, for all $\mathbf{z} \in K^m$.

The proof is similar to that of Lemma 6.1: we analyze the fundamental volume of the mn-dimensional lattice $\sigma(\Psi_q(\mathbf{a})) \subset H^m$, then use Minkowski's Theorem to bound its minimum distance.

First, $\Psi_q(\mathbf{a})$ is an additive subgroup of $\mathcal{O}_K^m$. Now consider the quotient group $G = \mathcal{O}_K^m / \Psi_q(\mathbf{a})$: two elements $\mathbf{z}, \mathbf{z}' \in \mathcal{O}_K^m$ represent the same residue in G iff $\sum a_j z_j = \sum a_j z'_j \bmod \langle q \rangle$. Therefore the size of G is at most the number of residue classes mod $\langle q \rangle$, which is $N(\langle q \rangle) = N(q) = q^n$. The determinant of the lattice $\sigma(\mathcal{O}_K^m)$ is $(\sqrt{\Delta_K})^m$, so we conclude that $\det \Psi_q(\mathbf{a}) \leq q^n \cdot (\sqrt{\Delta_K})^m$.

Now consider the mn-dimensional "closed unit cube" $\mathcal{C}^m = \{\mathbf{x} \in H^m \ : \ \|\mathbf{x}\|_\infty \leq 1\}$, where $\mathcal{C}$ is the unit cube in H. From the proof of Lemma 6.1, the volume of $\mathcal{C}^m$ is $2^{mn} \cdot (\pi/2)^{mr_2}$.

Now for any $\beta > q^{1/m} \cdot \sqrt{\mathcal{D}_K} \cdot (2/\pi)^{r_2/n}$, we have

$$\mathrm{vol}(\beta\mathcal{C}^m) = \beta^{mn} \cdot \mathrm{vol}(\mathcal{C}^m) > 2^{mn} \cdot q^n \cdot (\sqrt{\Delta_K})^m.$$

By Minkowski's Theorem (Proposition 4.1), $r\mathcal{C}$ contains a nonzero point of $\Psi_q(\mathbf{a})$. □

6.4 Sums of Discrete Gaussians

In our worst-case to average-case reduction, the output is distributed roughly according to the sum of m discrete Gaussians over the input ideal. Each discrete Gaussian is scaled by a factor $z_j \in \mathcal{O}_K$, where $\mathbf{z} = (z_1, \ldots, z_m) \in \mathcal{O}_K^m$ is a short solution to the random instance of our average-case problem. Therefore the length of the output, and hence the quality of the reduction, is largely determined by the behavior of such sums of discrete Gaussians. In this section, we give tail bounds on the ℓ_p length as a function of the length $\|\mathbf{z}\|_r$, for various values of p and r. The bounds follow from a general analysis of discrete Gaussians in a concurrent paper by Peikert [41].

Lemma 6.7. *Let $\mathcal{I}$ be a fractional ideal over a number field K of degree n. Let m be a positive integer, $\epsilon \leq 1/(2m+1)$ be a positive real, $s \geq \eta_\epsilon(\mathcal{I})$, and $\mathbf{z}, \mathbf{c} \in K^m$. Let $y_j \sim D_{\mathcal{I},s,c_j}$ be independent samples from discrete Gaussians over $\mathcal{I}$, and define $v = \sum_{j\in[m]} z_j \cdot (y_j - c_j)$.*

If $\|\mathbf{z}\|_\infty \leq \beta$, then we have $\Pr\left[\|v\|_p > L\right] \leq \frac{1}{2}$, where

$$L = c_p \cdot s \cdot \beta \cdot \sqrt{m} \cdot \begin{cases} n^{1/p} & \text{for} \quad p \in [1, \infty) \\ \sqrt{\log n} & \text{for} \quad p = \infty \end{cases}$$

and c_p is a constant depending only on p.

In the context of module lattices and our average-case problem, a requirement of the form $\|\mathbf{z}\|_\infty \leq \beta$ (as in the lemma above) is qualitatively the *strictest* one we can impose. In light of Lemma 6.6, it is also interesting (and may be useful in cryptographic contexts) to consider relaxed requirements of the form $\|\mathbf{z}\|_r \leq \beta \cdot (mn)^{1/r}$ for other ℓ_r lengths. For such $\mathbf{z}$, we can show a nearly identical bound on the sum of discrete Gaussians in any ℓ_p length for $p \leq r$. The only effective difference from Lemma 6.7 is an $m^{1/r}$ factor rather than $\sqrt{m}$, when $r < 2$.

Lemma 6.8. *Take the same notation described in Lemma 6.7.*

If $\|\mathbf{z}\|_r \leq \beta \cdot (mn)^{1/r}$ for some $r \in [1, \infty)$, then for any $p \leq r$ we have $\Pr\left[\|v\|_p > L\right] \leq \frac{1}{2}$, where

$$L = c_p \cdot s \cdot \beta \cdot m^{\max\{1/2, 1/r\}} \cdot n^{1/p}.$$

The proofs of Lemmas 6.7 and 6.8 are somewhat technical (but relatively straightforward), and are given in Appendix A.

7 Computational Problems on Ideal Lattices

In this section we define several computational problems (both worst-case and average-case) relating to number fields and ideal lattices.

7.1 Preprocessing Number Fields

All of the problems we define are parameterized by a fixed choice of number field K (or, in their asymptotic versions, an infinite family $\mathcal{K}$ of number fields). Because the number field is fixed for all time in advance, an adversary can perform computations on it for an arbitrarily long time, and use what it has learned when finally presented with a specific instance over K to solve. This is an example of a general notion called *preprocessing*, which also applies to problems in coding, lattices, and cryptography (see, e.g. [14, 35, 21]).

All of the problems we define in this section should be interpreted as problems with preprocessing of the number field. That is, any algorithm for solving a problem over K receives a polynomially-long (in the representation of K) auxiliary input which can depend arbitrarily on K. We refer to this auxiliary input as "advice about K" in any of our reductions that use it. Alternately, one may imagine a specific circuit designed to solve a problem over a specific number field. Similar comments apply for families $\mathcal{K}$ of number fields and sequences of advice strings or circuit families.

7.2 Worst-Case Problems

Here we define several worst-case problems on ideal lattices. By scaling, it will suffice to define these problems only for integral (rather than fractional) ideals $\mathcal{I} \subseteq \mathcal{O}_K$.

In all of the computational problems below, p is any value in $[1, \infty]$, γ is a fixed positive real, and ϕ is some arbitrary function on lattices (one may imagine $\phi = \lambda_1^p$ or $\phi = \eta_\epsilon$ for concreteness). For now, all of the problems are defined over a fixed number field K.

Definition 7.1 (Ideal Generalized/Shortest Vector Problem)**.** An input to K-$\mathsf{IGVP}_\gamma^{p,\phi}$ is an ideal $\mathcal{I} \subseteq \mathcal{O}_K$. The goal is to output a nonzero $x \in \mathcal{I}$ such that $\|x\|_p \leq \gamma \cdot \phi(\mathcal{I})$. The *ideal shortest vector problem*, denoted K-ISVP_γ^p, is the special case where $\phi = \lambda_1^p$.

We next define an *incremental* version of IGVP, which will be the actual worst-case problem we reduce to our average-case problem. The purpose of introducing this incremental problem is to simplify the worst-case to average-case reduction down to its most essential ideas.

Definition 7.2 (Incremental IGVP)**.** An input to K-$\mathsf{IncIGVP}_\gamma^{p,\phi}$ is a pair $(\mathcal{I}, x)$ where $\mathcal{I}$ is an ideal in $\mathcal{O}_K$ and $x \in \mathcal{I}$ such that $\|x\|_p > \gamma \cdot \phi(\mathcal{I})$. The goal is to output a nonzero $x' \in \mathcal{I}$ such that $\|x'\|_p \leq \|x\|_p / 2$.

It is straightforward to show that there is a standard reduction from K-$\mathsf{IGVP}_\gamma^{p,\phi}$ to K-$\mathsf{IncIGVP}_\gamma^{p,\phi}$ which makes a polynomial number of calls to its oracle.

We can also define a generalized promise problem on ideal lattices, which captures the problem of estimating any particular lattice parameter (such as λ_1 or the covering radius μ).

Definition 7.3 (Generalized Parameter Problem)**.** An input to K-$\mathsf{GapIGPP}_\gamma^\phi$ is a pair $(\mathcal{I}, R)$ where $\mathcal{I}$ is an ideal in $\mathcal{O}_K$ and $R \in \mathbb{R}$. It is a YES instance if $\phi(\mathcal{I}) \leq R$, and is a NO instance if $\phi(\mathcal{I}) > \gamma \cdot R$.

