Limits on the Hardness of Lattice Problems in ℓ_p Norms*

Chris Peikert†

Abstract

Several recent papers have established limits on the computational difficulty of lattice problems, focusing primarily on the ℓ_2 (Euclidean) norm. We demonstrate close analogues of these results in ℓ_p norms, for every $2 < p \leq \infty$. In particular, for lattices of dimension n:

- Approximating the closest vector problem, the shortest vector problem, and other related problems to within $O(\sqrt{n})$ factors (or $O(\sqrt{n \log n})$ factors, for $p = \infty$) is in coNP.
- Approximating the closest vector and bounded distance decoding problems *with preprocessing* to within $O(\sqrt{n})$ factors can be accomplished in deterministic polynomial time.
- Approximating several problems (such as the shortest independent vectors problem) to within $\tilde{O}(n)$ factors in the *worst case* reduces to solving the *average-case* problems defined in prior works (Ajtai, STOC 1996; Micciancio and Regev, SIAM J. on Computing 2007; Regev, STOC 2005).

Our results improve prior approximation factors for ℓ_p norms by up to $\sqrt{n}$ factors. Taken all together, they complement recent reductions from the ℓ_2 norm to ℓ_p norms (Regev and Rosen, STOC 2006), and provide some evidence that lattice problems in ℓ_p norms (for $p > 2$) may not be substantially harder than they are in the ℓ_2 norm.

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

1 Introduction

An n-dimensional *lattice* $\Lambda \subset \mathbb{R}^n$ is a periodic "grid" of points generated by all integer linear combinations of n linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^n$, which form what is called a *basis* of Λ. Over the past two centuries, lattices have emerged as fundamental objects in many areas of mathematics. More recently in computer science, lattices have been at the center of several celebrated results in algorithms, complexity theory, and cryptography (e.g., [LLL82, Ajt04]).

Two of the central computational problems on lattices are the *shortest vector* problem SVP and the *closest vector* problem CVP. The goal of SVP, given an arbitrary basis of a lattice, is to find a (nonzero) lattice point that is closest to the origin. The goal of CVP, given an arbitrary basis of

*An edited version of this paper appears in Computational Complexity 17(2):300-351 (2008). A preliminary version appeared in the 22nd Annual IEEE Conference on Computational Complexity (CCC 2007). This is the author's version.

†SRI International, 333 Ravenswood Ave, Menlo Park, CA, 94025. Email: cpeikert@alum.mit.edu. This material is based upon work supported by the National Science Foundation under Grant CNS-0716786. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

a lattice and a target point $\mathbf{v} \in \mathbb{R}^n$, is to find a lattice point closest to $\mathbf{v}$. In these problems and others, the distance is measured relative to some fixed norm on $\mathbb{R}^n$. Generally this is the ℓ_2 (i.e., Euclidean) norm, but it is also common to use the ℓ_p norm for some $1 \leq p \leq \infty$.[1]

As with most optimization problems, it is interesting to consider *approximation* versions of lattice problems. For SVP (respectively, 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 $\gamma \geq 1$. Generally, γ is taken to be some function $\gamma(n)$ of the dimension of the lattice.

Known polynomial-time algorithms for approximating SVP and CVP (in the ℓ_2 norm), such as the LLL algorithm, achieve approximation factors $\gamma(n)$ that are only mildly better than *exponential* in n [LLL82, Bab86, Sch87, AKS01]. To approximate SVP or CVP (in any ℓ_p norm) to within even polynomial factors, known algorithms require time and space that are exponential in n [AKS01, AKS02, BN07]. (We add that these algorithms can achieve any *constant* approximation factor $\gamma(n) = 1 + \epsilon$, and that the algorithm of [AKS01] can even solve SVP in the ℓ_2 norm *exactly*.)

There are a number of other seemingly difficult computational problems relating to geometric properties of lattices; we informally describe a few of them here.

