Electronic Colloquium on Computational Complexity, Report No. 148 (2006)

Limits on the Hardness of Lattice Problems in ℓ_p Norms

Chris Peikert*

December 3, 2006

Abstract

We show that for any $p \geq 2$, lattice problems in the ℓ_p norm are subject to all the same limits on hardness as are known for the ℓ_2 norm. In particular, for lattices of dimension n:

- Approximating the shortest and closest vector in the ℓ_p norm to within $\tilde{O}(\sqrt{n})$ factors is contained in coNP.
- Approximating the length of the shortest vector in the ℓ_p norm to within $\tilde{O}(n)$ factors reduces to the same average-case problems that have been studied in related works (Ajtai, STOC 1996; Micciancio and Regev, FOCS 2004; Regev, STOC 2005).

Each of these results improves upon the current understanding of ℓ_p norms by up to a $\sqrt{n}$ factor. Taken together, they can be seen as a partial converse to recent reductions from lattice problems in the ℓ_2 norm to corresponding problems in ℓ_p norms (Regev and Rosen, STOC 2006).

One of our main technical contributions is a general analysis of *sums* of independent Gaussian distributions over lattices, which may be of independent interest. Our proofs employ analytical techniques of Banaszczyk which, to our knowledge, have yet to be exploited in computer science.

1 Introduction

The last two decades have seen an explosion of interest in the complexity of computational problems on *lattices*. A lattice is a periodic "grid" of points in $\mathbb{R}^n$ generated by all integer combinations of some *basis* of linearly independent vectors. The two central problems on lattices are the *shortest vector* problem SVP and the *closest vector* problem CVP. In SVP, the goal is to find a (nonzero) lattice point which is closest to the origin, given a basis for the lattice. In CVP, the goal is to find a lattice point which is closest to some given target point in $\mathbb{R}^n$. In these problems, it is common to measure distances in the ℓ_p norm for some $1 \leq p \leq \infty$, usually $p = 2$.[1]

As with many optimization problems, it is interesting to consider *approximation* versions of SVP and CVP (and others). For SVP (resp., CVP), the goal then becomes to find a lattice point whose distance from the origin (resp., target) exceeds the optimal distance by no more than a certain approximation factor $\gamma \geq 1$. Known polynomial-time algorithms for SVP and CVP, such as the celebrated LLL algorithm, achieve approximation factors essentially exponential in the dimension n [20, 7, 29, 5]. The best algorithm for solving the *exact* version of SVP runs in randomized $2^{O(n)}$ time [5].

*SRI International, cpeikert@alum.mit.edu

[1] For $p < \infty$, the ℓ_p norm of $\mathbf{x} \in \mathbb{R}^n$ is $\|\mathbf{x}\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$, and the ℓ_∞ norm is $\|\mathbf{x}\|_\infty = \max_i |x_i|$.

ISSN 1433-8092

Hardness. There are several hardness results that reinforce the apparent difficulty of lattice problems for small approximation factors. For the sake of brevity, we describe the highlights.

SVP in any ℓ_p norm is hard to approximate to within any constant factor (and more, for $p = \infty$), assuming $\mathsf{NP} \not\subseteq \mathsf{RP}$ [19, 12, 28]. CVP in any ℓ_p norm is NP-hard to approximate to within almost-polynomial $n^{O(1/\log\log n)}$ factors [13, 12]. Other problems such as the *closest vector with preprocessing* problem CVPP and the *shortest independent vectors* problem SIVP in any ℓ_p norm are hard for any constant factor as well, assuming $\mathsf{NP} \not\subseteq \mathsf{RP}$ [6, 11, 28]. (Several of these constant factors can be improved to almost-polynomial in n, assuming $\mathsf{NP} \not\subseteq \mathsf{RTIME}(2^{\mathrm{poly}(\log n)})$.)

Of special interest is a recent result of Regev and Rosen [28]. Using an elegant application of norm embeddings, they showed (essentially) that lattice problems are *easiest* in the ℓ_2 norm for any given approximation factor.

Limits on hardness. Given the difficulty of designing efficient approximation algorithms for even moderately sub-exponential factors, one might hope to significantly improve upon the known hardness. However, there seem to be strict limits on any such improvements. We explain below.

For certain approximation factors $\gamma(n)$ as small as $O(\sqrt{n})$, it turns out that many interesting lattice problems in the ℓ_2 norm (including SVP and CVP) are in coAM or even coNP. This implies that these approximation problems are not NP-hard, unless the polynomial hierarchy collapses. Results of this type can be found in [16, 2, 18].

One of the most remarkable features of lattice problems are the worst-case to average-case reductions first demonstrated by Ajtai [3], which have seen many extension and improvements in recent years (see, e.g., [4, 24, 27]). Such a reduction is usually taken as positive evidence for the *hardness* of the average-case problem. At the same time, though, the reduction also *limits the hardness* of the worst-case problem, by showing that it is *as easy* as the average-case problem (which is in, say, distributional-NP [21, 10]). The state of the art for worst-case/average-case connections is represented by the works of Micciancio and Regev [24] and Regev [27], who obtained reductions from several lattice problems in the ℓ_2 norm for almost-linear $\tilde{O}(n)$ approximation factors.

All the results limiting the hardness of lattice problems are focused primarily on the ℓ_2 norm. Using standard relations between norms, one can obtain limits for other ℓ_p norms, but the approximation factors suffer. For example, the factors become $O(n^{1/2+|1/2-1/p|}) = O(n)$ for the containments in coNP, and $\tilde{O}(n^{1+|1/2-1/p|}) = \tilde{O}(n^{1.5})$ for the worst-case/average-case reductions.

Summary. In terms of ℓ_p norms, the landscape looks as follows: by Regev and Rosen [28], we know that for a given approximation factor, as p increases from 2 to ∞ (or decreases from 2 to 1), lattice problems become *no easier*. In addition, the known limits on their hardness become *weaker*. But do the problems actually get *strictly harder*? This is the main question motivating our work.

1.1 Our Results

Stated informally, our results are that for any $p \geq 2$, lattice problems in the ℓ_p norm are subject to *all the same limits* on hardness as are known for the ℓ_2 norm, for essentially the same asymptotic approximation factors. Specifically, for any $2 \leq p \leq \infty$ we show that:

- For certain $\tilde{O}(\sqrt{n})$ approximation factors, the following problems in the ℓ_p norm are contained in coNP: CVP, SVP, SIVP, and CRP.

- For certain $\tilde{O}(\sqrt{n})$ approximation factors, decisional CVPP in the ℓ_p norm is *easy* (i.e., in P).
- For certain $\tilde{O}(n)$ connection factors, the following worst-case problems in the ℓ_p norm reduce to the average-case problems from [24, 27]: SIVP, decisional SVP, and others.

Each of these results improves upon the current best approximation factors for the ℓ_p norm by up to a $\sqrt{n}$ factor, and essentially matches the current state of the art for the ℓ_2 norm.

On the technical side, one of our main contributions is a very general analysis of *sums* of Gaussian distributions over lattices. This may be of independent interest and utility elsewhere. Indeed, this analysis has also been applied in a concurrent work by Peikert and Rosen [25] on worst-case/average-case reductions for special classes of algebraic lattices.

Our results also have cryptographic implications. Until now, the hardness of average-case problems has always been based on worst-case problems in the ℓ_2 norm. Because lattice problems are in fact *easiest* in the ℓ_2 norm (for similar approximation factors), the security of the resulting cryptographic primitives was therefore based on the *strongest* worst-case assumption of its kind. Our results imply that security can be based on the possibly weaker assumption that lattice problems are hard in *some* ℓ_p norm, $2 \leq p \leq \infty$.

We remark that the factors hidden by the $\tilde{O}$-notation above do depend mildly on the choice of norm. In all of the quantities, there is a constant factor proportional to $\sqrt{p}$ for $p < \infty$, and a factor proportional to $\sqrt{\log n}$ for $p = \infty$. These are in addition to any small logarithmic factors which may already exist in the known results for the ℓ_2 norm.

1.2 Techniques

One way of obtaining all our results (and more) would be to give approximation-preserving reductions from lattice problems in the ℓ_p norm to problems in the ℓ_2 norm. While reductions in the reverse direction are known [28], reducing from the ℓ_p norm to the ℓ_2 norm appears to be much more challenging.

We will instead obtain our results by *directly demonstrating* the requisite coNP proof systems, worst-case to average-case reductions, etc. Remarkably, we are able to use, *without modification*, the exact same constructions that were designed for the ℓ_2 norm [2, 24, 27]! Our main technical contribution is a *novel analysis* of these algorithms for ℓ_p norms. We rely crucially on harmonic analysis techniques of Banaszczyk that were initially developed to prove transference theorems for lattices, first for the ℓ_2 norm [8], and later for arbitrary ℓ_p norms [9]. Ideas from the former paper stimulated many of the advances achieved by [2, 24, 27]. To the best of our knowledge, this is the first time that techniques from the latter paper have been applied in computational complexity.

For showing that CVP (among other problems) in the ℓ_p norm is in coNP for $\tilde{O}(\sqrt{n})$ approximation factors, we directly apply certain *measure inequalities* from [9] to the framework laid out by Aharonov and Regev [2] for the ℓ_2 norm. The main tool in [2] is a function f which distinguishes points which are close to a lattice from those which are very far from the lattice. The measure inequalities (see Section 3) will guarantee that $f(\mathbf{x})$ is *small* for any point $\mathbf{x}$ whose ℓ_p distance from the lattice is at least $n^{1/p}$. At the same time, $f(\mathbf{x})$ is guaranteed to be *large* for any point $\mathbf{x}$ within ℓ_p distance at most $\sim n^{1/p-1/2}$ from the lattice, by standard properties of norms. These facts are the essence of the $\tilde{O}(\sqrt{n})$ gap for the resulting coNP proof system.