The problem K-$\mathsf{GapISVP}_\gamma^p$ is defined by setting $\phi = \lambda_1^p$. The problem K-$\mathsf{GapICRP}_\gamma^p$ is defined by setting $\phi = \mu^p$.

The following are the ideal lattice variants of the closest vector problem in its search and decision versions (respectively):

Definition 7.4 (Ideal Closest Vector Problem)**.** An input to K-ICVP_γ^p is a pair $(\mathcal{I}, t)$ where $\mathcal{I}$ is an ideal in $\mathcal{O}_K$ and $t \in K$. The goal is to output a $v \in \mathcal{I}$ such that $\|t - v\|_p \leq \gamma \cdot \mathrm{dist}^p(t, \mathcal{I})$.

Definition 7.5 (Gap ICVP)**.** An input to K-$\mathsf{GapICVP}^p_\gamma$ is a tuple $(\mathcal{I}, t, R)$ where $\mathcal{I}$ is an ideal in $\mathcal{O}_K$, $t \in K$, and $R \in \mathbb{R}$. It is a YES instance if $\mathrm{dist}^p(t, \mathcal{I}) \leq R$, and is a NO instance if $\mathrm{dist}^p(t, \mathcal{I}) > \gamma \cdot R$.

Asymptotics. In order to speak meaningfully about asymptotic hardness, we parameterize all of the above problems by an infinite family $\mathcal{K} = \{K_n\}_{n \in T}$ of number fields (for some infinite set $T \subseteq \mathbb{N}$), where K_n has degree n.[6] This is analogous to the formulation of computational problems for particular infinite families of error-correcting codes (e.g., Reed-Solomon codes).

For an infinite family $\mathcal{K}$ of number fields and a function $\gamma : \mathbb{Z}^+ \to \mathbb{R}^+$, for any problem K-P_γ above, we define $\mathcal{K}$-P_γ to be the ensemble of instances from K_n-$\mathsf{P}_{\gamma(n)}$. When reducing from a problem $\mathcal{K}$-P to another problem $\mathcal{K}$-P', we say that the reduction is *number field-preserving* if, for every input instance of the problem K_n-P, the reduction only issues queries on instances of the problem K_n-P'.

7.3 Average-Case Problem

We now define our average-case problem, whose goal is to find a nontrivial $\mathbf{z} = (z_1, \ldots, z_m) \in \Psi_q(\mathbf{a})$ that is short in ℓ_r length, for random $\mathbf{a}$.

Definition 7.6 (Short Algebraic Integer Solution)**.** For a number field K, positive $q, m \in \mathbb{Z}$, positive $\beta \in \mathbb{R}$, and $r \in [1, \infty]$, an input to the problem K-$\mathsf{SAIS}^r_{q,m,\beta}$ is a vector $\mathbf{a} \in \mathcal{O}_K^m$. The goal is to find a nonzero $\mathbf{z} \in \Psi_q(\mathbf{a})$ such that $\|\mathbf{z}\|_r \leq \beta \cdot (mn)^{1/r}$.

For an infinite family $\mathcal{K} = \{K_n\}$ of number fields, $r \in [1, \infty]$, and functions $q(n)$, $m(n)$, $\beta(n)$, define $\mathcal{K}$-$\mathsf{SAIS}^r_{q,m,\beta}$ to be the probability ensemble over instances $\mathbf{a} = (a_1, \ldots, a_{m(n)})$ of K_n-$\mathsf{SAIS}_{q(n),m(n),\beta(n)}$ where the a_i are chosen independently and uniformly from a canonical set of residues representing $\mathcal{O}_{K_n}/\langle q(n) \rangle$.

Implicit in the definition of K-SAIS is the assumption that a short enough solution $\mathbf{z} \in \Psi_q(\mathbf{a})$ exists (otherwise the problem is trivially unsolvable). Lemma 6.6 gives a relationship among the parameters that guarantees the existence of a solution. We require that β be no less than the bound from Lemma 6.6, otherwise we say that the problem is undefined.

8 Complexity of Worst-Case Problems

8.1 Easy Estimation Problems

In this subsection we show that it is easy to estimate various parameters of ideal lattices to within small factors, by showing that their corresponding promise problems are in P. These results all follow very easily from the bounds we derived in Section 6.1 (plus other tools on lattices such as transference theorems), therefore we will omit detailed explanations. We remark that none of these results seem to have any impact on the apparent difficulty of *search* problems on ideal lattices.

Lemma 8.1. *Let $\mathcal{K} = \{K_n\}$ be a family of number fields. The problem $\mathcal{K}$-$\mathsf{GapIGPP}^\phi_\gamma$ is easy (i.e., in P) for the following choices of the lattice parameter ϕ and approximation factor $\gamma(n)$:*

- *$\phi = \lambda_1^p$ and $\gamma(n) = \sqrt{\mathcal{D}_{K_n}}$.*

[6] In fact, for any number field K of *fixed* degree, all of the above problems are solvable *exactly* (i.e., for $\gamma = 1$) in time polynomial in the instance size, e.g. using [7]. Of course, the running time grows exponentially in the degree.

- $\phi = \lambda_n^p$ *or* $\phi = g^p$ *(the basis minimum in* ℓ_p *length) and* $\gamma(n) = \mathcal{D}_{K_n} \cdot O(\sqrt{\log n})$.
- $\phi = \mu^p$ *(the covering radius in* ℓ_p *length) and* $\gamma(n) = \sqrt{\mathcal{D}_{K_n}} \cdot O(n\sqrt{\log n})$.

8.2 Reductions Between Problems

In this subsection we provide some (worst-case to worst-case) reductions between problems on ideal lattices, such as from SIVP to SCVP. For all of our results, analogous reductions are known to exist for *general* lattices [24], but those reductions are not valid for *ideal* lattice problems because they query their oracles on non-ideal lattices. Nevertheless, the reductions we demonstrate here use techniques similar to those for general lattice problems.

The essential technique from [24] for reducing, say, SVP to CVP can be abstracted in the following way: for an instance involving a lattice Λ, construct a carefully-chosen set of sublattices $\Lambda_i \subseteq \Lambda$ such that (1) the quotient groups Λ/Λ_i are all small, and (2) the intersection of all the Λ_is cannot contain a shortest vector of Λ. For general lattices, Λ_i is constructed simply by doubling the ith basis vector of Λ, and leaving the remaining basis vectors unchanged. This makes the size of the quotient groups $|\Lambda/\Lambda_i| = 2$, and the intersection $\bigcap \Lambda_i = 2\Lambda$. While this technique satisfies the conditions from above, in our setting it may not yield *ideal* sublattices.

Instead, we will generate subideals of the input ideal $\mathcal{I}$ by *multiplying* $\mathcal{I}$ by a collection of appropriately-chosen (fixed) ideals. We also generalize the above structure, constructing several *chains* of subideals $\mathcal{I}_{i,e} \subset \cdots \subset \mathcal{I}_{i,1} \subset \mathcal{I}_{i,0} = \mathcal{I}$ such that (1) the quotient groups $\mathcal{I}_{i,j-1}/\mathcal{I}_{i,j}$ are all small, and (2) the intersection of all the $\mathcal{I}_{i,j}$s cannot contain a shortest vector in $\mathcal{I}$.

Our reductions will rely on an integer prime $q \in \mathbb{Z}$ that "splits well" over $\mathcal{O}_K$ into prime ideal divisors having small norm. That is, if $\langle q \rangle$ factors in $\mathcal{O}_K$ as $\langle q \rangle = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_L^{e_L}$, we will need $N(\mathfrak{q}_i)$ to be small for all i. One way (but not the only way) of satisfying this condition is to let q be a prime that *splits completely* in $\mathcal{O}_K$, namely, $\langle q \rangle$ factors into n distinct prime ideals, each of norm q. Guruswami [25] demonstrated that there exist infinite families of number fields, all having the same (constant) root discriminant, for which some fixed q splits completely in every member of the family. Another way is to let q be a prime that is *fully ramified* in $\mathcal{O}_K$, namely, $\langle q \rangle = \mathfrak{q}^n$ where $N(\mathfrak{q}) = q$. We do not know if there is an efficient way, given an arbitrary number field, to find a prime q that splits completely, is fully ramified, or otherwise splits well for our purposes.