- The goal of the *shortest independent vectors problem* SIVP is to find (given an arbitrary lattice basis) a set of n linearly independent lattice vectors, where the length of the longest such vector is minimized (more generally, is within a $\gamma(n)$ factor of optimal).
- The goal of the *covering radius problem* CRP is to compute (to within a $\gamma(n)$ factor, given a basis) the covering radius of the lattice, which is defined as the maximum distance from a point $\mathbf{t}$ to the lattice, taken over all $\mathbf{t} \in \mathbb{R}^n$.
- Motivated by coding theory, the closest vector problem *with preprocessing* CVPP is a version of CVP in which the lattice is *fixed*, and arbitrary succinct advice about the lattice (generated by an unbounded preprocessing phase) may be used to find a lattice vector close to a given target point.
- In a variant of CVPP called *bounded distance decoding* BDD (with preprocessing), the distance from the target point to the lattice is guaranteed to be within a certain factor of the minimum distance of the lattice (i.e., the length of its shortest nonzero vector).

Hardness. The apparent difficulty of lattice problems is reinforced by several hardness results. As is standard in hardness of approximation, these results are expressed in terms of decision problems satisfying a *promise*, or "gap" (by convention, the problem names start with Gap). The goal is to decide whether the relevant quantity is bounded from above by (say) 1, or from below by γ. For example, the goal of GapSVP is to determine whether the length of the shortest nonzero vector in a lattice (given by an arbitrary basis) is at most 1 or greater than γ (otherwise, any answer is acceptable). Hardness for a gap problem generally implies analogous hardness for the corresponding optimization and search problems, by standard reductions. Below, we briefly describe the state of the art for hardness of lattice problems, which applies for all ℓ_p norms (in many cases, via a general result of [RR06]) unless stated otherwise.

To summarize, some form of hardness is known for all the problems we have introduced. The approximation factors range from specific constants to as large as $n^{c/\log\log n}$ for some constant $c > 0$.

[1] For $1 \leq p < \infty$, 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 $\|\mathbf{x}\|_\infty = \max_i |x_i|$.

However, no hardness is known for any polynomial approximation factor n^ϵ, under any standard complexity assumption.

For GapSVP and $p < \infty$, there is no (randomized) polynomial-time algorithm for any constant approximation factor unless NP $\subseteq$ RP, nor for any $2^{(\log n)^{1-\epsilon}}$ factor unless NP $\subseteq$ RTIME$(2^{\mathrm{poly}(\log n)})$, nor for some $n^{c/\log\log n}$ factor (where $c > 0$) unless NP $\subseteq$ RSUBEXP $= \cap_{\delta>0}$RTIME(2^{n^δ}) [Ajt98, CN99, Mic00, Kho05, HR07]. For $p = \infty$, GapSVP is NP-hard to approximate to within a factor $n^{c/\log\log n}$ for some constant $c > 0$ [vEB81, Din02].

GapCVP is NP-hard to approximate to within $n^{c/\log\log n}$ factors for some constant $c > 0$ [ABSS97, DKRS03, Din02]. For GapCVPP with $2 \leq p < \infty$, there is no (randomized) polynomial-time algorithm for any $(\log n)^{1/2-\epsilon}$ factor unless NP $\subseteq$ RTIME$(2^{\mathrm{poly}(\log n)})$; for $p = \infty$, there is no such algorithm for any constant factor unless NP $\subseteq$ RP, nor for any $(\log n)^{1/2-\epsilon}$ factor unless NP $\subseteq$ RTIME$(2^{\mathrm{poly}(\log n)})$ [FM04, Reg04a, AKKV05, RR06]. For GapBDD, there is no polynomial-time algorithm for any factor $\sqrt{2} - \epsilon$ unless NP $\subseteq$ P/poly [LLM06].

GapSIVP is NP-hard to approximate to within any constant factor, and no polynomial-time algorithm exists for any $2^{(\log n)^{1-\epsilon}}$ factor unless NP $\subseteq$ DTIME$(2^{\mathrm{poly}(\log n)})$. For all sufficiently large $p \leq \infty$, there is a constant $c_p > 1$ such that GapCRP in the ℓ_p norm is Π_2-hard to approximate to within any factor less than c_p (for $p = \infty$, the constant $c_\infty = 1.5$) [HR06].