For analyzing the worst-case to average-case reductions of Micciancio and Regev [24] and Regev [27], we will need to derive new facts about the *discrete Gaussian* probability distributions over lattices that emerge in their reductions. Specifically, we analyze *sums* of independent samples from discrete Gaussians, and show that in several important respects they behave just like continuous Gaussians (see Section 5 for details). This significantly generalizes prior analyses by Micciancio and Regev [24] and by Lyubashevsky and Micciancio [22], while providing an arguably more tractable proof.

1.3 Open Questions

The complexity of lattice problems in ℓ_p norms for $1 \leq p < 2$ is less well understood. Our techniques do not seem to be as useful for these norms, due to an asymmetry in how they relate to the ℓ_2 norm. We are unable to conclude anything other than what is already implied by basic relations among norms, i.e.: problems in coNP for $\tilde{O}(n^{1/p})$ factors, and worst-case to average-case reductions with $\tilde{O}(n^{1/2+1/p})$ connection factors. We remark that there is a natural correspondence between the ℓ_p norm and its *dual* ℓ_q norm, where $1/p + 1/q = 1$ and $1 \leq p \leq 2 \leq q \leq \infty$. It seems plausible that lattice problems in the ℓ_p norm could be related to problems in the ℓ_q norm via this duality.

Another important question is whether there are approximation-preserving reductions from problems in the ℓ_p norm to problems in the ℓ_2 norm. Combined with the results of [28], such reductions would imply that all ℓ_p norms are (essentially) equally hard for *any* approximation factor, not just those that emerge in proof systems or worst-case/average-case reductions.

1.4 Organization

In Section 2 we review lattices, computational problems, and Gaussian measures. In Section 3 we explain some of Banaszczyk's previously unexploited measure inequalities and their immediate implications. In Section 4 we use these inequalities to demonstrate that several lattice problems in ℓ_p norms are in coNP. In Section 5 we develop new tools for analyzing sums of discrete Gaussian distributions. In Section 6 we apply these tools by extending the analysis of prior worst-case to average-case reductions to other ℓ_p norms.

2 Preliminaries

2.1 Notation

We denote set of real numbers by $\mathbb{R}$ and the integers by $\mathbb{Z}$. For a positive integer n, $[n]$ denotes $\{1, \ldots, n\}$. The function log will always denote the natural logarithm. Extend any function $f(\cdot)$ to a countable set A in the following way: $f(A) = \sum_{x \in A} f(x)$.

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 $\mathbb{R}^n$ 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. $\mathbf{X}$. The ith column vector of $\mathbf{X}$ is denoted $\mathbf{x}_i$. We denote the standard inner product between $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ as $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i \in [n]} x_i y_i$. For simplicity, we sometimes write $\mathbf{x}^2$ for $\langle \mathbf{x}, \mathbf{x} \rangle$.

It is well-known that for any $\mathbf{x} \in \mathbb{R}^n$ and any $p \in [2, \infty]$, we have $n^{1/p-1/2} \left\|\mathbf{x}\right\|_2 \leq \left\|\mathbf{x}\right\|_p \leq \left\|\mathbf{x}\right\|_2$, whereas for any $p \in [1, 2]$, we have $\left\|\mathbf{x}\right\|_2 \leq \left\|\mathbf{x}\right\|_p \leq n^{1/p-1/2} \left\|\mathbf{x}\right\|_2$. For any $\mathbf{t} \in \mathbb{R}^n$ and set $V \subseteq \mathbb{R}^n$, define $\operatorname{dist}^p(\mathbf{t}, V) = \inf_{\mathbf{v} \in V} \left\|\mathbf{t} - \mathbf{v}\right\|_p$. Let $\mathcal{B}_n^p = \{\mathbf{x} \in \mathbb{R}^n \ : \ \left\|\mathbf{x}\right\|_p \leq 1\}$ denote the n-dimensional unit ball under the ℓ_p norm.

Let $\Gamma(z)$ denote the Euler Gamma function for real $z > 0$, defined as $\Gamma(z) = 2\int_{r=0}^{\infty} r^{2z-1} e^{-r^2}\, dr$. We write $\text{poly}(\cdot)$ for some unspecified polynomial function in its parameter. A function $f(n)$ is *negligible* in n if it decreases faster than the inverse of any polynomial in n.

2.2 Lattices

A *lattice* in $\mathbb{R}^n$ is

$$\Lambda = \mathcal{L}(\mathbf{B}) = \{\mathbf{Bc} \,:\, \mathbf{c} \in \mathbb{Z}^n\}\,,\ \mathbf{B} \in \mathbb{R}^{n \times n}$$

where the columns $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^n$ of $\mathbf{B}$ are linearly independent.[2] The matrix $\mathbf{B}$ is a *basis* of the lattice, and the columns $\mathbf{b}_i$ are *basis vectors.* A given lattice Λ has infinitely-many bases, which are related by unimodular transformations.

The *minimum distance* in ℓ_p norm of a lattice Λ, denoted $\lambda_1^p(\Lambda)$, is the length of its shortest nonzero element (in ℓ_p norm): $\lambda_1^p(\Lambda) = \min_{0 \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\|_p$. More generally, the *ith successive minimum* in ℓ_p norm $\lambda_i^p(\Lambda)$ is the smallest radius r such that the ball $r\mathcal{B}_n^p$ contains i linearly independent points of Λ. The *covering radius* in ℓ_p norm of Λ, denoted $\mu^p(\Lambda)$, is the smallest radius r such that balls $r\mathcal{B}_n^p$ centered at all points of Λ cover all of $\mathbb{R}^n$, i.e. $\mu^p(\Lambda) = \max_{\mathbf{x} \in \mathbb{R}^n} \text{dist}^p(\mathbf{x}, \Lambda)$.

The *dual lattice* of Λ, denoted Λ^*, is defined to be $\Lambda^* = \{\mathbf{x} \in \mathbb{R}^n \,:\, \forall\, \mathbf{v} \in \Lambda, \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbb{Z}\}$.

2.3 Problems on Lattices

Here we define some standard worst-case problems on lattices. See [23, 24] for motivation and discussion of these problems. All of the following are approximation problems parameterized by a positive function $\gamma = \gamma(n)$ of the dimension.

We define the following problems in their decisional *promise* versions.

Definition 2.1 (Shortest Vector Problem)**.** An input to GapSVP_γ^p is a pair $(\mathbf{B}, d)$ where $\mathbf{B}$ is an n-dimensional lattice basis and $d \in \mathbb{R}$. It is a YES instance if $\lambda_1^p(\mathcal{L}(\mathbf{B})) \leq d$, and is a NO instance if $\lambda_1^p(\mathcal{L}(\mathbf{B})) > \gamma(n) \cdot d$.

Definition 2.2 (Closest Vector Problem)**.** An input to GapCVP_γ^p is a tuple $(\mathbf{B}, \mathbf{v}, d)$ where $\mathbf{B}$ is an n-dimensional lattice basis, $\mathbf{v} \in \mathbb{R}^n$, and $d \in \mathbb{R}$. It is a YES instance if $\text{dist}^p(\mathbf{v}, \mathcal{L}(\mathbf{B})) \leq d$, and is a NO instance if $\text{dist}^p(\mathbf{v}, \mathcal{L}(\mathbf{B})) > \gamma(n) \cdot d$.

We informally define the *closest vector with preprocessing* problem $\mathsf{GapCVPP}$, whose goal is to solve GapCVP on some *fixed* lattice for a given target point, allowing for the use of some arbitrary short advice about the lattice. See [14] for motivation and a formal definition.

Definition 2.3 (Covering Radius Problem)**.** An input to GapCRP_γ^p is a pair $(\mathbf{B}, d)$ where $\mathbf{B}$ is an n-dimensional lattice basis and $d \in \mathbb{R}$. It is a YES instance if $\mu^p(\mathcal{L}(\mathbf{B})) \leq d$ and is a NO instance if $\mu^p(\mathcal{L}(\mathbf{B})) > \gamma(n) \cdot d$.

We define the following problems in their search versions.

Definition 2.4 (Shortest Independent Vectors Problem)**.** An input to SIVP_γ^p is an n-dimensional lattice basis $\mathbf{B}$. The goal is to output a set of n linearly independent lattice vectors $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{S}\|_p \leq \gamma(n) \cdot \lambda_n^p(\mathcal{L}(\mathbf{B}))$.

[2] Technically, this is the definition of a *full-rank* lattice, which is the only kind of lattice we will be concerned with.

The *guaranteed distance decoding* problem GDD and its *incremental* version IncGDD are variants of the closest vector problem. In this work we only need them to state an intermediate result on worst-case/average-case reductions, therefore we omit their precise definitions. See [24] for details.

2.4 Gaussian Measures

Our review of Gaussian measures over lattices follows the development by prior works [2, 26, 24]. 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 $D_{s,\mathbf{c}}(\mathbf{x}) = s^{-n} \cdot \rho_{s,\mathbf{c}}(\mathbf{x})$.