No matter how q splits in $\mathcal{O}_K$, the ideals $\mathfrak{q}_1^{e_1}, \ldots, \mathfrak{q}_L^{e_L}$ are pairwise relatively prime, so for any fractional ideal $\mathcal{I}$ we have

$$\bigcap_{i \in [L]} (\mathfrak{q}_i^{e_i} \mathcal{I}) = \left(\bigcap_{i \in [L]} \mathfrak{q}_i^{e_i} \right) \cdot \mathcal{I} = q \cdot \mathcal{I},$$

which cannot contain a shortest vector of $\mathcal{I}$.

For the remainder of this section, let K be a number field for which prime $q \in \mathbb{Z}$ factors as $\langle q \rangle = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_L^{e_L}$. All of the reductions we give below are between problems on the fixed number field K (not a family $\mathcal{K}$), because their efficiency depends only on the input size and the splitting behavior of q, and not (explicitly) on the dimension n.

Reducing SIVP to SCVP. Here we show that for ideal lattices, and for any ℓ_p length, approximating the shortest vector is no harder than approximating the closest vector, with no loss in approximation ratio.

Proposition 8.2. *For any γ and any $p \in [1, \infty]$, there is a deterministic non-adaptive Cook reduction from K-ISVP^p_γ (resp., K-$\mathsf{GapISVP}^p_\gamma$) to K-ICVP^p_γ (resp., K-$\mathsf{GapICVP}^p_\gamma$). The reduction makes $\sum_{i \in [L]} e_i \cdot (N(\mathfrak{q}_i) - 1)$ oracle queries, and runs in time $\mathrm{poly}(S) \cdot \sum_{i \in [L]} e_i \cdot N(\mathfrak{q}_i)$, where S is the input size.*

Proof. We provide a reduction between the search problems, which can be easily adapted for the decision versions. The advice about K needed by the reduction is the value of q and the factorization of $\langle q \rangle$ into prime ideals.

Suppose oracle $\mathcal{A}$ solves K-ICVP^p_γ in the worst case. Then our reduction proceeds as follows: on input an ideal $\mathcal{I} \subseteq \mathcal{O}_K$,

1. For each $i \in [L]$ and each $j \in \{0, \ldots, e_i\}$, let $\mathcal{I}_{i,j} = \mathfrak{q}_i^j \mathcal{I}$.
2. For each $i \in [L]$, $j \in [e_i]$, and each *nonzero* $t_{i,j,k} \in \mathcal{I}_{i,j-1}/\mathcal{I}_{i,j}$, let $v_{i,j,k} \leftarrow \mathcal{A}(\mathcal{I}_{i,j}, t_{i,j,k})$.
3. Among all vectors $t_{i,j,k} - v_{i,j,k}$, output one whose ℓ_p length is minimal.

We first analyze the running time of the reduction. Given bases for $\mathcal{I}$ and each $\mathfrak{q}_i$ we can efficiently compute a basis for $\mathcal{I}_{i,j} = \mathfrak{q}_i^j \mathcal{I}$ by performing $j \leq n$ multiplications of ideals. We can also enumerate over $\mathcal{I}_{i,j-1}/\mathcal{I}_{i,j}$ in time $N(\mathfrak{q}_i) \cdot \mathrm{poly}(S)$. The size of $\mathcal{I}_{i,j-1}/\mathcal{I}_{i,j}$ is $N(\mathfrak{q}_i)$, so the number of calls to $\mathcal{A}$ is $\sum_{i \in [L]} e_i \cdot (N(\mathfrak{q}_i) - 1))$.

We now prove that the reduction is correct. First, we see that $0 \neq t_{i,j,k} - v_{i,j,k} \in \mathcal{I}$ for every i, j, k, because both $t_{i,j,k}, v_{i,j,k} \in \mathcal{I}_{i,j-1} \subset \mathcal{I}$, but $t_{i,j,k} \notin \mathcal{I}_{i,j}$ while $v_{i,j,k} \in \mathcal{I}_{i,j}$. Therefore the reduction outputs a nonzero element of $\mathcal{I}$.

Now let $w \in \mathcal{I}$ be such that $\|w\|_p = \lambda_1^p(\mathcal{I})$. Then $w \notin \bigcap_{i \in [L]} (\mathfrak{q}_i^{e_i} \mathcal{I}) = q\mathcal{I}$. By the Chinese Remainder Theorem, there exists an $i \in [L]$ such that $w \neq 0 \bmod \mathcal{I}_{i,e_i}$. Then there exists a $j \in [e_i]$ such that $w \neq 0 \bmod \mathcal{I}_{i,j}$ but $w = 0 \bmod \mathcal{I}_{i,j-1}$. Therefore there exists some k such that $w = t_{i,j,k} \bmod \mathcal{I}_{i,j}$. Therefore $\mathrm{dist}^p(t_{i,j,k}, \mathcal{I}_{i,j}) = \mathrm{dist}^p(w, \mathcal{I}_{i,j}) \leq \|w\|_p = \lambda_1^p(\mathcal{I})$. By assumption on $\mathcal{A}$, $\|t_{i,j,k} - v_{i,j,k}\|_p \leq \gamma \cdot \mathrm{dist}^p(t_{i,j,k}, \mathcal{I}_{i,j}) \leq \gamma \cdot \lambda_1^p(\mathcal{I})$, so the reduction solves ISVP^p_γ. □

Reducing ICVP_1 to $\mathsf{GapICVP}_1$. Here we show a reduction from search to decision for the *exact* versions of the closest vector problem on ideal lattices. Just as for general lattices, we do not know of a reduction to the *approximation* version of the decision problem (for factors $\gamma > 1$).

Proposition 8.3. *For any γ and any $p \in [1, \infty]$, there is a deterministic adaptive Cook reduction from K-ICVP^p_1 to K-$\mathsf{GapICVP}^p_1$. The reduction runs in time $\mathrm{poly}\left(S \cdot \sum_{i \in [L]} e_i \cdot N(\mathfrak{q}_i)\right)$, where S is the input size.*

Proof. The advice about K needed by the reduction is the value of q, its factorization into prime ideals, and a set of coefficients for performing Chinese remaindering mod $\langle q \rangle$, specifically: for every $i \in [L]$, an element $r_i \in \mathcal{O}_K$ such that $r_i = 1 \bmod \mathfrak{q}_i^{e_i}$ and $r_i \in \mathfrak{q}_k^{e_k}$ for every $k \neq i$.

On an instance $(\mathcal{I}, t)$, let $v \in \mathcal{I}$ be some closest lattice point to t. It will suffice for the reduction to compute $w = v \bmod q\mathcal{I}$. Then we can iterate the reduction with $\mathcal{I}' = q\mathcal{I}$ and $t' = t - w$, which will output $w' = (v - w) \bmod q^2\mathcal{I}$, etc. After a polynomial number of iterations, we can reconstruct all of $v \in \mathcal{I}$. We defer the details to the full version.

In order to compute $v \bmod q\mathcal{I}$, the reduction will progressively find, for every i and increasing values of j up to e_i, the residue $v \bmod (\mathfrak{q}_i^j \mathcal{I})$. Using the final values $v_i = v \bmod (\mathfrak{q}_i^{e_i} \mathcal{I})$, it will then reconstruct $v \bmod q\mathcal{I}$ using the Chinese remaindering coefficients.

Suppose oracle $\mathcal{A}$ solves K-$\mathsf{GapICVP}_1^p$ in the worst case. The reduction proceeds as follows: on input $(\mathcal{I}, t)$ where $\mathcal{I}$ is an ideal of $\mathcal{O}_K$ and $t \in K$,

1. Compute the distance $d = \mathrm{dist}^p(t, \mathcal{I})$ from t to the lattice via binary search. (Details omitted.)
2. For $i \in [L]$ and $j \in \{0, \ldots, e_i\}$, let $\mathcal{I}_{i,j} = \mathfrak{q}_i^j \mathcal{I}$ be as in the reduction of Proposition 8.2.
3. For each $i \in [L]$, let $v_i = 0$ and $t' = t$. Then for each $j = 1, \ldots, e_i$ do:
   (a) Find (by enumeration) some $x \in \mathcal{I}_{i,j-1}/\mathcal{I}_{i,j}$ for which $\mathcal{A}(\mathcal{I}_{i,j}, t' - x, d) = \text{YES}$.
   (b) Let $t' = t' - x$, and $v_i = v_i + x$.
4. Output $\sum_{i \in [L]} v_i \cdot r_i \bmod q\mathcal{I}$.

