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

Chris Peikert* Alon Rosen†

November 26, 2006

Abstract

We demonstrate an *average-case* problem which 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 class of lattices was $\gamma(n) = \tilde{O}(n)$.

To obtain our results, we focus on 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 assumption we rely on is that in some ℓ_p length, it is hard to find approximate shortest vectors in these lattices, under an appropriate form of preprocessing of the number field. Our results build upon prior works by Micciancio (FOCS 2002), Peikert and Rosen (TCC 2006), and Lyubashevsky and Micciancio (ICALP 2006).

For the connection factors $\gamma(n)$ we achieve, the corresponding *decisional* promise 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 the current best of $\tilde{O}(n)$.

*SRI International, cpeikert@alum.mit.edu

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

1 Introduction

In 1996, Ajtai established a remarkable connection between the worst-case hardness and average-case hardness of certain lattice problems [2]. Ajtai showed that there is some polynomial $\gamma(n)$ for which $\gamma(n)$-approximating the length of the shortest vector in n-dimensional lattices in the *worst-case* reduces to solving a related computational problem *on the average*.

Soon thereafter, Ajtai also showed that the shortest vector problem is NP-hard under randomized reductions [3]. This spawned the hope that average-case hardness could someday be based on the assumption that P $\neq$ NP. The plan was to decrease the worst-case/average-case *connection factor* $\gamma(n)$ to a point at which the corresponding approximation problem is NP-hard. In pursuit of this goal, the connection factor was successively tightened [15, 35], and NP-hardness was established for increasingly large approximation factors [16, 32].

The current state of the art is defined by two powerful results: the first, by Micciancio and Regev, establishes a connection factor of $\gamma(n) = \tilde{O}(n)$ [36]. The second, by Khot, establishes the NP-hardness of approximating the shortest vector to within *any* constant factor [27]. The latter result already approaches the perceived limits on the hardness of approximating the shortest vector, as NP-hardness beyond a certain $\Omega(\sqrt{n})$ factor would imply that NP $\subseteq$ coNP [1] (or NP $\subseteq$ coAM, for a certain $\Omega(\sqrt{n/\log n})$ factor [21]).

Worst-case/average-case connections are also useful in arenas outside complexity theory. Ajtai's result and its successors go *beyond* average-case hardness, in that they actually yield cryptographic one-way functions and collision-resistant hash functions [22]. Even public-key encryption is attainable from certain worst-case hardness assumptions on lattices [4, 41, 42].

These cryptographic applications introduce another, more pragmatic motivation for tightening the connection factor. The best known polynomial-time shortest vector algorithms produce only a $2^{\tilde{\Omega}(n)}$-approximate solution [29, 45], whereas the best algorithm for finding an optimal solution takes $2^{O(n)}$ time [5]. In addition, there are algorithms that allow trade-offs between running time and quality of approximation [45, 28]. In practice, then, a loose connection factor may fail to guarantee security for realistic values of the dimension. Indeed, one of the critiques of lattice-based cryptography is that the known lattice algorithms require the use of prohibitively large concrete parameters.

It is thus clear that tightening the connection factor is an important goal, from both a practical and theoretical point of view. One should keep in mind that it would be meaningful to obtain average-case hardness from any worst-case problem whose actual time complexity is large (e.g., exponential), *even if the problem is not NP-hard*.[1] In light of this, any approach that would yield a tighter connection factor, without compromising on the concrete hardness of the worst-case problem, would be interesting and useful.

1.1 Our Results

We open a new avenue for obtaining worst-case/average-case lattice reductions with *very small* connection factors. To obtain these reductions, we shift the focus from general lattices to certain 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. Our worst-case assumption is that in some ℓ_p length, it is hard to find approximate shortest vectors in these lattices, under an appropriate form of preprocessing of the number field.

[1] In fact, giving up on NP-hardness might even be necessary for constructing certain cryptographic primitives [12, 6].

For the connection factors we achieve, the corresponding *decisional* promise 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 [2] and the improvements proposed by Micciancio and Regev [36], as well as the works of Micciancio [34], Peikert and Rosen [40], and Lyubashevsky and Micciancio [31]. The latter works obtained efficient cryptographic primitives by generalizing the role of the integers in prior reductions, replacing them with elements from some larger ring. (See Section 2 for details.) We show that by substituting the integers with the more general notion 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.

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

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 [44], though no *efficient* construction is yet known. 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 that admit a connection factor better than the current best of $\tilde{O}(n)$.

1.2 Explicitness

One apparent drawback of our work is that it is still unknown how to *efficiently* compute families of number fields with 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 *fixed* degrees (see, e.g. [18]). As far as we can tell, the problem of *efficiently constructing* good *asymptotic* families of number field has not received nearly as much attention. The best we know of is an infinite family of *cyclotomic* number fields having root discriminants as small as $O(n(\log\log n)/(\log n))$ [46]. As mentioned above, families having root discriminants even up to $O(n^{2/3-\epsilon})$ would yield improved connection factors.

In some sense, the current state of affairs is not unlike the early days of coding theory, or even the era in which expander graphs had many promising applications in theoretical computer science, but explicit constructions were yet to be discovered. Just as with these examples, we are hopeful that with enough effort, explicit constructions of good number fields will eventually be found.

1.3 Uniformity

Our reductions also 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 [13, 33, 20], 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 this mapping appears to require some special information about the number field.

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 possibly enabling cryptographic hardness. This is indeed the case for the cyclotomic number fields mentioned above.

1.4 The Worst-Case Assumption

Our worst-case assumption is that it is asymptotically hard to *find* approximately shortest non-zero vectors in ideal lattices over certain families of number fields of increasing degree. This should hold even in the face of arbitrary preprocessing of the *number field*.

We note that due to the algebraic structure of ideal lattices over number fields, it is actually trivial to closely approximate the *length* λ_1 of a shortest vector. Fortunately, there is no known reduction from the search problem to the corresponding (easy) decisional approximation problem.

Finding short elements 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 LLL algorithm [29]. This 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 (e.g., this is a reason why the recent quantum algorithm of Hallgren [26] is limited to *fixed* degree). Any efficient algorithm for finding a short element in ideal lattices in the worst case would be considered a major breakthrough in computational number theory [46, 10]. Finally, the LLL and related algorithms for general lattices seem to perform no better on ideal lattices.

It is hard to qualitatively compare our results with the known results on general lattices. On the one hand, our worst-case assumption is restricted to a subclass of lattices and hence could be seen as a stronger assumption than on general lattices. On the other hand, the approximation factor in our assumption is substantially smaller.

1.5 Additional Contributions

We additionally give reductions among 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 exact decisional version of CVP. Analogous results were already known for general lattices [23], however these reductions do not preserve the "ideal" structure of their input lattices. (That is, the instances generated by the reduction are not necessarily ideal lattices, even if the input lattice is ideal.) Our new reductions rest upon 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 minima, and covering radius, in arbitrary ℓ_p lengths. We also give a new bound on the smoothing parameter which, for lattices over number fields with small root discriminant, is significantly stronger than a prior bound [36].

We also point out that number fields with constant root discriminant give rise to a large collection of lattices which exemplify the tightness (up to constant factor) of known *transference theorems* on lattices, in *all* ℓ_p lengths [8, 9]. This gives an alternative to a prior example by Conway and Thompson [37] for ℓ_2 lengths.

1.6 Related Work

The idea of imposing special structure on lattices is not new. Some of the results in Ajtai's original paper [2] are based on lattices whose shortest vectors are "unique" in some formal sense, and the same applies for cryptosystems of Ajtai and Dwork [4] and Regev [41].

Micciancio exploited *cyclic* lattices to obtain a very efficient one-way function [34]. In later independent works, Peikert and Rosen [40] and Lyubashevsky and Micciancio [31] extended Micciancio's result to obtain efficient collision-resistant hash functions, by recognizing and exploiting the underlying algebraic structure of cyclic lattices.

All of the above results were primarily focused on obtaining cryptographic primitives with additional functionality or efficiency. The best connection factor achieved by any of them is $\tilde{O}(n)$. To the best of our knowledge, we are the first to propose a plausible worst-case assumption under which one can obtain sub-linear (and even dramatically lower) connection factors.

Our lattices are somewhat related to a generalization of cyclic lattices proposed by Lyubashevsky and Micciancio [31]. They called these objects ideal lattices, and pointed out some connections to algebraic number theory. Ideal lattices were actually already conceived of in the realm of algebraic number theory, though under a different definition. This work employs the latter definition, which will prove crucial in understanding the geometric structure of ideal lattices, and in obtaining our improved connection factors. See Section 2 for details.