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 [24] proposed a new lattice quantity which they called the *smoothing parameter*:

Definition 2.5 ([24])**.** 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 formal statement and proof of this fact can be found in [24].) The smoothing parameter is closely related to the successive minima of the lattice:

Lemma 2.6. *For any n-dimensional lattice Λ, real $p \in [2, \infty]$, and real $\epsilon > 0$, we have*

$$\eta_\epsilon(\Lambda) \quad \leq \quad \lambda_n(\Lambda) \cdot \sqrt{\frac{\log(2n(1+1/\epsilon))}{\pi}} \quad \leq \quad \lambda_n^p(\Lambda) \cdot n^{1/2-1/p} \cdot \sqrt{\frac{\log(2n(1+1/\epsilon))}{\pi}}.$$

Proof. The first inequality is due to [24]. The second follows by $\|\mathbf{x}\|_2 \leq n^{1/2-1/p} \|\mathbf{x}\|_p$ for $p \geq 2$. □

The smoothing parameter also influences the behavior of discrete Gaussians over the lattice. In our new analysis of discrete Gaussians we will rely upon the following simple lemma:

Lemma 2.7 ([24], implicit)**.** *For any $s \geq \eta_\epsilon(\Lambda)$, real $\epsilon \in (0,1)$, and $\mathbf{c} \in \mathbb{R}^n$, we have*

$$\tfrac{1-\epsilon}{1+\epsilon} \cdot \rho_s(\Lambda) \leq \rho_{s,\mathbf{c}}(\Lambda) \leq \rho_s(\Lambda).$$

For the worst-case to average-case reduction from GapSVP we will need the following lemma:

Lemma 2.8 ([24, implicit in Lemma 4.5])**.** *For any n-dimensional lattice Λ, real $\epsilon \in (0,1)$, real $s \geq \eta_\epsilon(\Lambda)$, and $\mathbf{c}, \mathbf{v} \in \mathbb{R}^n$, we have:*

$$\left| \mathop{\mathrm{E}}_{\mathbf{x} \sim D_{\Lambda,s,\mathbf{c}}} \left[e^{2\pi i \langle \mathbf{x}, \mathbf{v} \rangle} \right] \right| \leq \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\rho_{1/s}(\Lambda^* - \mathbf{v})}{\rho_{1/s}(\Lambda^*)}.$$

3 Measure Inequalities for ℓ_p Norms

In this section we review some inequalities developed by Banaszczyk [9] and a few of their immediate consequences for our applications.

The goal of these inequalities is to bound the total Gaussian measure $\rho((\Lambda + \mathbf{v}) \setminus r\mathcal{B}_n^p)$ assigned to those points of a shifted lattice $\Lambda + \mathbf{v}$ whose ℓ_p norm exceeds a certain radius r. The measure is typically normalized by the total measure $\rho(\Lambda)$ on the entire unshifted lattice, yielding a ratio between 0 and 1. This ratio has proven to be a crucial quantity in obtaining transference theorems for lattices [8, 9], and in the study of the computational complexity of lattice problems [1, 2, 24].

In a prior work of Banaszczyk [8], it was shown that for $p = 2$ and radius $r = \sqrt{n}$, the ratio described above is *exponentially small* in n. The results below are generalizations of this statement to arbitrary ℓ_p norms. Loosely speaking, they show that for some suitable constant C, the normalized measure is small for $r = C \cdot n^{1/p}$. The ratio is *not*, generally speaking, exponentially small, but for our applications we will only need it to be a small constant. The important fact is that as a function of p, we can obtain a ratio bounded away from 1 using a radius $r \sim n^{1/p}$.

Lemma 3.1 ([9, Lemma 2.9]). *For any n-dimensional lattice Λ, $p \in [1, \infty)$, $\mathbf{v} \in \mathbb{R}^n$, and real $r > 0$, we have:*

$$\frac{\rho((\Lambda + \mathbf{v}) \setminus rB_n^p)}{\rho(\Lambda)} < p\pi^{-p/2}\Gamma\left(\frac{p}{2}\right) \cdot n \cdot r^{-p}.$$

Corollary 3.2. *For any $p \in [1, \infty)$, there is a constant $c_p \approx \sqrt{p}$ such that for any n-dimensional lattice Λ and $\mathbf{v} \in \mathbb{R}^n$,*

$$\frac{\rho((\Lambda + \mathbf{v}) \setminus c_p n^{1/p} \cdot B_n^p)}{\rho(\Lambda)} < 1/4.$$

Proof. Follows immediately from Lemma 3.1 by setting

$$r = \left(4\pi^{-p/2} \cdot p \cdot \Gamma\left(\frac{p}{2}\right)\right)^{1/p} \cdot n^{1/p} \approx \sqrt{p} \cdot n^{1/p}. \qquad \square$$

Lemma 3.3 ([9, Lemma 2.10]). *For any n-dimensional lattice Λ, $\mathbf{v} \in \mathbb{R}^n$, and real $r > 0$, we have:*

$$\frac{\rho((\Lambda + \mathbf{v}) \setminus rB_n^\infty)}{\rho(\Lambda)} < 2ne^{-\pi r^2}.$$

Corollary 3.4. *There is a constant c such that for any n-dimensional lattice Λ and $\mathbf{v} \in \mathbb{R}^n$,*

$$\frac{\rho((\Lambda + \mathbf{v}) \setminus c\sqrt{\log n} \cdot B_n^\infty)}{\rho(\Lambda)} < 1/4.$$

Proof. Follows immediately from Lemma 3.3 by setting $r = \sqrt{\frac{\log(8n)}{\pi}} \leq c\sqrt{\log n}$ for some c. □

3.1 Smoothing Parameter

The measure inequalities from the previous section yield bounds on the smoothing parameter relative to lattice minima in ℓ_p norms. We will need these bounds for showing the worst-case to average-case reductions for GapSVP in ℓ_p norms, later in the paper.

Lemma 3.5. *For any n-dimensional lattice Λ, $p \in [1, \infty]$, and real $\epsilon > 0$, we have:*

$$\eta_\epsilon(\Lambda) \leq \frac{n^{1/p} \cdot \sqrt{\log(2n/\epsilon)/\pi}}{\lambda_1^p(\Lambda^*)}.$$

Proof. Let $s > \frac{n^{1/p} \cdot \sqrt{\log(2n/\epsilon)/\pi}}{\lambda_1^p(\Lambda^*)}$. Then

$$\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) = \rho(s\Lambda^* \setminus \{\mathbf{0}\}) = \rho(s\Lambda^* \setminus s\lambda_1^p(\Lambda^*) \cdot \mathcal{B}_n^p) \leq \rho(s\Lambda^* \setminus s\lambda_1^p(\Lambda^*) \cdot n^{-1/p} \cdot \mathcal{B}_n^\infty) < \epsilon,$$

where we have used that $\mathcal{B}_n^p \supseteq n^{-1/p} \cdot \mathcal{B}_n^\infty$, and the final inequality follows from Lemma 3.3. □

4 Problems in coNP

In this section, we show that for certain $\gamma(n) = \tilde{O}(\sqrt{n})$ approximation factors, the following decisional "gap" problems in ℓ_p norm are contained in coNP: the shortest vector problem GapSVP_γ^p, the closest vector problem GapCVP_γ^p, the covering radius problem GapCRP_γ^p, and the shortest independent vectors problem $\mathsf{GapSIVP}_\gamma^p$. This implies that for these problems are not NP-hard unless the polynomial hierarchy collapses (see [16, 15] for a discussion of some subtleties concerning promise problems and the polynomial hierarchy). For similar approximation factors, we also show that the closest vector with preprocessing problem GapCVPP in ℓ_p norm is *easy* (i.e., in P).

These results are intended partly as a warm-up for the more complicated analysis of discrete Gaussians later in the paper. Indeed, the results in this section are a simple application of the measure inequalities from Section 3 to prior work by Aharonov and Regev [2], who developed the main techniques for the ℓ_2 norm. Therefore we will omit many technical details, and direct the reader to [2] for full discussions.

4.1 Closest Vector Problem

The main result we need is the containment $\mathsf{GapCVP}_\gamma^p \in \mathsf{coNP}$ for $p \in [2, \infty]$ and certain choices of $\gamma(n) = \tilde{O}(\sqrt{n})$. (The coNP verifier for GapCVP will also play a role in the worst-case to average-case reductions of Section 6.) The remaining results will follow by known reductions to GapCVP, which work for arbitrary ℓ_p norms and approximation factors.

Here we give a brief informal overview of the main proof technique of Aharonov and Regev [2] for GapCVP. It is shown in [2] that for any n-dimensional lattice Λ, there is a positive function $f : \mathbb{R}^n \to [0, 1]$ which indicates whether an arbitrary point $\mathbf{v} \in \mathbb{R}^n$ is close to, or far from, the lattice (in ℓ_2 norm): when $\mathbf{v}$ is close, $f(\mathbf{v})$ is large; when $\mathbf{v}$ is *very* far, $f(\mathbf{v})$ is small. More precisely, when $\mathbf{v}$ is within, say, distance $1/100$ of the lattice, $f(\mathbf{v}) \geq 1/2$; when $\mathbf{v}$ is more than $\sqrt{n}$ away from the lattice, $f(\mathbf{v})$ is exponentially small in n. The precise definition of f is the (normalized) sum of Gaussians centered at every lattice point, i.e. $f(\mathbf{v}) = \rho(\Lambda + \mathbf{v})/\rho(\Lambda)$.

It is also shown in [2] that f can be *succinctly approximated*, by choosing elements from the dual lattice Λ^* under an appropriate distribution. This leads to an NP proof system for the fact that $\mathbf{v}$ is far from the lattice. The witness is a succinct representation of a function $\tilde{f} \approx f$. The verifier accepts if $\tilde{f}(\mathbf{v})$ is small, and if $\tilde{f}$ is a good enough approximation to f (this is more technical, but can also be done efficiently). On the other hand, if $\mathbf{v}$ is actually close to the lattice, then $\tilde{f}(\mathbf{v})$ will always be large for any acceptable $\tilde{f}$, and the verifier rejects.