Using arguments similar to those in [24], we can show that in Step (3a) there is always an x that makes $\mathcal{A}$ output YES. It is also not hard to show that the final values of v_i are as described above, and that the output is $v \bmod q\mathcal{I}$ by the Chinese remainder theorem. We defer the details. □

9 Worst-Case to Average-Case Reduction

In this section we give a reduction from solving the problem IGVP (actually, its equivalent incremental version) in the worst case to solving SAIS on the average. We complete the section by connecting ISVP to IGVP for concrete choices of the family $\mathcal{K}$ of number fields.

Theorem 9.1 (Main Reduction)**.** *For any infinite family of number fields $\mathcal{K} = \{K_n\}$, $p \in [1, \infty]$, $m(n), q(n), \beta(n) = \mathrm{poly}(n)$ and $\gamma(n)$ that satisfy the conditions below, there is a polynomial-time number field-preserving reduction from solving $\mathcal{K}$-$\mathsf{IncIGVP}_\gamma^{p,\eta_\epsilon}$ (equivalently, $\mathcal{K}$-$\mathsf{IGVP}_\gamma^{p,\eta_\epsilon}$) in the worst case to solving $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}^\infty$ on the average with non-negligible probability.*

The conditions on the parameters are as follows:

1. *For $p \in [1, \infty)$, $\gamma(n) \geq c_p \cdot \beta(n) \cdot \sqrt{m(n)} \cdot n^{1/p}$ for a constant c_p depending only on p;*
   For $p = \infty$, $\gamma(n) \geq c_\infty \cdot \beta(n) \cdot \sqrt{m(n)} \cdot \sqrt{\log n}$ for a universal constant c_∞.
2. *$q(n) \geq 2 \cdot \beta(n) \cdot m(n) \cdot n \cdot g^\infty(\mathcal{O}_{K_n})$.*

Proof. The parameters ϵ, γ, m, q, and β are all functions of n, and the number fields K_n are from the family $\mathcal{K}$ indexed by n. For notational clarity we will often omit this dependence on n.

Advice about $\mathcal{K}$. For instances of K-$\mathsf{IncIGVP}$ where $K = K_n$ is a number field in $\mathcal{K}$, the advice about K needed by the reduction is an integral basis $B = \{b_1, \ldots, b_n\} \subset \mathcal{O}_K$ of K that is as short as possible in ℓ_∞ length. By Lemma 6.3, there exists such a basis B with $\|B\|_\infty = g^\infty(\mathcal{O}_K) \leq C \cdot \mathcal{D}_K \sqrt{\log n}$ for some constant C. Actually, it suffices merely to have $\|B\|_\infty = \mathrm{poly}(n)$, if we require $q(n)$ to be correspondingly larger.

Rounding in K. Our reduction will need to "round off" elements in $K = K_n$ to nearby (but not necessarily nearest) algebraic integers in $\mathcal{O}_K$. The rounding algorithm will take a $w \in K$ and the integral basis B of K discussed above, and will output some algebraic integer denoted $\lfloor w \rceil_B \in \mathcal{O}_K$. This can be accomplished by taking the representation of w in the basis B, $w = \sum_{i \in [n]} c_i b_i$ where $c_i \in \mathbb{Q}$, and rounding each coefficient to the nearest integer: $\lfloor w \rceil_B = \sum_{i \in [n]} \lfloor c_i \rceil b_i$. This algorithm outputs $\lfloor w \rceil_B \in \mathcal{O}_K$ such that $\|w - \lfloor w \rceil_B\|_\infty \leq \frac{n}{2} \cdot \|B\|_\infty$ (by the triangle inequality). We write $\lfloor w \rceil = \lfloor w \rceil_B$ when B is clear from context.

The reduction. Suppose oracle $\mathcal{F}$ solves the average-case problem $\mathcal{K}\text{-}\mathsf{SAIS}_{q,m,\beta}$ with non-negligible probability. We construct an algorithm to solve $\mathcal{K}\text{-}\mathsf{InclGVP}_{\gamma}^{p,\eta_\epsilon}$ as follows:

On input $(\mathcal{I}, x)$ where $\mathcal{I}$ is an ideal of $\mathcal{O}_K$ and $x \in \mathcal{I}$ with $\|x\|_p > \gamma \cdot \eta_\epsilon(\mathcal{I})$,

1. For $j = 1$ to m,
   - Sample a uniform $v_j \in \mathcal{I}/\langle x \rangle$.
   - Sample $y_j \sim D_s^K$, where $s = 2\|x\|_p / \gamma \geq 2\eta_\epsilon(\mathcal{I})$. Let $y'_j = y_j \bmod \mathcal{I}$.
   - Let $w_j = qx^{-1}(v_j + y'_j) \bmod \langle q \rangle$. Let $a_j = \lfloor w_j \rceil_B \bmod \langle q \rangle$.
2. Let $\mathbf{a} = (a_1, \ldots, a_m)$ and let $\mathbf{z} = (z_1, \ldots, z_m) \leftarrow \mathcal{F}(A)$. Output

$$x' = \sum_{j \in [m]} \left(\frac{x(w_j - \lfloor w_j \rceil)}{q} - y_j \right) \cdot z_j. \tag{1}$$

Analysis. The correctness of the reduction follows from several claims, which we state and prove in turn. In all of the claims, probabilities are taken over the randomness of the reduction and of $\mathcal{F}$.

Claim 9.2. *The probability that $\mathbf{z} \in \Psi_q(\mathbf{a})$ is non-negligible in n.*

Proof. It suffices to bound the statistical distance $\Delta(\mathbf{a}, \mathsf{U}^m(\mathcal{O}_K/\langle q \rangle))$ by $m \cdot \epsilon/2 = \nu(n)$. Each a_j is independent, so by the triangle inequality it suffices to bound $\Delta(a_j, \mathsf{U}(\mathcal{O}_K/\langle q \rangle))$ by $\epsilon/2$.

First, by Lemma 4.3, $\Delta(y'_j, K/\mathcal{I}) \leq \epsilon/2$ (i.e., y'_j is almost uniform over $K/\mathcal{I}$), and because v_j is uniform over $\mathcal{I}/\langle x \rangle$, we have $\Delta(v_j + y'_j, \mathsf{U}(K/\langle x \rangle)) \leq \epsilon/2$. Because $w_j = qx^{-1}(v_j + y'_j)$, we have $\Delta(w_j, \mathsf{U}(K/\langle q \rangle)) \leq \epsilon/2$. It follows by the description of the rounding algorithm that $\Delta(a_j, \mathsf{U}(\mathcal{O}_K/\langle q \rangle)) \leq \epsilon/2$, as desired. □

Claim 9.3. *If $\mathbf{z} \in \Psi_q(\mathbf{a})$, then $x' \in \mathcal{I}$.*

Proof. From Equation (1), we can rewrite x' as:

$$x' = \sum_{j \in [m]} \left(\frac{xw_j}{q} - y_j \right) \cdot z_j - x \cdot \sum_{j \in [m]} \frac{\lfloor w_j \rceil z_j}{q} \tag{2}$$

We start by analyzing the second term of Equation (2). By construction,

$$\sum_{j \in [m]} \lfloor w_j \rceil z_j = \sum_{j \in [m]} a_j z_j \bmod \langle q \rangle.$$

By hypothesis $\mathbf{z} \in \Psi_q(\mathbf{a})$, so $\sum a_j z_j \in \langle q \rangle$, and we conclude $x \cdot \sum \frac{\lfloor w_j \rceil z_j}{q} \in \langle x \rangle \subseteq \mathcal{I}$.

Now we turn to the first term of Equation (2). By definition of w_j,

$$\frac{xw_j}{q} = (v_j + y'_j) \bmod \langle x \rangle.$$

Therefore

$$\left(\frac{xw_j}{q} - y_j \right) \cdot z_j = (v_j + y'_j - y_j) \cdot z_j \bmod \langle x \rangle.$$

Both $v_j, y'_j - y_j \in \mathcal{I}$, and $z_j \in \mathcal{O}_K$ by hypothesis. Therefore $(v_j + y'_j - y_j) \cdot z_j \in \mathcal{I}$, and because $\langle x \rangle \subseteq \mathcal{I}$, we conclude that the first term of Equation (2) is also in $\mathcal{I}$. □

Claim 9.4. *Conditioned on* $\mathbf{z} \in \Psi_q(\mathbf{a})$, $\|x'\|_p \le \frac{\|x\|_p}{2}$ *with probability at least* $1/2$.