1.7 Future Work

Our work opens up many interesting questions. 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 an explicit description of the degree-n number field from the family. Such constructions would also have applications in coding theory [30, 24]. It would be even nicer to find an explicit construction which provides, by design, the non-uniform advice that is needed by our reductions. A promising starting point is a construction due to Simon [47] of a family of *polynomials* of degree n whose root discriminants grow as $O(\sqrt{n})$. These polynomials are highly factorizable (not irreducible), hence they do not yield number fields, but products of number fields. This opens up at least two possibilities: either the construction can be adapted to yield irreducible polynomials, or our construction can be meaningfully adapted to products of number fields.

Another important problem is to better understand the worst-case hardness of the *search* problems we rely upon. The situation seems to be quite different from that of general lattices, because in our case the corresponding *decisional* problems are easy. See Section 1.4 for a discussion.

Our bound on the smoothing parameter of ideal lattices is most useful when the root discriminant is $O(\sqrt{n})$. Beyond that point, the prior bound relating the smoothing parameter to λ_n may be stronger [36]. 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.

2 Overview of Techniques

In this section we review the basic concepts behind worst-case/average-case reductions for lattices, the prior work most similar to ours, and how we obtain our improvements.

2.1 Ajtai's Framework

Similarly to most previous works on worst-case/average-case reductions [22, 15, 35, 41], we follow the framework initiated by Ajtai [2]. This framework shows how to reduce worst-case instances of lattice problems to finding "small" solutions to random instances of certain linear group equations. In order to perform such a reduction, one must sample random elements from the group in a way that is related to the original lattice problem. The heart of this is a method for sampling pairs consisting of a group element along with a corresponding short "offset" vector. Sampling several pairs yields a set of group elements that define the equation to be solved. The key property of the sampling procedure is this: linearly combining the offset vectors via *any* solution to the group equation yields a point in the original lattice.

In Ajtai's case, the group is $\mathbb{Z}_q^n$, where n corresponds to the dimension of the lattice and q is a (small) modulus. A solution to the group equation is simply a vector of integer coefficients. The length of the resulting lattice point is therefore governed by two quantities: the size of the coefficients in the solution vector, and the lengths of the offset vectors.

Micciancio and Regev [36] proposed a very powerful and elegant method of implementing the sampling procedure, which yields very short offset vectors. This method is based on Gaussian measures over lattices, and its performance depends on a lattice quantity which they called the *smoothing parameter*. This parameter is essentially bounded by the nth successive minimum λ_n of the lattice (times a small extra factor), and the offset vectors generated by the sampling procedure have length essentially $\sqrt{n} \cdot \lambda_n$. Because the solution vectors also have length essentially $\sqrt{n}$, the resulting lattice points have length $\tilde{O}(n) \cdot \lambda_n$. Using several additional ideas, these points can be used to approximate the length λ_1 of the shortest vector to within an $\tilde{O}(n)$ factor.

2.2 Ideal Lattices

The use of ideal lattices in our reduction is best understood in the context of a series of prior works. Micciancio [34] proposed generalizing Ajtai's framework by taking the average-case solution vectors to be over the larger ring $\mathbb{Z}_q[x]/\langle x^n - 1\rangle$, rather than $\mathbb{Z}_q$. This yielded "compact" average-case hardness and efficient one-way functions based on *cyclic* lattices.

In later independent works, Peikert and Rosen [40] and Lyubashevsky and Micciancio [31] observed that cyclic lattices are actually *ideals* in the ring $R = \mathbb{Z}[x]/\langle x^n - 1\rangle$, and obtained efficient cryptographic hash functions by exploiting additional algebraic structure of this ring. The latter work also suggested generalizing the ring R to be $\mathbb{Z}[x]/\langle f(x)\rangle$ for any monic irreducible integer polynomial $f(x)$ of degree n having small "expansion factor." The relevant lattices still correspond to ideals in R, where the lattice points are the coefficient vectors of the polynomial residues in the ideal. The solution vectors of the average-case problem are over the ring $R/\langle q\rangle$. All of these works achieved connection factors of $\tilde{O}(n)$, with slight variations in the hidden constants.

2.3 This Work

We retain the same basic structure as above, using a different kind of ring R and, just as importantly, a different method of constructing lattices from ideals. We take R to be the ring $\mathcal{O}_K$ of *algebraic integers* in a number field K of degree n. Ideals in $\mathcal{O}_K$ correspond to n-dimensional lattices via the n *embeddings* of the number field K into the complex numbers $\mathbb{C}$. This notion of an ideal lattice is actually a natural and standard one from algebraic number theory and the geometry of numbers;

see e.g. [19, Chapter 8], [11]. More importantly, this formulation will prove crucial to understanding the geometric structure of these lattices, and in obtaining much tighter connection factors.

From the discussion above, we can now define our average case problem. An instance is given by a uniformly random vector $(a_1, \ldots, a_m)$ where $a_i \in R/\langle q \rangle = \mathcal{O}_K/\langle q \rangle$. This defines an equation

$$\sum_{i=1}^{m} a_i z_i = 0 \bmod \langle q \rangle$$

in the ring $\mathcal{O}_K/\langle q \rangle$. The problem is to find a "small" solution to this equation, where the size of the solution is defined by the ℓ_∞ lengths of the embeddings of the elements z_i.

To obtain the tightest connection factors for our reduction, we will employ number fields K of degree n whose *root discriminants* $\mathcal{D}_K$ are as small as possible as a function of n. The small root discriminant will yield gains in two separate aspects of the reduction:

- *Shorter average-case solutions.* The average-case problem will admit solutions which are very "densely-packed" in space. This allows us to show that there exist *very short* solutions — even as short as (almost) constant in length. This is in contrast to all prior work, in which the lengths of the solutions were $\tilde{\Omega}(\sqrt{n})$ due to their inherent spareseness.

- *An improved bound on the smoothing parameter.* We show that in ideal lattices over number fields with small root discriminant, the smoothing parameter is actually an $\Omega(\sqrt{n})$ factor *smaller than* than the first successive minimum λ_1. This is a significant improvement over the prior bound of essentially λ_n [36].

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. A basic analysis of our reduction, using a result by Lyubashevsky and Micciancio [31] on the moments of discrete Gaussians, allows us to obtain an $\mathcal{D}_K^{1.5} \cdot O(\log^2 n)$ bound on the connection factor. A tighter analysis, relying on a result from concurrent work by Peikert [39] on the moments of *sums* of discrete Gaussians, yields an $\mathcal{D}_K^{1.5} \cdot O(\sqrt{\log n})$ bound on the connection factor.

Our treatment of general ℓ_p lengths is partly motivated by a recent result of Regev and Rosen [43], who showed that worst-case lattice problems are, at least for general lattices, at their easiest in ℓ_2 length. In light of this fact, obtaining reductions for arbitrary ℓ_p lengths under a unified connection factor is much more desirable. In addition, certain notions in algebraic number theory correspond to different ℓ_p lengths, e.g. ℓ_∞ length for "height" and often ℓ_1 length for "size."

3 Preliminaries

3.1 Notation

The set of real numbers is denoted by $\mathbb{R}$, the complex numbers by $\mathbb{C}$, the rationals by $\mathbb{Q}$, and the integers by $\mathbb{Z}$. For a positive integer n, $[n]$ denotes $\{1, \ldots, n\}$. For a real r, $\lfloor r \rceil = \lfloor r + \frac{1}{2} \rfloor$ denotes a closest integer to r. The function log will always denote the natural logarithm.

For a real a, we write $[a, \infty]$ for the set $[a, \infty) \cup \{\infty\}$. For simplicity, we use the following conventions: $\sqrt[\infty]{n} = 1$ for any positive n; $1/\infty = 0$; and $1/0 = \infty$.

A vector in $\mathbb{R}^n$ or $\mathbb{C}^n$ is represented in column form, and written as a bold lower-case letter, e.g. $\mathbf{x}$. For a vector $\mathbf{x}$, the ith component of x will be denoted by x_i, or when such notation would be confusing, by $(\mathbf{x})_i$. Matrices are written as bold capital letters, e.g. $\mathbf{B}$. The ith column vector of $\mathbf{B}$ is denoted $\mathbf{b}_i$.