Limits on hardness. Given the difficulty of designing efficient algorithms for even moderately subexponential approximation factors, one might hope to significantly increase the factors in the hardness results above. However, there seem to be strict limits to any such improvements.

One general approach for limiting hardness is to show, for some approximation factor $\gamma(n)$, that a certain problem is in some complexity class that is not believed to contain NP, such as coNP or coAM. Putting aside some subtleties, this in turn implies that the problem (for the particular factor $\gamma(n)$) is not NP-hard, assuming that the polynomial-time hierarchy does not collapse.

Lagarias, Lenstra, and Schnorr [LLS90] showed that for $\gamma(n) = n^{1.5}$, GapCVP and GapSVP are in coNP. In other words, there is a succinct witness proving that a given point is *far* from a given lattice, and a witness proving that a given lattice has *no* short nonzero vectors. The factor was improved to $\gamma(n) = n$ by Banaszczyk [Ban93]. Goldreich and Goldwasser showed that GapCVP and GapSVP are in coAM for some $\gamma(n) = O(\sqrt{n/\log n})$ [GG00]. Recently, Aharonov and Regev improved the containment to coNP, for a slightly relaxed factor $\gamma(n) = O(\sqrt{n})$ [AR05]. Building upon these most recent works, other problems such as GapSIVP and GapCRP have been placed in coNP for some $\gamma(n) = O(\sqrt{n})$ factor, and in coAM for some $\gamma(n) = O(\sqrt{n/\log n})$ factor [GMR05], and the problems GapCVPP and BDD have been shown to be computable in polynomial time (not including the unlimited preprocessing stage) for $\gamma(n) = O(\sqrt{n/\log n})$ [AR05, LLM06].

One of the most remarkable features of lattice problems is their *worst-case/average-case reducibility*, first demonstrated by Ajtai [Ajt04] and studied extensively since then (e.g., by [CN97, Mic04, MR07, Reg05]). Worst-case to average-case reductions are typically taken as evidence for the *hardness* of the average-case problem, which often has applications in cryptography. At the same time, though, these reductions also *limit the hardness* of the underlying worst-case problem, by showing that it is no harder than the average-case problem (which is in, say, distributional-NP [Lev86, BDCGL92]). The state of the art in this area is represented by the works of Micciancio and Regev [MR07] and Regev [Reg05], who demonstrated reductions from worst-case lattice problems in the ℓ_2 norm for almost-linear $\gamma(n) = \tilde{O}(n)$ factors to certain average-case problems. Interestingly, the latter result of Regev is a *quantum* reduction, which is non-trivial because quantum computing

is not known to confer any advantage over classical algorithms for lattice problems.

A final intriguing limit on hardness comes from a recent paper of Regev and Rosen [RR06]. By applying embeddings from the ℓ_2 norm to other ℓ_p norms, they showed (essentially) that lattice problems in the ℓ_2 norm are *no harder* than they are in the ℓ_p norm, for any approximation factor $\gamma(n)$ and any $p \in [1, \infty]$.

We emphasize that all of these results have had the primary effect of limiting the hardness of lattice problems *in the ℓ_2 norm*. Using standard relations between norms, one can trivially obtain limits for other ℓ_p norms, but the approximations suffer by up to $\sqrt{n}$ factors. For example, the approximation factors become $O(n^{1/2+|1/2-1/p|}) = O(n)$ for the problems in coNP, and become $\tilde{O}(n^{1+|1/2-1/p|}) = \tilde{O}(n^{1.5})$ for the worst-case/average-case reductions.

Summary. Focusing on the relationship between norms, the landscape looks as follows: in certain cases (such as for GapSVP and GapCRP), the known hardness results for large values of p are stronger than those for, say, $p = 2$. Problems in ℓ_p norms are essentially *no easier* than those in the ℓ_2 norm, and furthermore, the known limits on their hardness are *weaker* by factors as large as $\sqrt{n}$. Therefore, most of the evidence indicates that lattice problems in ℓ_p norms could be *strictly harder* than those in the ℓ_2 norm — but is this actually true? This is the main question motivating our work.