Analysis for ℓ_p norms. We now consider arbitrary ℓ_p norms, $p \geq 2$. It turns out that we can use *exactly the same* verifier and witness as in [2]; only the analysis is different. We make the following observations: if $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p-1/2}/100$, then $\mathrm{dist}^2(\mathbf{v}, \Lambda) \leq 1/100$ by properties of ℓ_p norms. In such a case, we already know that the verifier always rejects. On the other hand, if $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c_p n^{1/p}$ for some appropriate constant c_p, then the measure inequalities for ℓ_p norms guarantee that $\tilde{f}(\mathbf{v}) \approx f(\mathbf{v})$ is a small constant, and the verifier accepts. The resulting gap factor is therefore $O(n^{1/p}/n^{1/p-1/2}) = O(\sqrt{n})$.

We conclude this informal overview with a discussion of ℓ_p norms, $1 \leq p < 2$. For these norms, completeness still holds when $\mathrm{dist}^p(\mathbf{v}, \Lambda) > Cn^{1/p}$. However, soundness is compromised: if $\mathrm{dist}^p(\mathbf{v}, \Lambda) = n^{1/p-1/2}$, it may still be the case that $\mathrm{dist}^2(\mathbf{v}, \Lambda) = n^{1/p-1/2} \gg 1$. The only way to guarantee that $\mathbf{v}$ is close enough to Λ in ℓ_2 norm is to require, say, $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq 1/100$. This yields an approximation factor of $O(n^{1/p})$, which was already known from [2] using the relations between ℓ_p norms. We do not know if there is an alternate proof system which improves upon this factor.

We now proceed to the detailed statement of the theorem and its proof.

Theorem 4.1. *For any $p \in [2, \infty)$, there is a constant $c_p \approx \sqrt{p}$ such that $\mathsf{GapCVP}^p_{c_p\sqrt{n}} \in \mathsf{NP} \cap \mathsf{coNP}$. For $p = \infty$, there is a constant c such that $\mathsf{GapCVP}^\infty_{c\sqrt{n \log n}} \in \mathsf{NP} \cap \mathsf{coNP}$.*

Proof. The containment in NP is trivial, as well as the proof for $n = 1$. Thus it suffices to prove GapCVP is in coNP, assuming $n \geq 2$. That is, we must show a polynomial-time algorithm which, given a lattice basis $\mathbf{B}$, a point $\mathbf{v}$, and a witness, verifies that $\mathbf{v}$ is *far* from the lattice $\mathcal{L}(\mathbf{B})$.

The verifier $\mathcal{V}$ we use is the same one from from [2]; we recall it here. The input to $\mathcal{V}$ is an instance $(\mathbf{B}, \mathbf{v}, d)$ of GapCVP, plus a witness matrix $\mathbf{W} \in \mathbb{R}^{n \times N}$, for sufficiently large N. Let $\Lambda = \mathcal{L}(\mathbf{B})$. The verifier algorithm performs the following tests, and accepts if all three hold (otherwise it rejects):

1. Check that $f_{\mathbf{W}}(\mathbf{v}) < 1/2$, where $f_{\mathbf{W}}$ is the function $f_{\mathbf{W}}(\mathbf{v}) = \frac{1}{N} \sum_{i \in [N]} \cos(2\pi \langle \mathbf{w}_i, \mathbf{v} \rangle)$.

2. Check that $\mathbf{w}_i \in \Lambda^*$ for all $i \in [N]$, i.e. that $\mathbf{w}_i$ are dual lattice vectors.

3. Check that the largest eigenvalue of the matrix $\mathbf{W}\mathbf{W}^T$ is at most $N/(2\pi d)^2$.

As argued in [2], $\mathcal{V}$ can be implemented in polynomial time.

We now demonstrate the correctness of the verifier for all ℓ_p norms, $p \in [2, \infty]$. First, we perform the following rescaling: we map an instance $(\mathbf{B}, \mathbf{v}, d')$ of GapCVP^p to an instance $(\mathbf{B}, \mathbf{v}, d)$ where $d = d' \cdot n^{1/2-1/p}$, and invoke the verifier $\mathcal{V}$ on that instance (and the same witness $\mathbf{W}$).

Soundness. Suppose $(\mathbf{B}, \mathbf{v}, d')$ is a NO instance, i.e. $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq d'$. Then

$$\mathrm{dist}^2(\mathbf{v}, \Lambda) \leq d' \cdot n^{1/2-1/p} = d$$

by the properties of ℓ_p norms. In [2] it is shown that $\mathcal{V}$ always rejects in this case.

Completeness. Suppose now that $(\mathbf{B}, \mathbf{v}, d')$ is a YES instance, i.e. $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c_p\sqrt{n} \cdot d' = c_p n^{1/p} \cdot d$ for $p \in [2, \infty)$, or $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c\sqrt{n \log n} \cdot d' = c\sqrt{\log n} \cdot d$ for $p = \infty$.

In [2] it is shown that when the vectors $\mathbf{w}_i$ of $\mathbf{W}$ are chosen independently from a certain distribution (i.e., the discrete Gaussian $D_{\Lambda^*, 1/d}$ over the dual lattice Λ^*), Test 3 is satisfied with overwhelming probability, and Test 2 is always satisfied by definition of the distribution.

It remains to show that Test 1 is satisfied, i.e. $f_{\mathbf{W}}(\mathbf{v}) < 1/2$, with some positive constant probability over the choice of $\mathbf{W}$, which implies the existence of a $\mathbf{W}$ which makes $\mathcal{V}$ accept. The Pointwise Approximation Lemma of [2] (with appropriate scaling) states that with probability at least 3/4 over the random choice of $\mathbf{W}$, the difference

$$\left| f_{\mathbf{W}}(\mathbf{v}) - \frac{\rho_d(\Lambda + \mathbf{v})}{\rho_d(\Lambda)} \right| = \left| f_{\mathbf{W}}(\mathbf{v}) - \frac{\rho(\Lambda' + \mathbf{v}')}{\rho(\Lambda')} \right| \leq 1/n^2,$$

where $\Lambda' = \Lambda/d$ and $\mathbf{v}' = \mathbf{v}/d$ are a rescaled lattice and target point.

First take $p \in [2, \infty)$. Then $\mathrm{dist}^p(\mathbf{v}', \Lambda') > c_p n^{1/p}$, so $\rho(\Lambda' + \mathbf{v}') = \rho((\Lambda' + \mathbf{v}') \setminus c_p n^{1/p} \cdot \mathcal{B}_n^p)$. Then by Corollary 3.2 we have $f_{\mathbf{W}}(\mathbf{v}) < 1/4 + 1/n^2 \leq 1/2$ with probability at least 3/4.

Now take $p = \infty$. Then $\mathrm{dist}^\infty(\mathbf{v}', \Lambda') > c\sqrt{\log n}$. By a similar argument using Corollary 3.4, the proof is complete. □

4.2 Other Problems

Theorem 4.2. *For any $p \in [2, \infty)$, there is a constant $c_p \approx \sqrt{p}$ such that all of the following problems are in* coNP *for $\gamma(n) = c_p\sqrt{n}$:* GapSVP^p_γ, GapCRP^p_γ, *and* $\mathsf{GapSIVP}^p_\gamma$.

For $p = \infty$, there is a constant c such that all of the above are in coNP *for $\gamma(n) = c\sqrt{n \log n}$.*

Proof. The theorem follows from known approximation- and norm-preserving reductions to GapCVP. For GapSVP, there is a reduction from to GapCVP due to Goldreich *et al* [17]. For GapCRP, there is a simple nondeterministic reduction to GapCVP, due to Guruswami *et al* [18], in which the reduction guesses a "deep hole" (i.e., a point far from the lattice) which becomes the target point in the resulting GapCVP instance. This reduction suffices to show inclusion in coNP. For GapSIVP, there is a more complicated nondeterministic reduction, also due to Guruswami *et al* [18], to GapCVP. □

4.3 Closest Vector with Preprocessing

Aharonov and Regev also showed that GapCVPP in ℓ_2 norm is easy for some $\gamma(n) = c\sqrt{n/\log n}$ [2]. Applying a similar analysis as above, it is easy to extend this result to ℓ_p norms.

Theorem 4.3. *For any $p \in [2, \infty)$, there is a constant $c_p \approx \sqrt{p}$ such that* $\mathsf{GapCVPP}^p_{c_p\sqrt{n/\log n}} \in \mathsf{P}$. *For $p = \infty$, there is a constant c such that* $\mathsf{GapCVPP}^\infty_{c\sqrt{n}} \in \mathsf{P}$.

5 New Analysis of Discrete Gaussians

In this section, we develop new tools that are useful for analyzing worst-case to average-case reductions that rely on Gaussians. We will analyze *sums* of independent samples from discrete Gaussian distributions over lattices. More precisely, we analyze the *moments* of such sums when projected onto a subspace of $\mathbb{R}^n$. Our main result is that these moments are nearly identical to those for sums of *continuous* Gaussian distributions. This can be used to derive very "intuitive" bounds on expectations and tail probabilities for ℓ_p norms. In summary, our analysis further reinforces the idea that discrete Gaussians over lattices behave remarkably similarly to continuous Gaussians.

Our analysis seems a natural continuation of prior study into discrete Gaussians. Using ideas from Banaszczyk [8], Micciancio and Regev [24] analyzed a few low-order moments of a discrete Gaussian projected onto a one-dimensional subspace of $\mathbb{R}^n$. Lyubashevsky and Micciancio [22] extended this analysis to all the higher moments. Unfortunately, even at this stage the analysis becomes quite cumbersome, involving several pages of heavy manipulations.

We analyze *sums* of discrete Gaussians over lattices, projected onto *arbitrary subspaces* of $\mathbb{R}^n$. Despite these generalities, our analysis is actually simpler, due crucially to the use of techniques from a follow-up work of Banaszczyk [9].