For $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$, the standard Hermitian inner product is defined to be $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i \in [n]} x_i \overline{y_i}$, where $\overline{y_i}$ denotes the complex conjugate of y_i. When $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, the Hermitian inner product $\langle \mathbf{x}, \mathbf{y} \rangle$ specializes to the standard inner product.

For $p \in [1, \infty)$, the ℓ_p length of a vector $\mathbf{x} \in \mathbb{R}^n$, denoted $\|\mathbf{x}\|_p$, is defined as $(\sum_{i \in [n]} |x_i|^p)^{1/p}$.[2] For $p = \infty$, the ℓ_∞ length of $\mathbf{x}$ is defined as $\|\mathbf{x}\|_\infty = \max_{i \in [n]} |x_i|$.

Later we will also define ℓ_p lengths over domains other than $\mathbb{R}^n$. For any finite set $V = \{v_1, \ldots, v_n\}$ or tuple $V = (v_1, \ldots, v_n)$ of elements from a domain having an ℓ_p length, define $\|V\|_p = \max_{i \in [n]} \|v_i\|_p$. For any element t and set V from a domain having an ℓ_p length, define $\mathrm{dist}^p(t, V) = \inf_{v \in V} \|t - v\|_p$. We take $p = 2$ whenever it is omitted from any expression.

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 $\mathsf{U}(S)$.

3.2 Lattices

A *(full-rank) lattice* in $\mathbb{R}^n$ is the set of all integer combinations

$$\Lambda = \left\{ \sum_{i \in [n]} c_i \mathbf{b}_i \mid c_i \in \mathbb{Z} \text{ for } 1 \le i \le n \right\}$$

of n independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^n$. The set of vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n$, often written in matrix form as $\mathbf{B} = [\mathbf{b}_1 | \cdots | \mathbf{b}_n]$ with the basis vectors as columns, is called a *basis* for the lattice. The lattice generated by $\mathbf{B}$ is denoted $\mathcal{L}(\mathbf{B})$. For any basis $\mathbf{B}$, its *fundamental parallelepiped* $\mathcal{P}(\mathbf{B}) = \{\mathbf{B} \cdot \mathbf{x} : \mathbf{x} \in [0, 1)^n\}$. The n-dimensional volume $\mathrm{vol}(\mathcal{P}(\mathbf{B})) = \det \mathbf{B}$ is invariant over any basis $\mathbf{B}$ of Λ; this quantity is called the *fundamental volume* and is denoted by $\det \Lambda$.

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): $\lambda_1^p(\Lambda) = \min_{0 \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\|_p$. The following is a form of Minkowski's first theorem under ℓ_p lengths:

Proposition 3.1. *For any n-dimensional lattice Λ and any $p \in [1, \infty]$, we have*

$$\lambda_1^p(\Lambda) \le \sqrt[p]{n} \cdot (\det \Lambda)^{1/n}.$$

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 $\mathcal{B}_n^p = \{\mathbf{x} \in \mathbb{R}^n \ : \ \|\mathbf{x}\|_p \le 1\}$ 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. Define $g^p(\Lambda)$, which we call the *basis minimum* (in ℓ^p length), to 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.

[2] Usually the name ℓ_p *norm* is used, but as we will see, that term is be claimed by a notion from number theory.

The *dual lattice* of Λ, denoted Λ^*, is defined to be $\Lambda^* = \{\mathbf{x} \in \mathbb{R}^n : \forall\, \mathbf{v} \in \Lambda, \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbb{Z}\}$. Banaszczyk's *transference theorems* give relations between properties of lattices and their duals, in both the standard ℓ_2 length [8] and in general ℓ_p lengths [9]. Following Cai [14] in a straightfoward manner, we can slightly generalize Banaszczyk's results to relate the length of a shortest *basis* for Λ (under any ℓ_p length) to the minimum distance of Λ^* (under the *dual length* of ℓ_p):

Lemma 3.2 (Synthesis of [9] and [14])**.** *There is a constant C such that for any n-dimensional lattice Λ and any $p, q \in [1, \infty]$ with $1/p + 1/q = 1$,*

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

3.3 Gaussian Measures

Our review of Gaussian measures over lattices follows the development of prior works [1, 41, 36]. For any $s > 0$ define the Gaussian function centered at $\mathbf{c}$ with parameter s as:

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

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 $\mathbb{R}^n$ is s^n, therefore we can define a continuous Gaussian probability distribution as $D_{s,\mathbf{c}}(\mathbf{x}) = s^{-n} \cdot \rho_{s,\mathbf{c}}(\mathbf{x})$.

$D_{s,\mathbf{c}}$ is the sum of n orthogonal 1-dimensional Gaussian distributions, which can each be approximated and sampled arbitrarily well using standard algorithms. For simplicity, we will assume that algorithms can efficiently sample from $D_{s,\mathbf{c}}$ exactly.

For any $\mathbf{c} \in \mathbb{R}^n$, real $s > 0$, and lattice Λ, define the *discrete Gaussian distribution over* Λ as:

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

(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 $D_{s,\mathbf{c}}$ and conditioning on $\mathbf{x} \in \Lambda$.

The smoothing parameter. Micciancio and Regev [36] proposed a new lattice quantity which they called the *smoothing parameter*:

Definition 3.3 ([36])**.** For an n-dimensional 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 3.4 ([36])**.** *Let $\mathbf{B}$ be a lattice basis and $\Lambda = \mathcal{L}(\mathbf{B})$. For any $\epsilon > 0$, $\mathbf{c} \in \mathbb{R}^n$, and $s \geq \eta_\epsilon(\Lambda)$,*

$$\Delta(D_{s,\mathbf{c}} \bmod \mathcal{P}(\mathbf{B}), \mathsf{U}(\mathcal{P}(\mathbf{B}))) \leq \epsilon/2.$$

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

Lemma 3.5 ([36])**.** *For any n-dimensional lattice Λ, $\eta_\epsilon(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^*)$ where $\epsilon = 2^{-n}$.*

The smoothing parameter also influences the behavior of the discrete Gaussian distribution $D_{\Lambda,s,\mathbf{c}}$. We will need a few facts about this behavior: the first, shown by Peikert and Rosen [40], bounds the *maximum* value of $D_{\Lambda,s,\mathbf{c}}$ (i.e., the probability of the mode):

Lemma 3.6 ([40]). *Let Λ be a lattice in $\mathbb{R}^n$. For any $\epsilon > 0$, $s \geq 2 \cdot \eta_\epsilon(\Lambda)$, $\mathbf{x} \in \Lambda$, and $\mathbf{c} \in \mathbb{R}^n$,*

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

The next fact, due to Lyubashevsky and Micciancio [31], is a tail inequality on the *coordinates* of a sample $\mathbf{x} \sim D_{\Lambda,s,\mathbf{c}}$ (or more generally, the length of $\mathbf{x}$ when projected onto any unit vector). We will use this lemma in our basic analysis of the worst-case to average-case reduction to obtain a $O(\log^2 n)$ connection factor for any ℓ_p length, $p \in [1, \infty]$.

Lemma 3.7 ([31]). *For any n-dimensional lattice Λ, point $\mathbf{c} \in \mathbb{R}^n$, unit vector $\mathbf{u} \in \mathbb{R}^n$, real $s > 2\eta_\epsilon(\Lambda)$ where $\epsilon < n^{-2\log\log n}$,*

$$\Pr_{\mathbf{x} \sim D_{\Lambda,s,\mathbf{c}}}[|\langle \mathbf{x} - \mathbf{c}, \mathbf{u} \rangle| > s \log n] = \nu(n).$$

In order to obtain the most optimized connection factor, we will perform a parallel analysis of our reduction using a concurrent result of Peikert [39] on *sums* of independent samples from discrete Gaussians:

Lemma 3.8 ([39]). *Let $S \subseteq \mathbb{R}^n$ be a d-dimensional subspace, and for any $\mathbf{x} \in \mathbb{R}^n$, let $\mathbf{x}^S$ denote the projection of $\mathbf{x}$ onto S. Let $m = m(n) = \mathrm{poly}(n)$, let $\epsilon(n) \leq 1/(2m(n)+1)$, and for each $i \in [m]$ let Λ_i be an arbitrary n-dimensional lattice, let $\mathbf{c}_i \in \mathbb{R}^n$ an arbitrary center, and let $s_i \geq \eta_\epsilon(\Lambda_i)$.*

Then for any $p \in [1, \infty)$, there is a constant c_p such that for all sufficiently large n:

$$\mathop{\mathrm{E}}_{\mathbf{x}_i \sim D_{\Lambda_i,s_i,\mathbf{c}_i}} \left[\left\| \sum_{i \in [m]} (\mathbf{x}_i - \mathbf{c}_i)^S \right\|_p \right] \leq c_p \cdot \|\mathbf{s}\|_2 \cdot n^{1/p},$$

where the expectation is taken over the m independent samples $\mathbf{x}_i \sim D_{\Lambda_i,s_i,\mathbf{c}_i}$ for $i \in [m]$.