1.1 Our Results

We show (perhaps surprisingly, given the state of the art) that many known limits on hardness for the ℓ_2 norm carry over to the ℓ_p norm for any $p > 2$, for essentially the *same* asymptotic approximation factors. Specifically, for any $2 < p \leq \infty$ we show that:

- For certain $O(\sqrt{n})$ approximation factors (or $O(\sqrt{n \log n})$ factors for $p = \infty$), GapCVP in the ℓ_p norm is in coNP. By known relations among lattice problems, it also follows that GapSVP, GapSIVP, and GapCRP are in coNP for the same asymptotic factors.
- For certain $O(\sqrt{n})$ approximation factors, GapCVPP and the search variant BDD in the ℓ_p norm can be solved in deterministic polynomial time (with unlimited preprocessing).
- For certain $\tilde{O}(n)$ approximation factors, the worst-case problems SIVP and GapSVP (among others) in the ℓ_p norm reduce to the average-case problem first defined in [Ajt04]. The same holds true for the average-case problem defined in [Reg05], under an efficient *quantum* reduction.

Each of these results improves upon the prior known approximation factors for ℓ_p norms by up to a $\sqrt{n}$ factor, and matches (up to $O(\log n)$ factors or better) the current state of the art for the ℓ_2 norm. We remark that the factors hidden by the O-notation above do depend very mildly on the choice of norm and the specific lattice problem. For the problems in coNP (with p finite), the $O(\sqrt{n})$ expression hides a constant factor proportional to $\sqrt{p}$. For the problems with preprocessing, the approximation factor does not depend on p, but our bound is an $O(\sqrt{\log n})$ factor looser than those known for the ℓ_2 norm [AR05, LLM06]. For the worst-case/average-case reductions, our approximation factors are looser than those achieved in [MR07, Reg05] by any $\omega(\sqrt{\log n})$ factor.

Our results also have implications for cryptography. Until now, the hardness of lattice-based cryptographic primitives has always been based on worst-case problems *in the ℓ_2 norm*. Because

lattice problems are essentially *easiest* in the ℓ_2 norm (for any given approximation factor), the security of the resulting primitives has thus far been 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, $2 \leq p \leq \infty$. Viewed another way, our results imply that an adversary capable of breaking the cryptographic primitive is also capable of solving lattice problems in *all* such norms.

One of our main technical contributions is a very general analysis of Gaussian distributions over lattices, which we hope will be of independent interest and utility elsewhere. Indeed, our analysis has also been applied in work by Peikert and Rosen [PR07] to obtain very tight worst-case/average-case reductions for special classes of algebraic lattices.

Finally, we remark that our proofs can easily be adapted to essentially arbitrary norms $\|\cdot\|$. Generally speaking, our proofs exploit *upper bounds* in the ℓ_2 norm in one case, and *lower bounds* in the ℓ_p norm in the opposing case. Therefore it typically suffices to know some $R > 0$ such that for all $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|$ (without loss of generality, by scaling) and $\frac{1}{R}\|\mathbf{x}\| \leq \|\mathbf{x}\|_\infty$; the value R then contributes to the final approximation factor. Note that by the relationship between the ℓ_2 and ℓ_∞ norms, we always have $R \geq \sqrt{n}$.

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 [RR06], reducing from the ℓ_p norm to the ℓ_2 norm appears to be much more challenging.

We 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 the *exact same* algorithms and reductions that were initially designed for the ℓ_2 norm! Our results follow by a novel analysis of these constructions for ℓ_p norms. We rely on results and techniques of Banaszczyk that were initially developed to prove transference theorems for lattices, first for the ℓ_2 norm [Ban93], and later for norms defined by more general convex bodies, including ℓ_p norms [Ban95]. Ideas from the former paper have stimulated many recent advances in the understanding of lattices in computer science. To the best of our knowledge, this is the first time that techniques from the latter paper have been applied in computational complexity.