Proof. By rewriting Equation (1) and the triangle inequality, we have:

$$\|x'\|_p \le \sum_{j\in[m]} \left\| \frac{x(w_j - \lfloor w_j \rceil) z_j}{q} \right\|_p + \left\| \sum_{j\in[m]} y_j z_j \right\|_p. \tag{3}$$

We start by bounding the first summation of (3). For all $j \in [m]$, we have

$$\left\| \frac{x(w_j - \lfloor w_j \rceil) z_j}{q} \right\|_p \le \frac{1}{q} \|x\|_p \cdot \|w_j - \lfloor w_j \rceil\|_\infty \cdot \|z_j\|_\infty .$$

By the rounding algorithm, we have $\|w_j - \lfloor w_j \rceil\|_\infty \le \frac{n}{2} \|B\|_\infty$, and by hypothesis, $\|z_j\|_\infty \le \beta$. Then the first summation of (3) is at most

$$\|x\|_p \cdot \frac{\beta m n \|B\|_\infty}{2q}.$$

By hypothesis $q \ge 2\beta m n \|B\|_\infty$, so the quantity above is at most $\|x\|_p /4$ with certainty.

We now bound the second term from (3). By a now-standard argument [38, 36, 42, 33], conditioned on any value of y'_j, the value $y_j - y'_j$ is distributed according to $D_{\mathcal{I},s,-y'_j}$, and is independent of $\mathbf{a}$ and $\mathbf{z}$. We will bound the value $\left\|\sum_{j\in[m]} y_j z_j\right\|_p$ conditioned on *any* fixed values of y'_j, which by averaging implies the bound on the unconditioned value. By Lemma 6.7,

$$\Pr\left[\left\|\sum\nolimits_{j\in[m]} y_j z_j\right\|_p > L\right] = \Pr\left[\left\|\sum\nolimits_{j\in[m]} z_j \cdot ((y_j - y'_j) - (-y'_j))\right\|_p > L\right] \le 1/2,$$

where for $p \in [1, \infty)$,

$$L = c'_p \cdot s \cdot \beta \cdot \sqrt{m} \cdot n^{1/p}$$

(for $p = \infty$, the $n^{1/p}$ term is replaced by $\sqrt{\log n}$). Because $s = 2\|x\|_p /\gamma$ and by definition of γ, we have for $p \in [1, \infty)$,

$$L \le \frac{2c'_p \cdot \|x\|_p \cdot \beta \cdot \sqrt{m} \cdot n^{1/p}}{c_p \cdot \beta \cdot \sqrt{m} \cdot n^{1/p}} = \frac{\|x\|_p}{4}$$

for appropriate constant c_p, and similarly for $p = \infty$. Therefore $\left\|\sum_{j\in[m]} y_j z_j\right\|_p \le \|x\|_p /4$ with probability at least $1/2$, so by (3) the claim follows. □

Claim 9.5. *Conditioned on* $\mathbf{z} \in \Psi_q(\mathbf{a})$, $x' \ne 0$ *with overwhelming probability.*

Proof. The main idea is that $x' = 0$ if and only if a sample from $D_{\mathcal{I},s,c}$ hits *one particular* "bad" value. Lemma 4.6 guarantees that the probability of this event is negligibly small.

By definition of w_j,

$$\frac{x w_j}{q} = t_j + v_j + y'_j$$

for some $t_j \in \langle x \rangle$. Therefore

$$x' = 0 \iff \sum_{j\in[m]} \left(t_j + v_j + y'_j - y_j - \frac{x\lfloor w_j \rceil}{q} \right) \cdot z_j = 0.$$

Because $\mathbf{z} \neq 0$, there exists i such that $z_j \neq 0$; assume without loss of generality that $i = 1$. Then by rearranging, we get $x' = 0$ if and only if:

$$y_1 - y_1' = \left(t_1 + v_1 - \frac{x\lfloor w_j \rceil}{q}\right) + z_1^{-1} \sum_{i=2}^{m} \left(t_j + v_j + y_j' - y_j - \frac{x\lfloor w_j \rceil}{q}\right) \cdot z_j. \tag{4}$$

As in the proof of Claim 9.4, conditioned on the value of y_1', $y_1 - y_1'$ distributed according to $D_{\mathcal{I},s,-y_1'}$ and is independent of all other variables appearing in Equation (4). There are two cases: if the right-hand side of Equation (4) is not in $\mathcal{I}$, then the equation is satisfied with probability zero because the support of $D_{\mathcal{I},s,-y_1'}$ is $\mathcal{I}$. If the right-hand side of Equation (4) is in $\mathcal{I}$, then because $s \geq 2 \cdot \eta_\epsilon(\mathcal{I})$, Lemma 4.6 guarantees that the equation is only satisfied with probability $2^{-n} \cdot \frac{1+\epsilon}{1-\epsilon} = \nu(n)$. □

By Claims 9.2 through 9.5 and the union bound, the reduction solves $\mathcal{K}$-$\mathsf{IncIGVP}^{p,\eta_\epsilon}_{\gamma}$ with non-negligible probability. This can be amplified to overwhelming probability by standard repetition techniques for worst-case problems. Theorem 9.1 follows. □

Theorem 9.1 concerns the average-case problem $\mathsf{SAIS}^{\infty}_{q,m,\beta}$, whose solutions must have ℓ_∞ length at most β. We can also show that the more relaxed average-case problem $\mathsf{SAIS}^{r}_{q,m,\beta}$, whose solutions must have ℓ_r length at most $\beta \cdot (mn)^{1/r}$, is as hard as $\mathsf{ISVP}^{p}_{\gamma}$ for any $p \leq r$, for almost identical connection factors. This may make it easier to construct cryptographic applications.

Theorem 9.6. *Generalizing Theorem 9.1, there is a reduction from $\mathcal{K}$-$\mathsf{IncIGVP}^{p,\eta_\epsilon}_{\gamma}$ to $\mathcal{K}$-$\mathsf{SAIS}^{r}_{q,m,\beta}$ for any $1 \leq p \leq r < \infty$, subject to the following conditions on the parameters:*

1. *$\gamma(n) \geq c_p \cdot \beta(n) \cdot (m(n))^{\max\{1/2,1/r\}} \cdot n^{1/p}$ for a constant c_p depending only on p.*
2. *$q(n) \geq 2 \cdot \beta(n) \cdot (m(n) \cdot n)^{1+1/r} \cdot g^{\infty}(\mathcal{O}_{K_n})$.*

Proof. The reduction is the same as the one from the proof of Theorem 9.1, and most of the proof remains the same as well. The only change is to the proof of Claim 9.4.

To prove Claim 9.4, we need to analyze the two terms of (3). For the first term, we only need the fact that for all $j \in [m]$, $\|z_j\|_\infty \leq \|\mathbf{z}\|_\infty \leq \|\mathbf{z}\|_r$. By the definition of the problem $\mathcal{K}$-$\mathsf{SAIS}^{r}_{q,m,\beta}$, we have $\|\mathbf{z}\|_r \leq \beta \cdot (mn)^{1/r}$. By the (strengthened) hypothesis on q, the first term remains bounded by $\|x\|_p / 4$.

For the second term of (3), we use Lemma 6.8 rather than Lemma 6.7. By the hypothesis on γ, the second term also remains bounded by $\|x\|_p / 4$ with probability at least $1/2$. This completes the proof. □

9.1 Connection to ISVP

In this section we give a reduction from ISVP to SAIS (via IGVP), instantiating all the parameters from Theorem 9.1 asymptotically, and focusing especially on the role of the root discriminant.

Theorem 9.7. *For any infinite family $\mathcal{K} = \{K_n\}$ of number fields where $\mathcal{D}_{K_n} = \mathrm{poly}(n)$, any $p \in [1, \infty)$, and any $m(n) = \Theta(\log n)$, there exist*

$$q(n) = O(n \cdot \log^{1.5} n) \cdot \mathcal{D}_{K_n} \qquad \beta(n) = O(1) \cdot \sqrt{\mathcal{D}_{K_n}} \qquad \gamma(n) = O(\sqrt{\log n}) \cdot \mathcal{D}_{K_n}^{1.5}$$

such that solving $\mathcal{K}$-$\mathsf{ISVP}^{p}_{\gamma}$ in the worst case reduces to solving $\mathcal{K}$-$\mathsf{SAIS}^{\infty}_{q,m,\beta}$ on the average. For $p = \infty$, there exists $\gamma(n) = O(\log n) \cdot \mathcal{D}_{K_n}^{1.5}$ for which the same applies.