For $p = \infty$, there is a universal constant c such that for all sufficiently large n:

$$\Pr_{\mathbf{x}_i \sim D_{\Lambda_i,s_i,\mathbf{c}_i}} \left[\left\| \sum_{i \in [m]} (\mathbf{x}_i - \mathbf{c}_i)^S \right\|_p > c \cdot \|\mathbf{s}\|_2 \cdot \sqrt{\log n} \right] \leq 1/4.$$

4 Some Basic Algebraic Number Theory

In the following we review the necessary background in algebraic number theory. Due to lack of space, we will present most facts without proof (which may be found in any number of introductory books on algebraic number theory, e.g. [7, 38].)

An *algebraic number* is any root of some polynomial $p(x) \in \mathbb{Q}[x]$. The *minimal polynomial* of an algebraic number θ is the unique monic, irreducible polynomial $f(x) \in \mathbb{Q}[x]$ of minimal degree such that $f(\theta) = 0$. The *degree* of an algebraic number is the degree of its minimal polynomial. An *algebraic integer* is an algebraic number whose minimal polynomial has *integer* coefficients.

4.1 Number Fields

A *number field* of degree n is a field extension K that is constructed by adjoining a single degree-n algebraic integer to $\mathbb{Q}$. Formally, $K = \mathbb{Q}(\theta)$ for some algebraic integer $\theta \in \mathbb{C}$.

K is an n-dimensional vector space over $\mathbb{Q}$: θ has degree n, so the powers $P = \left\{1, \theta, \ldots, \theta^{n-1}\right\} \subset K$ are linearly independent over $\mathbb{Q}$. In addition, θ^i for $i \geq n$ can be expressed as a linear combination of elements in P using the minimal polynomial of θ, so P forms a basis of K as a vector space over $\mathbb{Q}$. This particular choice of basis is called a *power basis*.

Denote by $\mathcal{O}_K \subset K$ the ring of algebraic integers in K. This ring $\mathcal{O}_K$ is an integral domain, and is a free $\mathbb{Z}$-module of rank n. An *integral basis* for K is any $\mathbb{Z}$-basis $B = \{b_1, \ldots, b_n\} \subset \mathcal{O}_K$ of $\mathcal{O}_K$, i.e. any element of $\mathcal{O}_K$ can be written as an *integer* combination of elements from B. In general, a power basis of K is *not* an integral basis, nor vice-versa.

4.2 Ideals and Factorization

In any ring R (in this work, we will always have $R = \mathcal{O}_K$), an *ideal* $\mathcal{I} \subseteq R$ is a nontrivial (i.e., $\mathcal{I} \neq \{0\}$) set which forms a group under addition and which is closed under multiplication by R, i.e. $xr \in \mathcal{I}$ for all $x \in \mathcal{I}$ and all $r \in R$.[3] The product of two ideals $\mathcal{I}$ and $\mathcal{J}$ is another ideal that is the set of all *finite sums* of terms xy, where $x \in \mathcal{I}$ and $y \in \mathcal{J}$. A *fractional ideal* $\mathcal{I}$ is a generalization of an ideal: all of its elements can be written as fractions with some fixed denominator, i.e. there is some $d \in \mathcal{O}_K$ such that $d\mathcal{I} = \{dx : x \in \mathcal{I}\}$ is an ideal in $\mathcal{O}_K$. The set of fractional ideals is a group under multiplication; that is, fractional ideals can be inverted and multiplied.

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 a prime $q \in \mathbb{Z}$, the principal 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 $e_i = 1$.

4.3 Embeddings

An *embedding* is a ring homomorphism (i.e., one that preserves multiplication, addition, and identity elements). A number field K of degree n has exactly n embeddings (often called *conjugates*) $\{\sigma_j\}_{j \in [n]}$ into $\mathbb{C}$ that fix $\mathbb{Q}$. An embedding whose image lies in $\mathbb{R}$ is called a real embedding; otherwise it is called a complex embedding. For any complex embedding τ, there is also a conjugate embedding $\overline{\tau}$ defined by $\overline{\tau}(x) = \overline{\tau(x)}$ for all x. The number of real embeddings is denoted r_1, and the number of *pairs* of conjugate complex embeddings is denoted r_2, so $n = r_1 + 2r_2$. The pair (r_1, r_2) is called the *signature* of K. By convention, $\{\sigma_j\}_{j \in [r_1]}$ are the real embeddings, and the remaining complex embeddings are paired so that $\sigma_{j+r_1+r_2} = \overline{\sigma_{j+r_1}}$ for $j \in [r_2]$.

We define a *canonical embedding* $\sigma : K \to \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$, given by

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

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 defined component-wise. Due to the r_2 pairs of

[3] The nontriviality condition is non-standard; however, we will have no use for the zero ideal and its inclusion would only encumber the statements of our results.

conjugate embeddings, $\sigma(K)$ spans the n-dimensional subspace

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

Embedding into $\mathbb{R}^n$. Instead of working within the subspace $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ (which would be quite cumbersome and would necessitate re-proving many results in this subspace), it will be much more convenient to associate H with $\mathbb{R}^n$ via a unitary transformation $U \ : \ \mathbb{R}^n \to H$. Here we define a concrete choice of U and establish some important relations between the elements of K and their corresponding embeddings (via σ and U^{-1}) into $\mathbb{R}^n$.

Define an orthonormal basis $\{\mathbf{u}_i\}_{i \in [n]}$ of H according to the following, where unspecified coordinates of $\mathbf{u}_i$ are taken to be zero:

$$\begin{array}{lll} (\mathbf{u}_i)_i = 1 & \text{if} & 1 \le i \le r_1, \\ (\mathbf{u}_i)_i = \frac{1}{\sqrt{2}} = (\mathbf{u}_i)_{i+r_2} & \text{if} & r_1 < i \le r_1 + r_2, \\ (\mathbf{u}_i)_i = \frac{1}{\sqrt{2}}\sqrt{-1} = -(\mathbf{u}_i)_{i-r_2} & \text{if} & r_1 + r_2 < i \le n. \end{array}$$

The reader may verify that the $\mathbf{u}_i$s are orthonormal under the Hermitian inner product and are contained in H. Then letting $\{\mathbf{e}_i\}_{i \in [n]}$ be the standard basis for $\mathbb{R}^n$, we define the unitary linear transformation U so that $U(\mathbf{e}_i) = \mathbf{u}_i$ for all $i \in [n]$.

Now suppose that $x \in K$, and let $\mathbf{x} = U^{-1}(\sigma(x)) \in \mathbb{R}^n$ be the real vector associated with x via the canonical embedding and unitary transformation. By the above, the real part of $\sigma_i(x)$ is

$$\begin{array}{lll} \Re\sigma_i(x) = \sigma_i(x) = \langle \sigma(x), \mathbf{u}_i \rangle = \langle \mathbf{x}, \mathbf{e}_i \rangle = (\mathbf{x})_i & \text{if} & 1 \le i \le r_1, \\ \Re\sigma_i(x) = \frac{1}{2}(\sigma_i(x) + \sigma_{i+r_2}(x)) = \frac{1}{\sqrt{2}} \langle \sigma(x), \mathbf{u}_i \rangle = \frac{1}{\sqrt{2}} \langle \mathbf{x}, \mathbf{e}_i \rangle = \frac{1}{\sqrt{2}}(\mathbf{x})_i & \text{if} & r_1 < i \le r_1 + r_2, \\ \Re\sigma_i(x) = \Re\sigma_{i-r_2}(x) = \frac{1}{\sqrt{2}}(\mathbf{x})_{i-r_2} & \text{if} & r_1 + r_2 < i \le n \end{array}$$

A similar analysis applies to the imaginary part, yielding $\Im\sigma_i(x) = \frac{1}{\sqrt{2}}(\mathbf{x})_{i+r_2}$ for $r_1 < i \le r_1 + r_2$ and $\Im\sigma_i(x) = -\frac{1}{\sqrt{2}}(\mathbf{x})_i$ for $r_1 + r_2 < i \le n$.