5.1 Overview of Techniques

In this subsection we give a simplified and high-level overview of our techniques for analyzing sums of discrete Gaussians. Suppose without loss of generality that Λ is a sufficiently dense lattice in $\mathbb{R}^n$. Suppose that $\mathbf{x} \in \Lambda$ is a random variable distributed as the *sum* of m independent samples from the discrete Gaussian $D_{\Lambda,1}$.[3] We are interested in calculating the expected length of $\mathbf{x}$ in the ℓ_p norm, $\mathrm{E}[\|\mathbf{x}\|_p]$. By Jensen's inequality and linearity of expectation, this is at most

$$\left(\mathrm{E}\left[\|\mathbf{x}\|_p^p\right]\right)^{1/p} = \left(\sum\nolimits_{i \in [n]} \mathrm{E}\left[|x_i|^p\right]\right)^{1/p}, \tag{1}$$

so it suffices to bound $\mathrm{E}\left[|x_i|^p\right]$. The crucial tool we need is an exponential tail inequality on x_i:

Tail Inequality. *For any $r \geq 0$, the probability that $|x_i| > r$ decays exponentially with r^2/m:*

$$\Pr_{\mathbf{x}}[|x_i| > r] \approx \exp(-\Theta(r^2/m)).$$

This inequality is stated precisely and in full generality as Lemma 5.4 below. We remark that for sums of *continuous* Gaussians, proving this inequality is straightforward — a sum of Gaussians is just another Gaussian (with a larger variance), which can be analyzed by direct integration. However, for sums of *discrete* Gaussians the path is not so straightforward.

To prove the tail inequality and complete the analysis, we will draw upon techniques of Banaszczyk [8, 9]. First, view $\mathbf{x}$'s probability distribution as a positive function $D : \Lambda \to \mathbb{R}^+$. Then our goal is to bound the total measure assigned by D to $\Lambda^- = \{\mathbf{x} \in \Lambda : |x_i| > r\}$. The general strategy is to find some positive function $g : \Lambda \to \mathbb{R}^+$ satisfying two conditions: (1) the total measure $(D \cdot g)(\Lambda)$ only exceeds the total measure $D(\Lambda)$ by a "small" factor c, but (2) the total measure $(D \cdot g)(\Lambda^-)$ exceeds the total measure $D(\Lambda^-)$ by a "very large" factor C. Because D and g are positive, we get

$$C \cdot D(\Lambda^-) \leq (D \cdot g)(\Lambda^-) \leq (D \cdot g)(\Lambda) \leq c \cdot D(\Lambda) = c,$$

from which we conclude that the tail probability $D(\Lambda^-) \leq c/C$, a small quantity. It turns out that an appropriate choice for the function g is $g(\mathbf{x}) = \cosh(2\pi r\,|x_i|\,/m)$, where $\cosh(x) = \frac{1}{2}(e^x + e^{-x})$ is the hyperbolic cosine.

With the tail inequality in hand, the expectation $\mathrm{E}\left[|x_i|^p\right]$ can be written as an integral:

$$\begin{aligned}\sum_{\mathbf{x} \in \Lambda} |x_i|^p \cdot \Pr[\mathbf{x}] &= \sum_{\mathbf{x} \in \Lambda} \left(\int_{r=0}^{|x_i|} p r^{p-1}\, dr\right) \Pr[\mathbf{x}] = \int_{r=0}^{\infty} p r^{p-1} \left(\sum_{\mathbf{x} \in \Lambda,\, |x_i| > r} \Pr[\mathbf{x}]\right) dr \\ &= \int_{r=0}^{\infty} p r^{p-1} \cdot \Pr\left[|x_i| > r\right] dr \approx p \int_{r=0}^{\infty} r^{p-1} \exp(-\Theta(r^2/m))\, dr. \end{aligned} \tag{2}$$

[3] In the general case, the samples may be drawn from discrete Gaussians with different parameters $s \geq \eta_\epsilon$, different centers $\mathbf{c}$, and even defined over different lattices. In order to illuminate the key ideas, we focus on a much simpler case in this overview.

The final expression is a so-called *Gaussian integral*, which has a known closed form that evaluates to roughly $(\sqrt{pm})^p$. When plugged into Equation (1), this yields

$$\mathrm{E}[\|\mathbf{x}\|_p] \leq \sqrt{p} \cdot \sqrt{m} \cdot \sqrt[p]{n}.$$

In conclusion, for any fixed ℓ_p norm, the sum of m discrete n-dimensional Gaussians (with parameter $s = 1$) has expected norm that grows as $\sqrt{m}$ and $\sqrt[p]{n}$, just as with continuous Gaussians. For the ℓ_∞ norm, there is simply an extra $\sqrt{\log n}$ factor.

We devote the remainder of this section to the full statement of the result and its proof. We caution that the proof is technical and somewhat heavy in notation in places; the reader who is interested only in the applications may safely skip to Section 6.

5.2 Preliminaries and Notation

Let $\Lambda_1, \ldots, \Lambda_m$ be arbitrary lattices in $\mathbb{R}^n$, and let $\mathbf{\Lambda} = \Lambda_1 \times \cdots \times \Lambda_m$. Likewise, let $\mathbf{s} \in (\mathbb{R}^+)^m$ be a vector of positive parameters for m Gaussians, and let $\mathbf{C} \in \mathbb{R}^{n\times m}$ be the matrix of their m centers. Let $\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}} : \mathbf{\Lambda} \to \mathbb{R}$ be the function defined as the product of the corresponding Gaussian functions:

$$\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X}) = \prod_{i\in[m]} \rho_{s_i,\mathbf{c}_i}(\mathbf{x}_i).$$

(We may omit $\mathbf{C}$ when it is the all-zeros matrix.) It immediately follows from Lemma 2.7 that:

Lemma 5.1. *Let $\mathbf{\Lambda}$, $\mathbf{s}$, $\mathbf{C}$ be as above, and let $\epsilon > 0$ be such that $s_i \geq \eta_\epsilon(\Lambda_i)$ for all $i \in [m]$. Then:*

$$1 \leq \frac{\boldsymbol{\rho}_{\mathbf{s}}(\mathbf{\Lambda})}{\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{\Lambda})} \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)^m.$$

Denote by $D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}$ the joint distribution over $\mathbb{R}^{n\times m}$, and having support $\mathbf{\Lambda}$, given by sampling *independently* from $D_{\Lambda_i,s_i,\mathbf{c}_i}$ for each i. That is, for $\mathbf{X} \in \mathbf{\Lambda}$,

$$D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}(\mathbf{X}) = \prod_{i\in[m]} D_{\Lambda_i,s_i,\mathbf{c}_i}(\mathbf{x}_i) = \frac{\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X})}{\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{\Lambda})}.$$

For any $\mathbf{X} \in \mathbb{R}^{n\times m}$, define $h_{\mathbf{C}}(\mathbf{X}) = \sum_{i\in[m]}(\mathbf{x}_i - \mathbf{c}_i)$, the sum of the $\mathbf{x}_i$s shifted by the centers $\mathbf{c}_i$.

Finally, let $\mathbf{U} = \{\mathbf{u}_1, \ldots, \mathbf{u}_d\}$ be a set of $d \geq 1$ orthonormal vectors in $\mathbb{R}^n$. For concreteness (and indeed, for all of our applications in this work), we can take $d = 1$ and let $\mathbf{u}_1$ be one of the standard basis elements for $\mathbb{R}^n$. Define the "$\mathbf{U}$ norm" as $\|\mathbf{y}\|_{\mathbf{U}} = \sum_{i\in[d]} |\langle \mathbf{y}, \mathbf{u}_i\rangle|$ for any $\mathbf{y} \in \mathbb{R}^n$. In our concrete example of $\mathbf{U}$, then, $\|\mathbf{y}\|_{\mathbf{U}}$ is simply the absolute value of one of $\mathbf{y}$'s coordinates. More generally, the $\mathbf{U}$ norm is akin to the ℓ_1 norm within the subspace spanned by $\mathbf{U}$.

5.3 Moments of Gaussian Sums

Our main theorem concerns the sum of independent discrete Gaussians over lattices. The theorem bounds all the *moments* of the sum, where the moments are taken about the sum of the Gaussian's centers. Using our notation from above, then, we are concerned with the distribution of $h_{\mathbf{C}}(\mathbf{X})$, where $\mathbf{X} \sim D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}$. Of course, $h_{\mathbf{C}}(\mathbf{X})$ is distributed over $\mathbb{R}^n$, whereas moments usually refer to distributions over $\mathbb{R}$. We handle this mismatch by actually considering the moments of $\|h_{\mathbf{C}}(\mathbf{X})\|_{\mathbf{U}}$, i.e. the $\mathbf{U}$ norm of the sum of Gaussians.