Proof. Assume that $p \in [1, \infty)$, and let $\epsilon = 2^{-n}$. By Lemma 6.3, there are integral bases B_n (of K_n) with $\|B_n\|_\infty = O(\sqrt{\log n}) \cdot \mathcal{D}_{K_n}$. To satisfy the conditions in Theorem 9.1, we can choose some

$$\begin{aligned} q(n) &= O(\beta(n) \cdot n \cdot \log^{1.5} n) \cdot \mathcal{D}_{K_n} \\ \beta(n) &= O(q(n)^{1/m(n)}) \cdot \sqrt{\mathcal{D}_{K_n}} = O(\mathrm{poly}(n)^{1/\log n}) \cdot \sqrt{\mathcal{D}_{K_n}} = O(1) \cdot \sqrt{\mathcal{D}_{K_n}}. \end{aligned}$$

Applying Theorem 9.1, $\mathcal{K}\text{-}\mathsf{IGVP}^{p,\eta_\epsilon}_{\gamma'}$ reduces to $\mathcal{K}\text{-}\mathsf{SAIS}^{\infty}_{q,m,\beta}$ for some $\gamma'(n) = O(n^{1/p}\sqrt{\log n}) \cdot \sqrt{\mathcal{D}_{K_n}}$.

We now connect $\mathcal{K}$-IGVP to $\mathcal{K}$-ISVP. By Lemma 6.2 and Lemma 6.5,

$$\lambda_1^p(\mathcal{I}) \geq n^{1/p} \cdot N^{1/n}(\mathcal{I}) \geq \frac{n^{1/p}}{\mathcal{D}_{K_n}} \cdot \eta_\epsilon(\mathcal{I})$$

for any ideal $\mathcal{I} \subseteq \mathcal{O}_{K_n}$. Therefore solving $\mathcal{K}\text{-}\mathsf{IGVP}^{p,\eta_\epsilon}_{\gamma'}$ also solves $\mathcal{K}\text{-}\mathsf{ISVP}^{p}_{\gamma}$ for some $\gamma(n) = O(\sqrt{\log n}) \cdot \mathcal{D}^{1.5}_{K_n}$. For $p = \infty$, a similar analysis applies. □

Corollary 9.8. *There exists an infinite family $\mathcal{K} = \{K_n\}$ of number fields such that for any $p \in [1, \infty)$ and any $m(n) = \Theta(\log n)$, there exist*

$$q(n) = O(n \log^{1.5} n) \qquad \beta = O(1) \qquad \gamma(n) = O(\sqrt{\log n})$$

such that solving $\mathcal{K}\text{-}\mathsf{SAIS}^{\infty}_{q,m,\beta}$ on the average with non-negligible probability is at least as hard as solving $\mathcal{K}\text{-}\mathsf{ISVP}^{p}_{\gamma}$ in the worst case. For $p = \infty$, there exists $\gamma(n) = O(\log n)$ for which the same claim applies.

Proof. Follows from Theorem 9.7 by choosing $\mathcal{K}$ to be a family such that $\limsup_{n\to\infty} \mathcal{D}_{K_n} = C$ for some constant C. As we have mentioned before, such families exist by the theory of infinite towers of Hilbert class fields (cf. [46]). □

10 Future Work

Our work opens up many avenues of investigation. The most important open problem, in our view, is the explicit construction of families of number fields having small root discriminant. By "explicit construction" we mean an *efficient* algorithm which, given n, outputs a full description of the degree-n number field from the family. Such constructions would also have applications in coding theory [32, 25]. It would be even nicer to find an explicit construction which provides, by design, the non-uniform advice that is needed by our reductions.

Another important problem is to better understand the worst-case *search* version of SVP for ideal lattices over number fields. The situation seems to be quite different from that of general lattices, because in our case the *decision* version of SVP is easy for factors as small as $\sqrt{\mathcal{D}_K}$. It would also be interesting to consider quantum algorithms for search-SVP on ideal lattices.

Our bound on the smoothing parameter for ideal lattices is most useful when the root discriminant of the number field is at most $O(\sqrt{n})$. Beyond that point, the prior bound relating the smoothing parameter to λ_n may be stronger [38]. We leave it as an open problem to unify these two bounds for ideal lattices, for the full range of interesting values of the root discriminant.

In contrast to prior work on average-case hardness from lattice problems (e.g., [4, 23, 38, 42, 33]), we do not yet know how to obtain cryptographic hardness (e.g. collision-resistant hash functions) from ideal lattices over an arbitrary good number field K. The reason is that we seem to require an

efficient injective mapping from function inputs (bit strings) to sufficiently short vectors $\mathbf{z} \in \mathcal{O}_K^m$. This appears to require some additional advice about the number field.

For example, if we had an "almost-orthogonal" short basis for $\mathcal{O}_K^m$, then it would be possible to produce elements of $\mathcal{O}_K^m$ that are short enough in ℓ_2 length, simply by summing up subsets of the basis vectors. This would suffice for cryptographic hardness, assuming the worst-case difficulty of SVP on ideal lattices in the ℓ_p norm for some $p \leq 2$.[7] As mentioned above, it is possible that explicit constructions of number fields might come with the required advice as a side-effect.

A final interesting question is whether the public-key cryptosystem of Regev [44] can be adapted to work based on ideal lattices, with a corresponding improvement in its efficiency and connection factor. It seems plausible that this could be done without requiring the encryption and decryption algorithms to use any special advice about the number field (though the security reduction might still require it).

11 Acknowledgements

We gratefully acknowledge Eva Bayer-Fluckiger, Dan Boneh, Henri Cohen, Noam Elkies, Alex Healy, Hendrik Lenstra, and Denis Simon for their help regarding algebraic number theory and ideal lattices. We also thank Vadim Lyubashevsky, Daniele Micciancio, and Oded Regev for useful discussions, and the anonymous reviewers for their detailed and constructive feedback.

References

[1] *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*. ACM, 2005.

[2] *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*. ACM, 2006.

[3] D. Aharonov and O. Regev. Lattice problems in NP ∩ coNP. *J. ACM*, 52(5):749–765, 2005.

[4] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.

[5] M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19, 1998.

[6] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *STOC*, pages 284–293, 1997.

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

[8] A. Akavia, O. Goldreich, S. Goldwasser, and D. Moshkovitz. On basing one-way functions on NP-hardness. In *STOC* [2], pages 701–710.

[7] For cryptography based on the assumed hardness of SVP in some *arbitrary* ℓ_p length, we would need a way to produce elements of $\mathcal{O}_K$ that are extremely short in the ℓ_∞ length. This seems more difficult to do.

[9] S. Alaca and K. S. Williams. *Introductory Algebraic Number Theory*. Cambridge University Press, November 2003.

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

[11] W. Banaszczyk. Inequalites for convex bodies and polar reciprocal lattices in R^n. *Discrete & Computational Geometry*, 13:217–231, 1995.

[12] E. Bayer-Fluckiger. Personal communication.

[13] E. Bayer-Fluckiger. *A Panorama of Number Theory Or The View from Baker's Garden*, chapter 11, pages 168–184. Cambridge University Press, September 2002.

[14] J. Bruck and M. Naor. The hardness of decoding linear codes with preprocessing. *IEEE Transactions on Information Theory*, 36(2):381–385, 1990.

[15] J.-Y. Cai. A new transference theorem in the geometry of numbers and new bounds for Ajtai's connection factor. *Discrete Applied Mathematics*, 126(1):9–31, 2003.

[16] J.-Y. Cai and A. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *FOCS*, pages 468–477, 1997.

[17] J.-Y. Cai and A. Nerurkar. Approximating the SVP to within a factor $(1+1/\dim^{\epsilon})$ is NP-hard under randomized reductions. *J. Comput. Syst. Sci.*, 59(2):221–239, 1999.

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

[19] H. Cohen, F. D. y Diaz, and M. Olivier. A table of totally complex number fields of small discriminants. In *ANTS*, volume 1423 of *Lecture Notes in Computer Science*, pages 381–391. Springer, 1998.

[20] J. H. Conway and N. J. a Sloane. *Sphere Packings, Lattices and Groups*. Springer, December 1998.

[21] U. Feige and D. Micciancio. The inapproximability of lattice and coding problems with preprocessing. *J. Comput. Syst. Sci.*, 69(1):45–67, 2004.

[22] O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000.

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

[24] O. Goldreich, D. Micciancio, S. Safra, and J.-P. Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Inf. Process. Lett.*, 71(2):55–61, 1999.

[25] V. Guruswami. Constructions of codes from number fields. *IEEE Transactions on Information Theory*, 49(3):594–603, 2003.

[26] D. Gutfreund and A. Ta-Shma. New connections between derandomization, worst-case complexity and average-case complexity. Technical Report TR06-108, Electronic Colloquium on Computational Complexity, 2006.