4.4 Norm, Lengths, and Discriminant

The *(field) norm* of an element $x \in K$ is the product of all its conjugates: $N(x) = \prod_{i \in [n]} \sigma_i(x)$, which is always an element of $\mathbb{Q}$. If $x \in \mathcal{O}_K$, then $N(x) \in \mathbb{Z}$. The norm is multiplicative: $N(xy) = N(x)N(y)$; this is immediately due to the σ_i being ring homomorphisms.

The notion of norm also generalizes to (fractional) ideals. For any *integral* ideal $\mathcal{I}$ of $\mathcal{O}_K$, the quotient group $\mathcal{O}_K/\mathcal{I}$ is finite. The norm of $\mathcal{I}$, denoted $N(\mathcal{I})$, is defined to be $|\mathcal{O}_K/\mathcal{I}|$. For any $x \in \mathcal{O}_K$, it is the case that $|N(x)| = N(\langle x \rangle)$. For any *fractional* ideal $\mathcal{I}$ over K, we have some $d \in \mathcal{O}_K$ such that $d\mathcal{I}$ is an ideal of $\mathcal{O}_K$; the norm $N(\mathcal{I})$ is defined to be $N(d\mathcal{I})/N(d)$. The norm is therefore multiplicative for fractional ideals as well.

For any $x \in K$ and any $p \in [1, \infty]$, define the ℓ_p length of x to be $\|x\|_p = \|\sigma(x)\|_p$. Note that this definition of ℓ_p length is in relation to the subspace $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ (see Section 4.3). As always, we assume the ℓ_2 length when p is omitted. From these definitions and because the σ_i are ring homomorphisms, we can see that for any $x, y \in K$ and any $p \in [1, \infty]$:

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

The *discriminant* of any n-tuple of elements $x_1, \ldots, x_n \in K$, denoted $\mathrm{disc}(x_1, \ldots, x_n)$, is defined to be $(\det(A))^2$, where A is the $n \times n$ matrix having $A_{i,j} = \sigma_i(x_j)$. The discriminant is always rational, and is an integer if $x_i \in \mathcal{O}_K$ for all i. The discriminant of a number field K, denoted Δ_K, is $\mathrm{disc}(b_1, \ldots, b_n)$ where $\{b_1, \ldots, b_n\}$ is *any* integral basis for K (the discriminant is an invariant over any choice of integral basis). The *root discriminant* of K, denoted $\mathcal{D}_K$, is defined to be $|\Delta_K|^{1/n}$.

4.5 Ideal Lattices

Any (possibly fractional) ideal $\mathcal{I}$ over K is a free $\mathbb{Z}$-module having some basis $\{u_1, \ldots, u_n\}$. Then $\sigma(\mathcal{I})$ is a lattice spanning the subspace $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ having basis $\{\sigma(u_1), \ldots, \sigma(u_n)\}$. Using the unitary transformation U, $U^{-1}(\sigma(\mathcal{I}))$ is a full-rank lattice in $\mathbb{R}^n$ (see Section 4.3). Seen in either H or in $\mathbb{R}^n$, 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 an ideal $\mathcal{I}^*$ over an isomorphic number field $\overline{K} \equiv K$. We will not need the precise form of $\mathcal{I}^*$, but only the fact that $N(\mathcal{I}^*) = (N(\mathcal{I}) \cdot |\Delta_K|)^{-1}$.

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.

4.6 Distributions over Number Fields

In an analog to the definition of (discrete) Gaussian distributions over (lattices in) $\mathbb{R}^n$, we can define (discrete) Gaussian distributions in (fractional ideals over) a number field K. Just as with $\mathbb{R}^n$, the probability of an element x under a Gaussian with parameter $s > 0$ centered at $c \in K$ is proportional to $e^{-\pi\|(x-c)/s\|^2}$. In the continuous case $x \in K$, and in the discrete case $x \in \mathcal{I}$ for some fractional ideal $\mathcal{I}$.[4] We call the continuous probability distribution $D^K_{s,c}$, and it may be sampled in the following way: first, let $\mathbf{c} = U^{-1}(\sigma(c)) \in \mathbb{R}^n$, where U is the unitary transformation discussed in Section 4.3. Then choose an $\mathbf{x} \in \mathbb{R}^n$ according to $D_{s,\mathbf{c}}$, and let $x = \sigma^{-1}(U(\mathbf{x})) \in K$. The corresponding discrete distribution over $\mathcal{I}$ is called $D_{\mathcal{I},s,c}$ (omitting K because it is implicit).

Now suppose $x \in K$ is a random variable with some distribution D^K over K (e.g., a continuous or discrete Gaussian), and let $\mathbf{x}$ be the embedding of x into $\mathbb{R}^n$ with induced distribution D. By the discussion in Section 4.3, for the real embeddings σ_i, $1 \le i \le r_1$, we have $\sigma_i(x) = (\mathbf{x})_i$, i.e. $\sigma_i(x)$ is distributed as the projection of D onto a one-dimensional subspace of $\mathbb{R}^n$. For the complex embeddings σ_i, $r_1 < i \le r_1 + r_2$, we have $\sqrt{2} \cdot \sigma_i(x) = \sqrt{2}(\Re\sigma_i(x), \Im\sigma_i(x)) = ((\mathbf{x})_i, (\mathbf{x})_{i+r_2})$, i.e. $\sqrt{2}\sigma_i(x)$ is distributed as the projection of D onto a *two-dimensional* subspace of $\mathbb{R}^n$ (and likewise for $r_1 + r_2 < i \le n$).

4.7 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 computational operations in polynomial time. Our exposition mainly follows [26]; a detailed treatment of these issues can be found in [17, Sections 4.2–4.7] and [25, Section 2].

We say that a number field K and its ring of integers $\mathcal{O}_K$ are *explicitly given* if, for some *integral basis* for K, we have the integer matrices implementing multiplication by each basis element (relative

[4] Formally, K is not a continuous space, and therefore cannot support a continuous probability distribution; this can be overcome by standard mathematical techniques. In order to avoid excessive formalism, and because algorithms can only approximate Gaussian distributions anyway, we will remain content with this slight abuse.

to that basis). That is, K is given by an integral basis $B = \{b_1, \ldots, b_n\}$ for K, where each b_i is represented by a nondegenerate matrix $\mathbf{B}_i \in \mathbb{Z}^{n \times n}$. For any $y \in K$ represented as a rational vector $\mathbf{y} \in \mathbb{Q}^n$ relative to B, the prodcut $\mathbf{B}_i\mathbf{y}$ is the representation of $b_i y$ relative to B. Because B is a basis for K as a linear vector space, the matrices $\mathbf{B}_i$ also fully specify multiplication by any element $x \in K$. For measuring computational complexity, "polynomial" is taken to mean some polynomial in both n and $\log |\Delta_K|$. With this convention, each $\mathbf{B}_i$ can be represented using a polynomial number of bits. Addition, multiplication, and division within K, as well as the embeddings from K into $\mathbb{C}$ and their inverses, can all be performed in polynomial time.

An (integral) ideal $\mathcal{I} \subseteq \mathcal{O}_K$ has a $\mathbb{Z}$-basis $\{u_1, \ldots, u_n\} \subset \mathcal{O}_K$ and is given by the vector representations of each u_i relative to the integral basis B. Representing a *fractional* ideal $\mathcal{I}$ requires also specifying an element $d \in \mathcal{O}_K$ (also relative to B) for which $d\mathcal{I}$ is an integral ideal. The representations of the u_i generating the ideal can be kept in *Hermite Normal Form* (HNF), which makes the representation of the ideal unique (thus allowing efficient equality tests). It is possible to multiply an ideal by an element of K, to multiply two 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 to enumerate $\mathcal{I}/\mathcal{I}'$ in time polynomial in $|\mathcal{I}/\mathcal{I}'|$, $\log |\Delta_K|$, and n.

5 Properties of Ideal Lattices

In this section we develop several useful facts about ideal lattices. It is likely that some (if not many) of these facts are already known, but they play such an important role in our work that we prefer to present them and their proofs in full. Throughout the section, K denotes any number field of degree n.

5.1 Minima

Here we develop several useful facts about, and connections among, the various minima (successive minima, basis minimum) of ideal lattices.