Theorem 5.2 (Main Theorem: Moments of discrete Gaussian sums)**.** *Let* $\mathbf{\Lambda}$*,* $\mathbf{s}$*,* $\mathbf{C}$*, and* $\mathbf{U}$ *be as above, and let* $\epsilon > 0$ *be such that* $s_i \geq \eta_\epsilon(\Lambda_i)$ *for all* $i \in [m]$*. Then for any* $p > 0$*,*

$$\mathop{\mathrm{E}}_{\mathbf{X} \sim D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}} \left[\|h_{\mathbf{C}}(\mathbf{X})\|_{\mathbf{U}}^p\right] \leq 2^{d-1} \cdot \left(\frac{1+\epsilon}{1-\epsilon}\right)^m \cdot \left\|\mathbf{s}\sqrt{d/\pi}\right\|_2^p \cdot p \cdot \Gamma\left(\frac{p}{2}\right).$$

For the typical case where $d = 1$ (a concurrent work [25] also needs to consider $d = 2$, but no more), $m = m(n)$ is some small polynomial, and $\epsilon = \epsilon(n)$ is some small inverse polynomial, the following corollary is in a form more suitable for application:

Corollary 5.3. *Let* $\mathbf{\Lambda}$*,* $\mathbf{s}$*,* $\mathbf{C}$ *be as above, let* $m = m(n) = \mathrm{poly}(n)$*, and let* $\epsilon(n) \leq 1/(2m(n) + 1)$ *be such that* $s_i \geq \eta_\epsilon(\Lambda_i)$ *for all* $i \in [m]$*. For any* $p \in [1, \infty)$*, there is a constant* $c_p \approx \sqrt{p}$ *such that:*

$$\mathop{\mathrm{E}}_{\mathbf{X} \sim D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}} \left[\|h_{\mathbf{C}}(\mathbf{X})\|_p\right] \leq c_p \cdot \|\mathbf{s}\|_2 \cdot n^{1/p}.$$

For $p = \infty$*, there is a universal constant* c *such that:*

$$\Pr_{\mathbf{X} \sim D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}} \left[\|h_{\mathbf{C}}(\mathbf{X})\|_\infty > c \cdot \|\mathbf{s}\|_2 \cdot \sqrt{\log n}\right] \leq 1/4.$$

Proof. First we consider the case $p \in [1, \infty)$. Let $\{\mathbf{e}_i\}_{i \in [n]}$ be the standard basis of $\mathbb{R}^n$. We have:

$$\begin{aligned}
\mathrm{E}\left[\|h_{\mathbf{C}}(\mathbf{X})\|_p\right] &\leq \left(\sum_{i \in [n]} \mathrm{E}\left[|\langle h_{\mathbf{C}}(\mathbf{X}), \mathbf{e}_i\rangle|^p\right]\right)^{1/p} && \text{(Jensen's ineq, linearity of E)} \\
&\leq \left(n \cdot \left(\tfrac{1+\epsilon}{1-\epsilon}\right)^m \cdot \|\mathbf{s}\|_2^p \cdot p \cdot \Gamma\left(\frac{p}{2}\right)\right)^{1/p} && \text{(Theorem 5.2)} \\
&\leq c_p \cdot \|\mathbf{s}\|_2 \cdot n^{1/p} && \text{(appropriate } c_p\text{; } \left(\tfrac{1+\epsilon}{1-\epsilon}\right)^m \leq e\text{)}
\end{aligned}$$

Now we consider $p = \infty$. Here we actually use Lemma 5.4 (Tail Inequality) directly, choosing $r = c \cdot \|\mathbf{s}\| \cdot \sqrt{\log n}$ for appropriate constant c so that for any $i \in [n]$,

$$\Pr_{\mathbf{X} \sim D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}} \left[|\langle h_{\mathbf{C}}(\mathbf{X}), \mathbf{e}_i\rangle| > r\right] \leq 1/4n.$$

By the union bound, the proof is complete. □

Before proving the main theorem, we will need one more piece of notation. For any $r \geq 0$ and for $\mathbf{C}$ and $\mathbf{U}$ as above, define

$$Q_r = \{\mathbf{X} \, : \, \|h_{\mathbf{C}}(\mathbf{X})\|_{\mathbf{U}} \leq r\} \subseteq \mathbb{R}^{n \times m}.$$

We now proceed to the proof.

Proof of Theorem 5.2. The proof exactly follows the structure of Equation (2) from our overview, just with heavier notation. The main tool is the tail inequality from Lemma 5.4 below.

$$\begin{aligned}
\underset{\mathbf{X}\sim D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}}{\mathrm{E}}\left[\|h_{\mathbf{C}}(\mathbf{X})\|_{\mathbf{U}}^p\right] &= \sum_{\mathbf{X}\in\mathbf{\Lambda}} \|h_{\mathbf{C}}(\mathbf{X})\|_{\mathbf{U}}^p \cdot D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}(\mathbf{X}) && \text{(def. of E)}\\
&= p\int_{r=0}^{\infty} r^{p-1}\cdot \sum_{\mathbf{X}\in\mathbf{\Lambda}\setminus Q_r} D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}}(\mathbf{X})\,dr && \text{(calculus; see (2))}\\
&\le 2^d\cdot p\int_{r=0}^{\infty} r^{p-1}\exp(-\pi r^2/d\cdot\|\mathbf{s}\|_2^2)\cdot\frac{\boldsymbol{\rho}_{\mathbf{s}}(\mathbf{\Lambda})}{\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{\Lambda})}\,dr && \text{(Lemma 5.4)}\\
&\le 2^d\cdot\left(\tfrac{1+\epsilon}{1-\epsilon}\right)^m\cdot p\int_{r=0}^{\infty} r^{p-1}\exp(-\pi r^2/d\cdot\|\mathbf{s}\|_2^2)\,dr && \text{(Lemma 5.1)}\\
&= 2^{d-1}\cdot\left(\tfrac{1+\epsilon}{1-\epsilon}\right)^m\cdot\left\|\mathbf{s}\sqrt{d/\pi}\right\|_2^p\cdot p\cdot\Gamma\left(\tfrac{p}{2}\right). && \text{(integration)}
\end{aligned}$$
□

Lemma 5.4 (Tail Inequality)**.** *Let* $\mathbf{\Lambda}$, $\mathbf{s}$, $\mathbf{C}$ *be as above. Then for any* $r\ge 0$,

$$\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{\Lambda}\setminus Q_r) \quad\le\quad 2^d\cdot\exp\left(\frac{-\pi r^2}{d\cdot\|\mathbf{s}\|_2^2}\right)\cdot\boldsymbol{\rho}_{\mathbf{s}}(\mathbf{\Lambda}).$$

Proof. First, for any $t\in\mathbb{R}$ define the positive function $g_t:\mathbf{\Lambda}\to\mathbb{R}^+$ as:

$$g_t(\mathbf{X}) = \prod_{k\in[d]}\cosh(2\pi t\,\langle h_{\mathbf{C}}(\mathbf{X}),\mathbf{u}_k\rangle).$$

The proof will hinge on the following two inequalities (which we prove below):

Claim 5.5. *For any* $t\in\mathbb{R}$,

$$\sum_{\mathbf{X}\in\mathbf{\Lambda}}\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot g_t(\mathbf{X}) \quad\le\quad \exp(\pi t^2 d\cdot\|\mathbf{s}\|_2^2)\cdot\boldsymbol{\rho}_{\mathbf{s}}(\mathbf{\Lambda}).$$

Claim 5.6. *For any* $t\in\mathbb{R}$ *and* $r\ge 0$,

$$\sum_{\mathbf{X}\in\mathbf{\Lambda}\setminus Q_r}\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot g_t(\mathbf{X}) \quad\ge\quad \frac{\exp(2\pi t r)}{2^d}\cdot\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{\Lambda}\setminus Q_r).$$

Then we see that

$$\begin{aligned}
\frac{\exp(2\pi t r)}{2^d}\cdot\rho_{\mathbf{s},\mathbf{C}}(\mathbf{\Lambda}\setminus Q_r) \quad&\le\quad \sum_{\mathbf{X}\in\mathbf{\Lambda}\setminus Q_r}\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot g_t(\mathbf{X}) && \text{(Claim 5.6)}\\
&\le\quad \sum_{\mathbf{X}\in\mathbf{\Lambda}}\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot g_t(\mathbf{X}) && (\boldsymbol{\rho},\ g_t\text{ positive})\\
&\le\quad \exp(\pi t^2 d\cdot\|\mathbf{s}\|_2^2)\cdot\boldsymbol{\rho}_{\mathbf{s}}(\mathbf{\Lambda}). && \text{(Claim 5.5)}
\end{aligned}$$

Setting $t=\frac{r}{d\cdot\|\mathbf{s}\|_2^2}$ yields the lemma. □

We now justify the two claims from above.

Proof of Claim 5.5. Before beginning, we caution that the following proof is heavy in notation, though conceptually quite simple. It essentially consists of expanding $\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X})$ and $g_t(\mathbf{X})$ into their constituent parts, aligning the corresponding components, and re-combining. Except where explicitly noted, all of the steps are justified simply by the definition of the notation, or by linearity.

We start by analyzing terms of the following form, which will appear when we expand $g_t(\mathbf{X})$ according to its definition:

$$\begin{aligned}
&\boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X}) \cdot \exp\left(\textstyle\sum_{k\in[d]} 2\pi t \left\langle h_{\mathbf{C}}(\mathbf{X}), \pm\mathbf{u}_k \right\rangle\right) \\
&\quad= \boldsymbol{\rho}_{\mathbf{s},\mathbf{C}}(\mathbf{X}) \cdot \exp\left(\pi \textstyle\sum_{i\in[m]} 2\Big\langle \mathbf{x}_i - \mathbf{c}_i, t\textstyle\sum_{k\in[d]} \pm\mathbf{u}_k \Big\rangle\right) \\
&\quad= \exp\Bigg(-\pi\textstyle\sum_{i\in[m]}\Big((\mathbf{x}_i-\mathbf{c}_i)^2/s_i^2 - 2\Big\langle \mathbf{x}_i-\mathbf{c}_i, t\underbrace{\textstyle\sum_{k\in[d]}\pm\mathbf{u}_k}_{\mathbf{c}'_i}\Big\rangle\Big)\Bigg) \\
&\quad= \exp\Bigg(-\pi\textstyle\sum_{i\in[m]}\Big(\mathbf{x}_i - \underbrace{(\mathbf{c}_i + s_i t\mathbf{c}'_i)}_{\mathbf{c}''_i}\Big)^2/s_i^2 - (s_i t\mathbf{c}'_i)^2\Bigg) && (3) \\
&\quad= \exp\left(\pi t^2 d\cdot\|\mathbf{s}\|_2^2\right)\cdot\exp\left(-\pi\textstyle\sum_{i\in[m]}(\mathbf{x}_i-\mathbf{c}''_i)^2/s_i^2\right) && (4) \\
&\quad= \exp(\pi t^2 d\cdot\|\mathbf{s}\|_2^2)\cdot\rho_{\mathbf{s},\mathbf{C}''}(\mathbf{X})
\end{aligned}$$

Equation (3) is by completing the square. Equation (4) is by $(\mathbf{c}'_i)^2 = \|\mathbf{c}'_i\|_2^2 = d$, regardless of the pattern of ±'s, due to the orthonormality of $\{\mathbf{u}_k\}$.