[27] J. H. W. Lenstra. Algorithms in algebraic number theory. *Bulletin of the American Mathematical Society*, 26(2):211–244, April 1992.

[28] S. Hallgren. Fast quantum algorithms for computing the unit group and class group of a number field. In *STOC* [1], pages 468–474.

[29] I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *STOC*, 2007.

[30] S. Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005.

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

[32] H. W. Lenstra. Codes from algebraic number fields. In M. Hazewinkel, J. K. Lenstra, and L. G. L. T. Meertens, editors, *Mathematics and computer science II, Fundamental contributions in the Netherlands since 1945*, CWI Monograph 4, pages 95–104. Elsevier, North-Holland, Amsterdam, 1986.

[33] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155. Springer, 2006. Full version in ECCC Report TR05-142.

[34] D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035, 2000.

[35] D. Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Transactions on Information Theory*, 47(3):1212–1215, 2001.

[36] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *FOCS*, pages 356–365. IEEE Computer Society, 2002. Full version in ECCC TR04-095.

[37] D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to ajtai's connection factor. *SIAM J. Comput.*, 34(1):118–169, 2004.

[38] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. In *FOCS*, pages 372–381. IEEE Computer Society, 2004.

[39] J. Milnor and D. Husemoller. *Symmetric Bilinear Forms*. Springer, 1973.

[40] R. A. Mollin. *Algebraic Number Theory*. CRC Press, 1999.

[41] C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. In *IEEE Conference on Computational Complexity*, 2007. Full version in ECCC Report TR06-148.

[42] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006. Full version in ECCC TR05-158.

[43] O. Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004.

[44] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC* [1], pages 84–93.

[45] O. Regev and R. Rosen. Lattice problems and norm embeddings. In *STOC* [2], pages 447–456.

[46] P. Roquette. *On class field towers*, chapter IX, pages 231–249. Academic Press, 1967.

[47] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[48] D. Simon. Personal communication.

A Proofs

Here we give proofs of Lemmas 6.7 and 6.8 on the sums of discrete Gaussians. Throughout this section, for any $p \in [1, \infty]$ we let c_p denote an appropriate constant depending only on p which may vary from one expression to the next (to be concrete, c_p is proportional to $\sqrt{p}$ for finite p).

We start by stating our main tool, which is an analysis of discrete Gaussians from [41]. Let $\mathbf{U} = \{\mathbf{u}_1, \ldots, \mathbf{u}_d\} \subset \mathbb{C}^N$ denote a set of d orthonormal vectors for $d = 1$ or $d = 2$. Define the "$\mathbf{U}$ norm" on $\mathbb{C}^N$ as $\|\mathbf{x}\|_{\mathbf{U}} = \sum_{i \in [d]} |\langle \mathbf{x}, \mathbf{u}_i \rangle|$.

Proposition A.1 ([41])**.** *For any lattice Λ in $\mathbb{C}^N$, $p \in [1, \infty)$, $\mathbf{c} \in \mathbb{C}^n$, and $\mathbf{U}$ as above,*

$$\mathop{\mathrm{E}}_{\mathbf{x} \sim D_{\Lambda, \mathbf{c}}} \left[\|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}}^p \right] \leq c_p^p \cdot \frac{\rho(\Lambda)}{\rho_{\mathbf{c}}(\Lambda)}.$$

(Technically, the lemma from [41] is stated in terms of full-rank lattices in $\mathbb{R}^N$. However, the subspace $\mathrm{span}(\Lambda) \subseteq \mathbb{C}^N$ with inner product is isomorphic to $\mathbb{R}^{N'}$ for some N'; this is sufficient for the proof to go through in our setting.)

We now recall the notation from Lemmas 6.7 and 6.8, which we keep throughout this section: $\mathcal{I}$ is a fractional ideal over a number field K of degree n; m is a positive integer; $\epsilon \leq 1/(2m+1)$ and $s \geq \eta_\epsilon(\mathcal{I})$ are positive reals; $\mathbf{z}, \mathbf{c} \in K^m$ are arbitrary fixed vectors; $y_j \sim D_{\mathcal{I}, s, c_j}$ are independent samples from discrete Gaussians over $\mathcal{I}$, and we define

$$v = \sum_{j \in [m]} z_j \cdot (y_j - c_j).$$

The goal is to show tail bounds for $\|v\|_p$ assuming an appropriate upper bound on $\|\mathbf{z}\|_r$. For simplicity, we will assume without loss of generality that $s = 1 \geq \eta_\epsilon(\mathcal{I})$ throughout this section; the general results all follow by appropriate scaling.

Recall also that σ_i is the ith embedding of K, and $\sigma : K \to \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ is the canonical embedding of K, whose image is contained in the n-dimensional subspace $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$.

We start by deriving a bound on the moments of $|\sigma_i(v)|$, which will be needed in the proofs of both Lemma 6.7 and Lemma 6.8. For each $i \in [n]$, define $\mathbf{r}_i = (\sigma_i(z_1), \ldots, \sigma_i(z_m)) \in \mathbb{C}^m$.

Lemma A.2. *For any $p \in [1, \infty)$ and every $i \in [n]$, we have:*

$$\mathrm{E}\left[|\sigma_i(v)|^p\right] \leq c_p^p \cdot \|\mathbf{r}_i\|_2^p .$$

Proof. We first start by setting up some notation for the application of Proposition A.1. Just as in Section 6.3, define $\sigma : K^m \to (\mathbb{R}^{r_1} \times \mathbb{C}^{2r_2})^m$ via coordinate-wise application of σ. Then let $\Lambda' = \sigma(\mathcal{I} \times \cdots \times \mathcal{I}) \subset H^m$ be the Cartesian product of m copies of $\sigma(\mathcal{I})$, forming a lattice of rank mn. Likewise, let $\mathbf{c}' = \sigma(c_1, \ldots, c_m) \in H^m$, and $\mathbf{y}' = \sigma(y_1, \ldots, y_m) \in H^m$.

By some routine calculations (see [41, Proof of Corollary 5.2]), one can show that $\mathbf{y}'$ is distributed according to $D_{\Lambda', \mathbf{c}'}$, and that from the hypothesis $\epsilon \leq 1/(1+2m)$ we have

$$\frac{\rho(\Lambda')}{\rho_{\mathbf{c}'}(\Lambda')} \leq \left(1 + \tfrac{1}{m}\right)^m < \exp(1).$$

We now separate the proof into two cases, based on whether σ_i is a real embedding ($i \in [r_1]$) or a complex embedding ($r_1 < i \leq n$). The former case is somewhat simpler, so we start with it.

Define $\mathbf{w} \in H$ by $w_i = 1$, and $w_k = 0$ for all $k \neq i$. Then for any $x \in K$, we have $\sigma_i(x) = \langle \sigma(x), \mathbf{w} \rangle$. Therefore

$$\begin{aligned}
\sigma_i(v) &= \sum_{j \in [m]} \sigma_i(z_j) \cdot \sigma_i(y_j - c_j) \\
&= \sum_{j \in [m]} \langle \sigma(y_j - c_j), \sigma_i(z_j) \cdot \mathbf{w} \rangle \\
&= \Big\langle \mathbf{y}' - \mathbf{c}', \underbrace{(\sigma_i(z_1) \cdot \mathbf{w},\ \ldots\ , \sigma_i(z_m) \cdot \mathbf{w})}_{\tilde{\mathbf{w}} = r \cdot \mathbf{u} \in H^m} \Big\rangle \\
&= r \left\langle \mathbf{y}' - \mathbf{c}', \mathbf{u} \right\rangle ,
\end{aligned}$$

where $\mathbf{u} \in H^m$ is the unit vector parallel to $\tilde{\mathbf{w}}$, and $r^2 = \langle \tilde{\mathbf{w}}, \tilde{\mathbf{w}} \rangle = \sum_{j \in [m]} \sigma_i^2(z_j) = \|\mathbf{r}_i\|_2^2$. Now set $\mathbf{U} = \{\mathbf{u}\}$, so $|\sigma_i(v)| = \|\mathbf{r}_i\|_2 \cdot \|\mathbf{y}' - \mathbf{c}'\|_{\mathbf{U}}$. By Proposition A.1, we get the desired bound on $\mathrm{E}\left[|\sigma_i(v)|^p\right]$.