Lemma 5.1. *For any fractional ideal $\mathcal{I}$ over K and any $p \in [1, \infty]$, $\lambda_1^p(\mathcal{I}) \leq \sqrt[p]{n} \cdot N^{1/n}(\mathcal{I}) \cdot \sqrt{\mathcal{D}_K}$.*

Proof. Follows immediately by Proposition 3.1 and the fact that $\det \mathcal{I} = N(\mathcal{I})\sqrt{|\Delta_K|}$. □

Our next lemma and its implications provide one of two *crucial foundations* upon which our improved worst-case to average-case connection factor rests. In particular, it leads directly to our improved bound on the smoothing parameter (Lemma 5.6), and also to an essential tool (Lemma 5.4) that will allow us bound the length of a generating set for a principal ideal.

Lemma 5.2 (First Foundation)**.** *For any $x \in K$ and any $p \in [1, \infty]$, $\|x\|_p \geq \sqrt[p]{n} \cdot |N(x)|^{1/n}$.*

Proof. For $1 \leq p < \infty$, by the AM-GM inequality we get:

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

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

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

Corollary 5.3. *For any fractional ideal $\mathcal{I}$ over K and any $p \in [1, \infty]$, $\lambda_1^p(\mathcal{I}) \geq \sqrt[p]{n} \cdot N^{1/n}(\mathcal{I})$.*

Proof. For $x \in \mathcal{I}$, $N(\mathcal{I})$ divides $|N(x)|$, so for nonzero $x \in \mathcal{I}$, $|N(x)| \geq N(\mathcal{I})$. □

Recall that the basis minimum $g^p(\mathcal{I})$ is the minimal length of a basis (in ℓ_p length) for $\mathcal{I}$. Using the First Foundation Lemma, we can bound the value of $g^p(\mathcal{I})$:

Lemma 5.4. *There is a constant C such that for any fractional ideal $\mathcal{I}$ over K and any $p \in [1, \infty]$,*

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

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

Proof. Let $q = 1 - 1/p$. By Lemma 3.2, there is some C such that $g^p(\mathcal{I}) \cdot \lambda_1^q(\mathcal{I}^*) \leq C \cdot n\sqrt{\log n}$. Because $\mathcal{I}^*$ is a fractional ideal over a number field of degree n having root discriminant $\mathcal{D}_K$, Corollary 5.3 applies, yielding

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

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

Though we will not need this for any of our main results, from the basis minimum we can get a connection between the successive minima of any ideal lattice:

Lemma 5.5. *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. Consider an integral basis $B = \{b_1, \ldots, b_n\}$ 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 $\lambda_n^p(\mathcal{I}) \leq \|X\|_p \leq \|B\|_\infty \cdot \|x\|_p = g^\infty(\mathcal{O}_K) \cdot \lambda_1^p(\mathcal{I})$. □

5.2 Smoothing Parameter

Here we present a bound on the smoothing parameter for ideal lattices. While the proof is straightforward given our tools from above, for number fields with small root discriminant the bound is much stronger than a prior bound which related the smoothing parameter to λ_n [36].

Lemma 5.6. *For any fractional ideal $\mathcal{I}$ over K, $\eta_\epsilon(\mathcal{I}) \leq \mathcal{D}_K \cdot N^{1/n}(\mathcal{I})$, where $\epsilon = 2^{-n}$.*

Proof. We have

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

where the first inequality follows from Lemma 3.5 and the second is from Corollary 5.3. □

Let us take a moment to examine the implications of this bound, for simplicity restricting the discussion to the ℓ_2 length. Let $\mathcal{I}$ be a fractional ideal over K, having $N(\mathcal{I}) = 1$ without loss of generality. Then even for exponentially small $\epsilon = 2^{-n}$, $\eta_\epsilon(\mathcal{I}) \leq \mathcal{D}_K$ by the above lemma. On the other hand, $\lambda_1(\mathcal{I}) \geq \sqrt{n}$ by Corollary 5.3. Therefore η_ϵ is at most $\frac{\mathcal{D}_K}{\sqrt{n}} \cdot \lambda_1$, which is $O(\frac{1}{\sqrt{n}}) \cdot \lambda_1$ if

$\mathcal{D}_K$ is bounded by a constant as n grows. In constrast, a bound from [36] implies that η_ϵ is at most $f(n) \cdot \lambda_n$, where $f(n)$ grows very slowly with n, and $\epsilon(n)$ is a sufficiently large (but still negligible) function of n. Therefore Lemma 5.6 can yield at least a $\sqrt{n}$ factor or more improvement in the smoothing parameter.

We also remark that the above proof is essentially oblivious to the particular geometry of the lattice. The proof only depends on the norm of the ideal and the discriminant of the number field. We do not know if there is a stronger bound that uses more information about the lattice, even for a negligible $\epsilon(n)$ larger than 2^{-n}.

6 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. We also demonstrate several (worst-case to worst-case) reductions between these problems.

6.1 Preprocessing Number Fields

All of the problems we define are parameterized by a fixed choice of number field K (or, in their asymptotic versions, a 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. [13, 33, 20]).

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.

6.2 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 6.1 (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 6.2 (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 also can define decisional versions as promise ("gap") problems:

Definition 6.3 (Gap $\mathsf{IGVP}/\mathsf{ISVP}$)**.** An input to K-$\mathsf{GapIGVP}_\gamma^{\phi}$ 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 obtained by setting $\phi = \lambda_1^p$.

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

Definition 6.4 (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 6.5 (Gap ICVP)**.** An input to K-$\mathsf{GapICVP}_\gamma^{p}$ 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 the asymptotic hardness of these problems as a function of the degree n of the number fields, 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.[5] This is analogous to the formulation of computational problems and algorithms 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' for the same family $\mathcal{K}$, 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'.

6.3 Reductions Among Worst-Case Problems

In this section we provide some (worst-case to worst-case) reductions among problems on ideal lattices. For all of our results, analogous reductions are known to exist for *general* lattices [23], but those reductions are not valid for *ideal* lattice problems because (in general) they invoke their oracles on non-ideal lattices. Nevertheless, the reductions we construct here are in fact inspired by, and use similar techniques to, the reductions for general lattice problems.

The essential technique from [23] 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 small, and (2) the intersection of all Λ_is does not 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

[5] In fact, for any number field K of *fixed* degree, there is some constant approximation factor γ (which depends on the degree of K) for which all of the above problems are efficiently solvable in time polynomial in the instance size using the LLL algorithm and its variants [29, 45]. However, the factor γ is essentially exponential in the degree.

$|\Lambda/\Lambda_i| = 2$, 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 slightly 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 small, and (2) the intersection of all $\mathcal{I}_{i,j}$ does not contain a shortest vector in $\mathcal{I}$.

Our reductions will rely on the existence of an integer prime $q \in \mathbb{Z}$ for which all the prime ideal divisors of the principal ideal $\langle q \rangle$ have "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 perhaps 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 [24] demonstrated that there exist infinite families of number fields, all having the same (constant) root discriminant, for which some q splits completely in every member of the family.

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

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

Reducing ISVP to ICVP. 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. The efficiency of the reduction depends on the splitting of the prime q.

Proposition 6.6. *Let K be a number field for which prime $q \in \mathbb{Z}$ factors as $\langle q \rangle = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_L^{e_L}$. For any γ and any $p \in [1, \infty]$, there is a deterministic non-adaptive Cook reduction from K-ISVP_γ^p (resp., K-$\mathsf{GapISVP}_\gamma^p$) to K-ICVP_γ^p (resp., K-$\mathsf{GapICVP}_\gamma^p$). The reduction makes $\sum_{i \in [L]} e_i \cdot (N(\mathfrak{q}_i) - 1)$ queries to its oracle.*

Proof. We provide a reduction between the search problems, which can be easily adapted for the decisional 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_γ^p 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 $\mathcal{I}_{i,j-1}/\mathcal{I}_{i,j}$. 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_γ^p. □

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 6.7. *Let K be a number field for which prime $q \in \mathbb{Z}$ factors as $\langle q \rangle = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_L^{e_L}$. For any γ and any $p \in [1, \infty]$, there is a deterministic non-adaptive Cook reduction from K-ICVP_1^p to K-$\mathsf{GapICVP}_1^p$. The number of queries is polynomial in the input size times $\sum_{i \in [L]} e_i \cdot N(\mathfrak{q}_i)$.*

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. Our 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 proof of Proposition 6.6.
3. For each $i \in [L]$, let $v_i = 0$ and $t' = t$. 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 $\mathcal{A}(\mathcal{I}_{i,j}, t' - x, d) = \mathrm{YES}$.
   (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 [23], 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. □

6.4 Average-Case Problem

For a number field K, a positive integer $q \in \mathbb{Z}$, and $A = (a_1, \ldots, a_m) \in (\mathcal{O}_K/\langle q \rangle)^m$, define the set

$$\Psi(q, A) = \left\{ Z = (z_1, \ldots, z_m) \in \mathcal{O}_K^m \ : \ \sum_{i \in [m]} a_i z_i \in \langle q \rangle \right\}.$$

We remark that the set $\Psi(q, A)$ has a lattice-like structure: it is closed under (coordinate-wise) addition and multiplication by any element in $\mathcal{O}_K$. We now define our average-case problem, whose goal is to find a nontrivial $Z = (z_1, \ldots, z_m) \in \Psi(q, A)$ whose entries z_i all have bounded ℓ_∞ length.