We now analyze the expression that appears in the statement of Claim 5.5. Expanding the definition of g_t using $\cosh(x) = \frac{1}{2}(e^x + e^{-x})$, we see that the expression $\rho_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot g_t(\mathbf{X})$ contains 2^d terms of the form:

$$\frac{1}{2^d}\cdot\rho_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot\prod_{k\in[d]}\exp\left(\pm 2\pi t\left\langle h_{\mathbf{C}}(\mathbf{X}),\mathbf{u}_k\right\rangle\right) = \frac{1}{2^d}\cdot\rho_{\mathbf{s},\mathbf{C}}(\mathbf{X})\cdot\exp\left(\textstyle\sum_{k\in[d]}2\pi t\left\langle h_{\mathbf{C}}(\mathbf{X}),\pm\mathbf{u}_k\right\rangle\right),$$

which we analyzed above. Summed over all $\mathbf{X}\in\boldsymbol{\Lambda}$, each of these 2^d terms becomes:

$$\frac{\exp(\pi t^2 d\cdot\|\mathbf{s}\|_2^2)}{2^d}\cdot\rho_{\mathbf{s},\mathbf{C}''}(\boldsymbol{\Lambda}) \le \frac{\exp(\pi t^2 d\cdot\|\mathbf{s}\|_2^2)}{2^d}\cdot\rho_{\mathbf{s}}(\boldsymbol{\Lambda}),$$

where the inequality is due to Lemma 5.1. Combining all 2^d terms, Claim 5.5 follows. □

Proof of Claim 5.6. By the definition of g_t and the inequality $\cosh(x)\ge\frac{1}{2}\exp(|x|)$, we have

$$g_t(\mathbf{X}) \ge \frac{1}{2^d}\prod_{k\in[d]}\exp\left(2\pi t\left|\left\langle h_{\mathbf{C}}(\mathbf{X}),\mathbf{u}_k\right\rangle\right|\right) = \frac{1}{2^d}\cdot\exp\left(2\pi t\left\|h_{\mathbf{C}}(\mathbf{X})\right\|_{\mathbf{U}}\right).$$

Then because $\|h_{\mathbf{C}}(\mathbf{X})\|_{\mathbf{U}} \ge r$ for any $\mathbf{X}\in\boldsymbol{\Lambda}\setminus Q_r$, and by positivity $\boldsymbol{\rho}$, the claim follows. □

6 Worst-Case to Average-Case Reductions

In this section we provide a novel analysis, for ℓ_p norms, of two prior worst-case to average-case reductions that use Gaussians. The first, due to Micciancio and Regev [24], shows that solving random systems of modular linear equations is as hard as approximating several worst-case lattice problems in ℓ_2 norm to within $\tilde{O}(n)$ factors. We extend this result to all ℓ_p norms, $p \in [2, \infty]$, essentially maintaining the connection factor of the reduction.

The second reduction, due to Regev [27], shows that decoding random linear codes under a certain noise distribution is as hard as approximating worst-case lattice problems in ℓ_2 norm to within factors as small as $\tilde{O}(n)$, *using a quantum algorithm.* We also extend this result to all ℓ_p norms, $p \in [2, \infty]$, with essentially the same approximation factors.

Both of our extensions rely upon the analysis of discrete Gaussians we developed in Section 5, specifically, Theorem 5.2 and Corollary 5.3. In addition, we remark that a concurrent work of Peikert and Rosen [25] also uses our analysis of discrete Gaussians to obtain *sub-logarithmic* worst-case/average-case connection factors for special classes of algebraic lattices.

6.1 Random Modular Linear Equations

The average-case problem studied in [24] is to find small nonzero solutions to random linear systems of modular equations. This problem goes all the way back to Ajtai's seminal work [3], and can be used as a foundation for collision-resistant cryptographic hash functions. We use the following definition from [24], to which we refer the reader for a full discussion:

Definition 6.1. The *small integer solutions* problem in ℓ_2 norm, denoted SIS, is the following: for an integer q, matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, and real β, find a nonzero integer vector $\mathbf{z} \in \mathbb{Z}^m \setminus \{\mathbf{0}\}$ such that $\mathbf{Az} = 0 \bmod q$ and $\|\mathbf{z}\|_2 \leq \beta$. For functions $q(n)$, $m(n)$, $\beta(n)$, $\mathsf{SIS}_{q,m,\beta}$ is the ensemble over instances $(q(n), \mathbf{A}, \beta(n))$ where $\mathbf{A}$ is a uniformly random $n \times m(n)$ matrix mod $q(n)$.

When $\beta \geq \sqrt{m} \cdot q^{n/m}$, one can show that any SIS instance always has a nonzero solution [24, Lemma 5.2]. We will take β to be $\sqrt{m} \cdot q^{n/m}$ when it is omitted.

We now show that solving the average-case SIS problem is as hard as solving several worst-case lattice problems *in the ℓ_p norm.* The theorem below is an adaptation of the main theorem from [24], which reduces the incremental guaranteed distance decoding (IncGDD) problem to SIS. As we explain below, IncGDD is as hard as several more standard lattice problems.

Theorem 6.2. *For any $p \in [1, \infty)$, there is a constant c_p such that for any $g(n) > 0$, polynomially-bounded $m(n), \beta(n) = \text{poly}(n)$, and $q(n) \geq n \cdot g(n)\beta(n)\sqrt{m(n)}$, solving $\mathsf{SIS}_{q,m,\beta}$ on the average with non-negligible probability is as hard as solving $\mathsf{IncGDD}^{p,\eta_\epsilon}_{\gamma,g}$ in the worst case for sufficiently small $\epsilon(n) = 1/\text{poly}(n)$ and $\gamma(n) = 4c_p n^{1/p} \cdot \beta(n)$.*

For $p = \infty$, there is a constant c such that the same statement holds for $\gamma(n) = 2c \cdot \beta(n)\sqrt{\log n}$.

Proof. For simplicity of notation, we will omit the dependence on n for parameters m, β, etc. Let $\Lambda = \mathcal{L}(\mathbf{B})$, where $\mathbf{B}$ is the lattice basis of the input instance of IncGDD.

The statement of the theorem is virtually identical to one shown by Micciancio and Regev [24, Theorem 5.9]. The only difference is the generalization to ℓ_p norms, and the corresponding change of the approximation factor γ.[4] In addition, the reduction claimed in the theorem is *exactly* the one given in [24]; only the analysis is different. Fortunately, the bulk of the analysis from [24] is insensitive to the choice of p or the value of $\gamma(n)$.

The only part of the proof that differs in our case is the analysis of the length, in ℓ_p norm, of a sum of independent samples from discrete Gaussians. The precise form of the sum in question is $\|(\mathbf{X}-\mathbf{C})\mathbf{z}\|_p$, where the columns $\mathbf{x}_i$ of $\mathbf{X}$ are independent and distributed according to $D_{\Lambda,s,\mathbf{c}_i}$, $\mathbf{c}_i \in \mathbb{R}^n$ are fixed centers, $\mathbf{z} \in \mathbb{Z}^m$ is a fixed vector such that $\|\mathbf{z}\|_2 \leq \beta$, and $s \geq \eta_\epsilon(\Lambda)$. In order to complete the proof from [24], it suffices to show that

$$\|(\mathbf{X}-\mathbf{C})\mathbf{z}\|_p \leq s\gamma/2$$

with some positive constant probability. We do so by a straightforward application of the techniques developed in Section 5.

We first rescale the discrete Gaussian parameters in the following way: let $\Lambda_i = z_i\Lambda$, $\mathbf{s} = s\cdot\mathbf{z}$, and $\mathbf{C}' = \mathbf{C}\mathbf{z}$. Let $\mathbf{\Lambda} = (\Lambda_1, \ldots, \Lambda_m)$. By this rescaling, it is clear that $s_i \geq \eta_\epsilon(\Lambda_i)$, and $\|(\mathbf{X}-\mathbf{C})\mathbf{z}\|_p$ is distributed identically to $\|h_{\mathbf{C}'}(\mathbf{X}')\|_p$, where $\mathbf{X}'$ is sampled from $D_{\mathbf{\Lambda},\mathbf{s},\mathbf{C}'}$ (all of this notation is defined in Section 5.2).

Suppose $p \in [1,\infty)$. Then by Corollary 5.3, we have

$$\mathrm{E}\left[\left\|h_{\mathbf{C}'}(\mathbf{X}')\right\|_p\right] \leq c_p \left\|\mathbf{s}\right\|_2 \cdot n^{1/p} = c_p \cdot s \left\|\mathbf{z}\right\|_2 \cdot n^{1/p} \leq s \cdot c_p\beta \cdot n^{1/p} \leq s\gamma/4.$$

Then by Markov's inequality, $\|h_{\mathbf{C}'}(\mathbf{X}')\|_p \leq s\gamma/2$ with probability at least 1/2.