We now turn to the case that σ_i is a complex embedding. We may assume that $r_1 < i \leq r_1 + r_2$, for if $i > r_1 + r_2$, we have $\sigma_i(v) = \overline{\sigma_{i-r_2}(v)}$ and hence $|\sigma_i(v)| = |\sigma_{i-r_2}(v)|$. Now define $\mathbf{w} \in H$ to be the vector that "selects" the real part of the ith coordinate of any $\mathbf{x} \in H$ via Hermitian inner product, i.e., $\langle \mathbf{x}, \mathbf{w} \rangle = \mathrm{Re}\{x_i\}$. Concretely, $w_i = w_{i+r_2} = 1/2$ and $w_k = 0$ otherwise.

As above, we can show that $\mathrm{Re}\{\sigma_i(v)\} = \langle \mathbf{y}' - \mathbf{c}', \tilde{\mathbf{w}} \rangle = r \langle \mathbf{y}' - \mathbf{c}', \mathbf{u} \rangle$ where $\mathbf{u} \in H^m$ is the unit vector parallel to appropriate $\tilde{\mathbf{w}} \in H^m$ and

$$r^2 = \langle \tilde{\mathbf{w}}, \tilde{\mathbf{w}} \rangle = \sum_{j \in [m]} \left(\left| \frac{\sigma_i(z_j)}{2} \right|^2 + \left| \frac{\sigma_{i+r_2}(z_j)}{2} \right|^2 \right) = \sum_{j \in [m]} \frac{|\sigma_i(z_j)|^2}{2} = \frac{\|\mathbf{r}_i\|_2^2}{2}.$$

Similarly, let $\mathbf{w}' \in H$ be the vector that selects the imaginary part of the ith coordinate of $\mathbf{x} \in H$, i.e. $\langle \mathbf{x}, \mathbf{w}' \rangle = \mathrm{Im}\{x_i\}$. Concretely, $w'_i = -w'_{i+r_2} = \sqrt{-1}/2$, and $w_k = 0$ otherwise. Then as above, we have $\mathrm{Im}\{\sigma_i(v)\} = r \langle \mathbf{y}' - \mathbf{c}', \mathbf{u}' \rangle$, where $r^2 = \|\mathbf{r}_i\|_2^2 / 2$. Furthermore, one can check that $\mathbf{u}$ and $\mathbf{u}'$ are orthogonal (this stems from the fact that $\mathbf{w}$ and $\mathbf{w}'$ are orthogonal).

Now set $\mathbf{U} = \{\mathbf{u}, \mathbf{u}'\}$. We have $|\sigma_i(v)| \leq |\mathrm{Re}\{\sigma_i(v)\}| + |\mathrm{Im}\{\sigma_i(v)\}| = \frac{\|\mathbf{r}_i\|_2}{\sqrt{2}} \cdot \|\mathbf{y}' - \mathbf{c}'\|_{\mathbf{U}}$. Applying Proposition A.1, we get the desired bound on $\mathrm{E}\left[|\sigma_i(v)|^p\right]$. □

For analyzing $\|v\|_\infty$, we will also need the following:

Lemma A.3. *For every $i \in [n]$, we have*

$$\Pr\left[|\sigma_i(v)| > c_\infty \cdot \|\mathbf{r}_i\|_2 \cdot \sqrt{\log n}\right] \leq \frac{1}{2n}.$$

Proof. The proof is almost identical to the one above, but instead uses a tail inequality on discrete Gaussians [41, Lemma 5.3]. □

We now prove Lemmas 6.7 and 6.8, which we restate below (with $s = 1$) for convenience.

Lemma A.4 (Restatement of Lemma 6.7)**.** *If $\|\mathbf{z}\|_\infty \leq \beta$, then $\Pr[\|v\|_p > L] \leq \frac{1}{2}$, where*

$$L = c_p \cdot \beta \cdot \sqrt{m} \cdot \begin{cases} n^{1/p} & \text{for} \quad p \in [1, \infty) \\ \sqrt{\log n} & \text{for} \quad p = \infty \end{cases}$$

Proof. We start with the case $p \in [1, \infty)$. We have:

$$\begin{aligned} \mathrm{E}\left[\|v\|_p^p\right] &= \sum_{i \in [n]} \mathrm{E}\left[|\sigma_i(v)|^p\right] \\ &\leq c_p^p \cdot \sum_{i \in [n]} \|\mathbf{r}_i\|_2^p && \text{(Lemma A.2)} \\ &\leq c_p^p \cdot \sum_{i \in [n]} (\sqrt{m})^p \cdot \|\mathbf{r}_i\|_\infty^p && (\|\mathbf{r}_i\|_2 \leq \sqrt{m} \cdot \|\mathbf{r}_i\|_\infty) \\ &\leq c_p^p \cdot (\sqrt{m})^p \cdot \beta^p \cdot n. && (\|\mathbf{r}_i\|_\infty \leq \|\mathbf{z}\|_\infty \leq \beta) \end{aligned}$$

Therefore by Jensen's inequality, $\mathrm{E}\left[\|v\|_p\right] \leq c_p \cdot \beta \cdot \sqrt{m} \cdot n^{1/p}$. Finally, Markov's inequality yields the claim.

Now we consider the case $p = \infty$. Because $\|\mathbf{r}_i\|_2 \leq \sqrt{m} \cdot \|\mathbf{r}_i\|_\infty \leq \sqrt{m} \cdot \beta$, by Lemma A.3 we have $\Pr[|\sigma_i(v)| > L] \leq 1/2n$. By the union bound over all $i \in [n]$, the claim follows. □

Lemma A.5 (Restatement of Lemma 6.8)**.** *If $\|\mathbf{z}\|_r \leq \beta \cdot (mn)^{1/r}$ for some $r \in [1, \infty)$, then for any $p \leq r$ we have $\Pr[\|v\|_p > L] \leq \frac{1}{2}$, where*

$$L = c_p \cdot \beta \cdot m^{\max\{1/2, 1/r\}} \cdot n^{1/p}.$$

Proof. We recall Hölder's inequality, which implies that for $1 \leq p \leq r \leq \infty$ and any $\mathbf{x} \in \mathbb{C}^N$, we have $\|\mathbf{x}\|_p \leq N^{1/p - 1/r} \cdot \|\mathbf{x}\|_r$.

First, define $\mathbf{t} = (\|\mathbf{r}_1\|_r, \ldots, \|\mathbf{r}_n\|_r) \in \mathbb{R}^n$. Observe that

$$\|\mathbf{t}\|_r = \left(\sum_{i \in [n]} \|\mathbf{r}_i\|_r^r\right)^{1/r} = \left(\sum_{i \in [n], j \in [m]} |\sigma_i(z_j)|^r\right)^{1/r} = \|\mathbf{z}\|_r \leq \beta \cdot (mn)^{1/r}.$$

Now consider the case $r < 2$. By standard properties of norms, we have $\|\mathbf{r}_i\|_2 \leq \|\mathbf{r}_i\|_r$. Then:

$$\begin{aligned}
\mathrm{E}\left[\|v\|_p\right] &\leq \left(\mathrm{E}\left[\|v\|_p^p\right]\right)^{1/p} && \text{(Jensen's inequality)} \\
&= \left(\sum\nolimits_{i\in[n]} \mathrm{E}\left[|\sigma_i(v)|^p\right]\right)^{1/p} && \text{(linearity of E)} \\
&\leq c_p \cdot \left(\sum\nolimits_{i\in[n]} \|\mathbf{r}_i\|_2^p\right)^{1/p} && \text{(Lemma A.2)} \\
&\leq c_p \cdot \left(\sum\nolimits_{i\in[n]} \|\mathbf{r}_i\|_r^p\right)^{1/p} && \text{(properties of norms)} \\
&\leq c_p \cdot \|\mathbf{t}\|_p && \text{(definition of } \mathbf{t}\text{)} \\
&\leq c_p \cdot n^{1/p-1/r} \cdot \|\mathbf{t}\|_r && \text{(Hölder's inequality, } p \leq r\text{)} \\
&\leq c_p \cdot \beta \cdot m^{1/r} \cdot n^{1/p} && \text{(above bound on } \|\mathbf{t}\|_r = \|\mathbf{z}\|_r\text{)}
\end{aligned}$$

Now consider $r \geq 2$. By Hölder's inequality, we have $\|\mathbf{r}_i\|_2 \leq m^{1/2-1/r} \cdot \|\mathbf{r}_i\|_r$. By an argument nearly identical to the one above, we have

$$\mathrm{E}\left[\|v\|_p\right] \quad \leq \quad c_p \cdot m^{1/2-1/r} \cdot \|\mathbf{t}\|_p \quad \leq \quad c_p \cdot \beta \cdot m^{1/2} \cdot n^{1/p}.$$

Finally, the desired result follows by Markov's inequality. □