Definition 6.8 (Short Algebraic Integer Solution)**.** For a number field K, an input to the K-SAIS problem is a tuple (q, A, β) where $q \in \mathbb{Z}$, $A \in (\mathcal{O}_K/\langle q \rangle)^m$, and $\beta \in \mathbb{R}$. The goal is to find a nonzero $Z \in \Psi(q, A)$ such that $\|Z\|_\infty \leq 2\beta$.

For a family $\mathcal{K} = \{K_n\}$ of number fields and functions $q(n)$, $m(n)$, $\beta(n)$, define $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}$ to be the probability ensemble over instances $(q(n), A, \beta(n))$ of K_n-SAIS where A is chosen uniformly from $(\mathcal{O}_{K_n}/\langle q(n) \rangle)^{m(n)}$.

We will of course need to choose parameters q, m, and β so that K-SAIS instances admit a solution (otherwise the problem is trivially hard). This entails choosing a large enough β relative to m and q. The proper dependence also turns out to be governed by the discriminant of the number field: with a smaller discriminant, the points in $\Psi(q, A)$ are more "densely-packed," so a smaller bound β can suffice to guarantee a non-trivial solution. The following lemma makes this precise.

Lemma 6.9 (Second Foundation)**.** *Let K be a number field having signature (r_1, r_2) and let $m, q > 0$ be integers. For any $A \in (\mathcal{O}_K/\langle q \rangle)^m$ and any $\beta > 0$ such that*

$$\beta^n > \frac{\sqrt{|\Delta_K|}}{2^{r_1}\pi^{r_2}} \cdot q^{n/m},$$

there exists nonzero $Z \in \Psi(q, A)$ such that $\|Z\|_\infty \leq 2\beta$.

Proof. We employ a geometric argument that was used by Lenstra [30] and Guruswami [24] to construct error-correcting codes over number fields, whose essential idea goes back to Blichfeldt.

Recall the definition of the subspace $H \subseteq \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ from Section 4.3. Consider the closed "cube" $\mathcal{C} = \{\mathbf{x} \in H : \|\mathbf{x}\|_\infty \leq \beta\}$. The n-dimensional volume of $\mathcal{C}$, denoted $\mathrm{vol}(\mathcal{C})$, is computed (by integrating over each coordinate x_i, $i \in [r_1 + r_2]$ independently) to be $(2\beta)^{r_1}(\pi\beta^2)^{r_2} = 2^{r_1}\pi^{r_2}\beta^n$.

Now consider the lattice $\sigma(\mathcal{O}_K) \subset H$, whose n-dimensional fundamental volume is $\det \mathcal{O}_K = \sqrt{|\Delta_K|}$. Intuitively, by the ratio of $\mathrm{vol}(\mathcal{C})$ to $\det \mathcal{O}_K$, we would expect to have about

$$Q = \frac{2^{r_1}\pi^{r_2}\beta^n}{\sqrt{|\Delta_K|}}$$

points in $\sigma(\mathcal{O}_K) \cap \mathcal{C}$. An averaging argument [30, 24] can make this intuition rigorous: by shifting the cube $\mathcal{C}$, we can guarantee that Q lattice points lie in the shifted cube, i.e. for some $\mathbf{y} \in H$,

$$|(\mathcal{C} - \mathbf{y}) \cap \sigma(\mathcal{O}_K)| \geq Q.$$

Because $Q^m > q^n = |\mathcal{O}_K/\langle q \rangle|$, then by the pigeonhole principle there are distinct $X = (x_1, \ldots, x_m)$, $X' = (x'_1, \ldots, x'_m) \in (\mathcal{O}_K)^m$ such that $\sum a_i x_i = \sum a_i x'_i \bmod \langle q \rangle$ and $\sigma(x_i), \sigma(x'_i) \in (\mathcal{C} - \mathbf{y})$ for all $i \in [m]$. Letting $Z = X - X' \in 2\mathcal{C}$, we then have $Z \in \Psi(q, A)$ and $\|Z\|_\infty \leq 2\beta$, as desired. □

7 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 first state the result that emerges from a "basic" analysis of the reduction (this is Theorem 7.1). Immediately following that, we state the parameters that emerge from a tighter analysis which uses the techniques of [39] (this is Theorem 7.6). We complete the section by connecting ISVP to IGVP for concrete choices of the family $\mathcal{K}$ of number fields, obtaining our (sub-)logarithmic connection factors.

Theorem 7.1 (Main Reduction, Basic Analysis)**.** *For any infinite family of number fields $\mathcal{K} = \{K_n\}$, $p \in [1, \infty]$, $\epsilon(n) < n^{-2\log\log n}$, and $m(n)$, $q(n)$, $\beta(n)$, $\gamma(n)$ that satisfy the conditions below, there is a polynomial-time number field-preserving reduction from solving $\mathcal{K}$-$\mathsf{IncIGVP}^{p,\eta_\epsilon}_{\gamma}$ (or, equivalently, $\mathcal{K}$-$\mathsf{IGVP}^{p,\eta_\epsilon}_{\gamma}$) in the worst case to solving $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}$ on the average with non-negligible probability.*

The conditions on the parameters are as follows:

1. $\gamma(n) \geq 16 \cdot \beta(n) \cdot m(n) \cdot \sqrt[p]{n} \log n$,
2. $q(n) \geq 4n \cdot \beta(n) \cdot m(n) \cdot g^\infty(\mathcal{O}_{K_n})$, *and*
3. $m(n)$, $q(n)$, $\beta(n)$ *are all* $\mathrm{poly}(n)$ *and satisfy the condition in Lemma 6.9 for every* 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.[6] By Lemma 5.4, 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.

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 unique 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 (in polynomial time) a $\lfloor w \rceil_B \in \mathcal{O}_K$ 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 $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}$ with non-negligible probability. We construct an algorithm to solve $\mathcal{K}$-$\mathsf{IncIGVP}^{p,\eta_\epsilon}_{\gamma}$ 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 $A = (a_1, \ldots, a_m)$ and let $Z = (z_1, \ldots, z_m) \leftarrow \mathcal{F}(A)$. Output

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

[6] Actually, it suffices for B to have length $\|B\|_\infty = \mathrm{poly}(n)$, if we require $q(n)$ to be sufficiently large.