Now suppose $p = \infty$. By Corollary 5.3,

$$\Pr_{\mathbf{X}'}\left[\|h_{\mathbf{C}'}(\mathbf{X})\|_\infty > s\gamma/2 = c\cdot s\beta\cdot\sqrt{\log n}\right] \leq \Pr_{\mathbf{X}'}\left[\|h_{\mathbf{C}'}(\mathbf{X})\|_\infty > c\cdot\|\mathbf{s}\|_2\cdot\sqrt{\log n}\right] \leq 1/4.$$

This completes the analysis and the proof. □

Connection to other worst-case problems. As shown in [24, Section 5.3], the IncGDD problem (in ℓ_2 norm) is as hard as several more standard lattice problems (also in ℓ_2 norm), via straightforward worst-case to worst-case reductions. One can easily verify that these reductions also apply to any ℓ_p norm, as they only rely on simple properties of norms such as the triangle inequality.

When we instantiate $\beta(n)$ properly so that SIS solutions always exist, Theorem 6.2 and the reductions from [24] imply that, given an SIS oracle, we can find vectors of length $\tilde{O}(n^{1/2+1/p}) \cdot \eta_\epsilon$. By Lemma 2.6, the smoothing parameter η_ϵ is at most $\tilde{O}(n^{1/2-1/p}) \cdot \lambda_n^p$. Combining these two facts, we get an overall connection factor of $\tilde{O}(n)$, which we state precisely in the following corollary:

Corollary 6.3. *For any $p \in [2,\infty)$ and for any $m(n) = \Theta(n\log n)$, there exists some $q(n) = O(n^2\log n)$ such that solving $\mathsf{SIS}_{q,m}$ on the average with non-negligible probability is as hard as solving the following problems for some $\gamma(n) = \Theta(n\log n)$: SIVP^p_γ, GDD^p_γ, and GapCRP^p_γ.*

For $p = \infty$, the same applies for some $\gamma(n) = \Theta(n\log^{1.5} n)$.

[4] We have also slightly departed from [24] in the choice of $\epsilon(n) = 1/\operatorname{poly}(n)$ as an inverse polynomial, rather than negligible function. We observe that it is enough to choose $\epsilon(n)$ to be a small inverse polynomial related to the success probability of the SIS oracle. This is merely an optimization for obtaining the tightest possible reduction.

Proof. By Lemma 2.6, for any $\epsilon(n) = 1/\operatorname{poly}(n)$ there is some $\alpha(n) = \Theta(\sqrt{\log n})$ such that $\eta_\epsilon(\Lambda) \leq \alpha(n) \cdot n^{1/2-1/p} \cdot \lambda_n^p(\Lambda)$ for any n-dimensional lattice Λ. Therefore any algorithm that solves $\mathsf{IncGDD}_{\gamma',g}^{p,\eta_\epsilon}$ for some $\epsilon(n) = 1/\operatorname{poly}(n)$ also solves $\mathsf{IncGDD}_{\gamma,g}^{p,\lambda_n^p}$ for some $\gamma(n) = \Theta(\sqrt{\log n}) \cdot n^{1/2-1/p} \cdot \gamma'(n)$. Applying Theorem 6.2, we get an algorithm for $\mathsf{IncGDD}_{\gamma,g}^{p,\lambda_n}$, for some $\gamma(n) = \Theta(n \log n)$. Using the reductions from [24], we get the desired result. A similar argument applies for $p = \infty$. □

Comment on ℓ_p norms, $1 \leq p < 2$. We point out that Theorem 6.2 applies equally well to all ℓ_p norms for $1 \leq p \leq \infty$. The difficulty in concluding anything meaningful for $p < 2$ arises when we connect the smoothing parameter to λ_n^p. Unlike for $p \geq 2$, we cannot conclude that $\eta_\epsilon \leq \tilde{O}(n^{1/2-1/p}) \cdot \lambda_n^p$. Instead, the best bound we can obtain is $\eta_\epsilon \leq O(\sqrt{\log n}) \cdot \lambda_n^p$, which yields an overall approximation factor of $\gamma(n) = \tilde{O}(n^{1/2+1/p}) = \tilde{O}(n^{3/2})$ for the problems above.

Connection to the shortest vector problem. Just as in [24], the above results do not immediately imply a reduction that solves the shortest vector problem in the worst case. This is because the shortest vector in a lattice may be significantly shorter than the smoothing parameter η_ϵ, but the reduction from Theorem 6.2 may "stop working" once the Gaussian parameter drops below η_ϵ. In [24] a reduction is presented, using the ideas behind the coNP verifier for GapCVP from [2], which solves GapSVP_γ^2 for quasi-linear factors $\gamma(n) = O(n\sqrt{\log n})$. By applying the measure inequalities and their consequences to the techniques from [24], we obtain a reduction that solves GapSVP_γ^p for $\gamma(n) = O(n \log n)$.

Theorem 6.4. *For any $p \in [2, \infty]$, for any $m(n) = \Theta(n \log n)$, there exists* odd *$q(n) = O(n^{2.5} \log n)$ such that solving $\mathsf{SIS}_{q,m}$ on the average is as hard as solving GapSVP_γ^p for some $\gamma(n) = O(n \log n)$.*

Proof sketch. We give a brief sketch of the reduction, deferring the details to the full version.

As explained in [17, 24], GapSVP_γ^p reduces to a variant of GapCVP_γ^p. We reduce this latter problem to SIS as follows: on an instance $(\mathbf{B}, \mathbf{t}, d)$, create instances of SIS using a Gaussian parameter of

$$s = \frac{n^{1/p} \cdot O(\sqrt{\log n})}{\gamma \cdot d}$$

over the *dual* lattice $\mathcal{L}(\mathbf{B})^*$. The outputs of the SIS oracle are combined to yield vectors from the dual lattice, which are used as a witness $\mathbf{W}$ for running the GapCVP^p verifier algorithm $\mathcal{V}$ from Section 4.1 (with appropriate scaling of d). The reduction outputs the negation of $\mathcal{V}$'s output.

For YES instances of the GapCVP_γ^p variant, $\mathbf{t}$ is close to $\mathcal{L}(\mathbf{B})$, so the soundness property of $\mathcal{V}$ guarantees that it always rejects, and the reduction accepts. For NO instances, Lemma 3.5 guarantees that $s \geq \eta_\epsilon(\mathcal{L}(\mathbf{B})^*)$, so the instances of SIS are properly distributed, and the oracle yields sufficiently many samples from the dual lattice $\mathcal{L}(\mathbf{B})^*$, forming $\mathbf{W}$. As shown in [24], the matrix $\mathbf{W}\mathbf{W}^T$ passes $\mathcal{V}$'s eigenvalue test with overwhelming probability. In addition, using Lemma 2.8 and the measure inequalities from Section 3, we can show that the value $f_{\mathbf{W}}(\mathbf{t}) < 1/2$ with positive constant probability. Therefore all of $\mathcal{V}$'s tests are passed, and the reduction rejects with positive constant probability. Standard repetition techniques amplify this to overwhelming probability. □

6.2 Decoding Random Linear Codes

Regev demonstrated that decoding random linear codes mod q under a certain distribution of Gaussian noise is hard, unless there are efficient *quantum* algorithms for approximating the worst-case problems SIVP and GapSVP, in ℓ_2 norm, to within $\tilde{O}(n)$ factors [27]. While the exact nature of the decoding problem will not be important for us, we stress that it can be used as the basis for a semantically-secure public-key cryptosystem. We direct the interested reader to [27] for details.

The essence of Regev's reduction is a quantum strategy which, given a decoding oracle, generates samples from the discrete Gaussians $D_{\Lambda,s}$ for iteratively smaller values of s, all the way down to some $s = q(n) \cdot \eta_\epsilon(\Lambda)$, where $q(n)$ can be as small as $\Theta(\sqrt{n})$.

A straightforward application of our Corollary 5.3 (for the special case of a single discrete Gaussian) demonstrates that samples from $D_{\Lambda,s}$ are short in ℓ_p norm, which allows us to solve SIVP^p. Slightly modifying the reduction from Section 6.1, we can also obtain a reduction from GapSVP^p. This results in an adaptation of the main theorem from [27] to any ℓ_p norm, $p \in [2, \infty]$:

Theorem 6.5 (Informal). *Let $\alpha = \alpha(n) \in (0, 1)$ be a real number, and $q = q(n) \geq 2\sqrt{n}/\alpha(n)$ be an integer. For any $p \in [2, \infty]$, if there exists a (possibly quantum) polynomial-time algorithm that solves the decoding problem mod q, then there exist quantum algorithms that solve SIVP^p_γ and GapSVP^p_γ in the worst case for some $\gamma(n) = \tilde{O}(n/\alpha)$.*

7 Acknowledgements

We gratefully thank Alon Rosen for much encouragement and helpful advice in preparing this paper.

References

[1] D. Aharonov and O. Regev. A lattice problem in quantum NP. In *FOCS*, pages 210–219. IEEE Computer Society, 2003.

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

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

[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] M. Alekhnovich, S. Khot, G. Kindler, and N. K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing. In *FOCS*, pages 216–225. IEEE Computer Society, 2005.

[7] L. Babai. On lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.

[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] S. Ben-David, B. Chor, O. Goldreich, and M. Luby. On the theory of average case complexity. *J. Comput. Syst. Sci.*, 44(2):193–219, 1992.

[11] J. Blömer and J.-P. Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *STOC*, pages 711–720, 1999.

[12] I. Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. *Theor. Comput. Sci.*, 285(1):55–71, 2002.

[13] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003.

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

[15] O. Goldreich. Note available at `http://www.wisdom.weizmann.ac.il/~oded/p_lp.html`.

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

[17] 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.

[18] V. Guruswami, D. Micciancio, and O. Regev. The complexity of the covering radius problem on lattices and codes. In *IEEE Conference on Computational Complexity*, pages 161–173. IEEE Computer Society, 2004.

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

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

[21] L. A. Levin. Average case complete problems. *SIAM J. Comput.*, 15(1):285–286, 1986.

[22] 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.

[23] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, 2002.

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

[25] C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In submission., 2006.

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

[27] 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.

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

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