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 7.2. *The probability that $Z \in \Psi(q, A)$ is non-negligible in n.*

Proof. It suffices to bound the statistical distance $\Delta(A, \mathsf{U}^m(\mathcal{O}_K/\langle q \rangle))$ 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 3.3, $\Delta(y'_j, K/\mathcal{I}) \leq \epsilon/2$ (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 7.3. *If $Z \in \Psi(q, A)$, 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 $Z \in \Psi(q, A)$, 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 7.4. *Conditioned on $Z \in \Psi(q, A)$, $\|x'\|_p \leq \frac{\|x\|_p}{2}$ with overwhelming probability.*

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

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

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

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

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

$$\|x\|_p \cdot \frac{\beta mn \|B\|_\infty}{q}.$$

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

We now bound the second summation of Inequality (3). First, for all $j \in [m]$,

$$\|y_j z_j\|_p \le \|y_j\|_p \cdot \|z_j\|_\infty \le 2\beta \|y_j\|_p.$$

By a now-standard argument [36, 34, 40, 31], 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 A and Z. We now establish a tail inequality of $\|y_j\|_p \le s\sqrt[p]{n}\log n$ for all $j \in [m]$ (with overwhelming probability), conditioned on *any* fixed values of y'_j; the unconditioned inequality follows by averaging. To do so, it suffices to show that $|\sigma_i(y_j)| \le s\log n$ for all $i \in [n]$, $j \in [m]$ with overwhelming probability. There are two cases: if σ_i is a real embedding (i.e., $1 \le i \le r_1$), then by the discussion in Sections 4.3 and 4.6,

$$\sigma_i(y_j) = \sigma_i((y_j - y'_j) - (-y'_j)) = \left\langle D_{\Lambda,s,-\mathbf{y}'_j} - (-\mathbf{y}'_j), \mathbf{e}_i \right\rangle,$$

where $\Lambda = U^{-1}(\sigma(\mathcal{I})) \subset \mathbb{R}^n$ is the real lattice associated with $\mathcal{I}$, $\mathbf{y}'_j = U^{-1}(\sigma(y'_j)) \in \mathbb{R}^n$ is the center associated with y'_j, and $\mathbf{e}_i$ is the ith standard basis element of $\mathbb{R}^n$. Then by Lemma 3.7, $|\sigma_i(y_j)| \le s\log n$ with probability $1 - \nu(n)$. In the second case, σ_i is a complex embedding (i.e., $r_1 < i \le n$), and by the discussion in Sections 4.3 and 4.6, both $\Re\sigma_i(y_j)$ and $\Im\sigma_i(y_j)$ are distributed as $\frac{1}{\sqrt{2}}\left\langle D_{\Lambda,s,-\mathbf{y}'_j} - (-\mathbf{y}'_j), \pm\mathbf{e}_k \right\rangle$ for some appropriate basis vector $\mathbf{e}_k$. By Lemma 3.7 and $|z|^2 = (\Re z)^2 + (\Im z)^2$ for $z \in \mathbb{C}$, we again have $|\sigma_i(y_j)| \le s\log n$ with probability $1 - \nu(n)$. The desired claim holds for all i, j by the union bound.

We conclude that the second summation of Equation (3) is at most

$$2\beta ms\sqrt[p]{n}\log n = \|x\|_p \cdot \frac{4\beta m\sqrt[p]{n}\log n}{\gamma}.$$

By assumption, $\gamma \ge 16\beta m\sqrt[p]{n}\log n$, so the second summation is at most $\|x\|_p/4$, as desired. □

Claim 7.5. *Conditioned on $Z \in \Psi(q, A)$, $x' \ne 0$ with overwhelming probability.*

Proof. The main idea is that $x' = 0$ if and only if a sample from $D_{\mathcal{I},s,\mathbf{c}}$ hits a *single, particular* "bad" value. Lemma 3.6 guarantees that the probability of this event is negligibly small.

By definition of w_j,

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

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 $Z \neq 0$, 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:

$$y_1 - y'_1 = \left(t_1 + v_1 - \frac{x\lfloor w_j \rceil}{q}\right) + z_1^{-1} \sum_{i=2}^{m} \left(t_j + v_j + y'_j - y_j - \frac{x\lfloor w_j \rceil}{q}\right) \cdot z_j. \quad (4)$$

As in the proof of Claim 7.4, 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 3.6 guarantees that the equation is only satisfied with probability $2^{-n} \cdot \frac{1+\epsilon}{1-\epsilon} = \nu(n)$. □

By Claims 7.2 through 7.5 and the union bound over the negligible failure probabilities, 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 7.1 follows. □

7.1 A Tighter Analysis

Theorem 7.6 (Main Reduction, Tighter Analysis)**.** *The statement of Theorem 7.1 also holds for any negligible $\epsilon(n) = \nu(n)$, and for some $\gamma(n) = O(\beta(n) \cdot \sqrt{m(n)} \cdot \sqrt[p]{n})$ if $p \in [1, \infty)$, or for some $\gamma(n) = O(\beta(n) \cdot \sqrt{m(n) \log n})$ if $p = \infty$. The constants hidden by the $O(\cdot)$ depend only on p.*

Proof sketch. To prove the theorem, it is enough to re-establish Claim 7.4, which is the only claim that depends on $\gamma(n)$. Instead of obtaining a high-probability bound on $\|x'\|_p$, it is enough to prove that the expectation $\mathrm{E}\left[\|x'\|_p\right] \leq \frac{\|x\|_p}{4}$. Then by Markov's inequality we have $\Pr\left[\|x'\|_p > \frac{\|x\|_p}{2}\right] \leq 1/2$. All the other failure probabilities in the other claims are negligible, so the reduction succeeds with non-negligible probability by a union bound.

The analysis of $\mathrm{E}\left[\|x'\|_p\right]$ is an immediate application of Lemma 3.8. We defer the details. □

7.2 Connection to ISVP

We now give a reduction from ISVP, instantiating all the parameters from Theorem 7.6 asymptotically, and focusing especially on the role of the root discriminant.

Theorem 7.7. *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 $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}$ on the average with non-negligible probability is at least as hard as solving $\mathcal{K}$-ISVP^p_γ in the worst case. For $p = \infty$, there exists $\gamma(n) = O(\log n) \cdot \mathcal{D}_{K_n}^{1.5}$ for which the same claim applies.

Proof. Assume that $p \in [1, \infty)$, and let $\epsilon = 2^{-n}$. By Lemma 5.4, 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 7.6, 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 7.6, $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}$ is as hard as $\mathcal{K}$-$\mathsf{IGVP}^{p,\eta_\epsilon}_{\gamma'}$ for some $\gamma'(n) = O(\sqrt[p]{n}\sqrt{\log n})\cdot\sqrt{\mathcal{D}_{K_n}}$.

We now connect $\mathcal{K}$-IGVP to $\mathcal{K}$-ISVP. By Corollary 5.3 and Lemma 5.6,

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

for any ideal $\mathcal{I}\subseteq\mathcal{O}_{K_n}$. Therefore $\mathcal{K}$-$\mathsf{IGVP}^{p,\eta_\epsilon}_{\gamma'}$ is as hard as $\mathcal{K}$-ISVP^p_γ for some $\gamma(n) = O(\sqrt{\log n})\cdot\mathcal{D}^{1.5}_{K_n}$.

For $p=\infty$, a similar analysis applies. □

Corollary 7.8. *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 $\mathcal{K}$-$\mathsf{SAIS}_{q,m,\beta}$ on the average with non-negligible probability is at least as hard as solving $\mathcal{K}$-ISVP^p_γ in the worst case. For $p=\infty$, there exists $\gamma(n)=O(\log n)$ for which the same claim applies.

Proof. Follows from Theorem 7.7 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. [44]). □

8 Acknowledgements

We gratefully acknowledge Eva Bayer-Fluckiger, Dan Boneh, Henri Cohen, Noam Elkies, Alex Healy, and Denis Simon for their help regarding algebraic number theory and ideal lattices. We also thank Vadim Lyubashevsky and Daniele Micciancio for helpful discussions.

References

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

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

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

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

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

[6] A. Akavia, O. Goldreich, S. Goldwasser, and D. Moshkovitz. On basing one-way functions on NP-hardness. In J. M. Kleinberg, editor, *STOC*, pages 701–710. ACM, 2006.

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

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

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

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

[11] 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.

[12] G. Brassard. Relativized cryptography. In *FOCS*, pages 383–391. IEEE, 1979.

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

[14] 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.

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

[16] 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.

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

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

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

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

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

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

[23] 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.

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

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

[26] S. Hallgren. Fast quantum algorithms for computing the unit group and class group of a number field. In H. N. Gabow and R. Fagin, editors, *STOC*, pages 468–474. ACM, 2005.

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

[28] R. Kumar and D. Sivakumar. On polynomial-factor approximations to the shortest lattice vector length. *SIAM J. Discrete Math.*, 16(3):422–425, 2003.

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

[30] 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.

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

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

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

[34] 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.

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

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

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

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

[39] C. Peikert. Title to be determined. In submission., 2006.

[40] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In S. Halevi and T. Rabin, editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006.

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

[42] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In H. N. Gabow and R. Fagin, editors, *STOC*, pages 84–93. ACM, 2005.

[43] O. Regev and R. Rosen. Lattice problems and norm embeddings. In J. M. Kleinberg, editor, *STOC*, pages 447–456. ACM, 2006.

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

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

[46] D. Simon. Personal communication.

[47] D. Simon. Construction de polynômes de petits discriminants. *Comptes Rendus de l'Académie des Sciences - Série I - Mathematics*, pages 465–468, 1999.