As an illustration, let us summarize the proof that GapCVP in the ℓ_p norm is in coNP for some $O(\sqrt{n})$ approximation factor (where p is finite for simplicity). We apply certain *measure inequalities* from [Ban95] to the framework laid out by Aharonov and Regev [AR05]. Their main tool is a function $f : \mathbb{R}^n \to [0, 1]$ that is indicates whether its argument is close to, or far from, a given lattice (in the ℓ_2 norm). We show that the same function f also works in all ℓ_p norms, for $p \geq 2$. Specifically,

- For points $\mathbf{x}$ whose ℓ_p distance to the lattice is at least $c_p n^{1/p}$ (for some constant c_p depending only on p), the measure inequalities of [Ban95] guarantee that $f(\mathbf{x}) \leq \frac{1}{4}$.
- For points $\mathbf{x}$ whose ℓ_p distance to the lattice is at most $\frac{1}{100} \cdot n^{1/p-1/2}$, standard properties of norms combined with results from [AR05] guarantee that $f(\mathbf{x}) \geq \frac{1}{2}$.

These two facts are the essence of the $O(\sqrt{n})$ gap factor in the resulting coNP proof system. (Similar facts hold for $p = \infty$, yielding an $O(\sqrt{n \log n})$ gap.)

For our new analysis of prior worst-case to average-case reductions, we derive new facts about the *discrete Gaussian* probability distributions over lattices that emerge in the analysis of those reductions. Specifically, we show that in many important respects, these discrete Gaussians behave almost exactly like continuous Gaussians. In particular, the expected ℓ_p norm of a point sampled from an n-dimensional discrete Gaussian is proportional to $n^{1/p}$, and the sum of m independent discrete Gaussians behaves like a single discrete Gaussian whose standard deviation is a $\sqrt{m}$ factor larger. Our results extend prior analyses by Micciancio and Regev [MR07] and Lyubashevsky and Micciancio [LM06], while providing more modular and tractable proofs.

1.3 Open Questions

The case of $p < 2$. Our work does not say much about lattice problems in ℓ_p norms for $1 \leq p < 2$, due to their relationship with the ℓ_2 norm. We are unable to conclude anything other than what is trivially implied by basic relations among norms, i.e., problems in coNP for $O(n^{1/p})$ factors, and worst-case to average-case reductions with $\tilde{O}(n^{1/2+1/p})$ connection factors.

One way of approaching ℓ_p norms for $p < 2$ might be via *duality*, which defines a natural correspondence between not only pairs of lattices, but also pairs of norms. In particular, for $1 \leq p \leq 2 \leq q \leq \infty$ such that $1/p + 1/q = 1$, the ℓ_p norm and ℓ_q norm are dual to each other. It may be that lattice problems in the ℓ_p norm could be related to problems in the ℓ_q norm in this way.

We point out that any results going below the $n^{1/p}$ barrier (even for *just one* ℓ_p norm) would imply analogous non-trivial results for problems on *linear codes* over binary or ternary alphabets, such as the nearest codeword problem and the minimum distance problem. This follows from a standard transformation from codes to lattices (see, e.g., [FM04]), which converts a Hamming distance of d to an ℓ_p distance of $d^{1/p}$. Therefore, if CVP in some ℓ_p norm is in coNP for $\gamma(n) = n^{(1-\epsilon)/p}$, then the nearest codeword problem is in coNP for a sublinear approximation factor $n^{1-\epsilon}$. As far as we are aware, nothing of the sort is known about codes, and this may explain the difficulty in obtaining better bounds for lattice problems when $p < 2$. (For the nearest codeword problem, there are polynomial-time algorithms achieving $\Omega(n)$ approximation factors without randomization, and $\Omega(n/\log n)$ factors with randomization [BK02].)

coNP versus coAM. Another interesting question is whether our results for coNP can be tightened by a $\sqrt{\log n}$ factor, by relaxing the containments to coAM. This question is motivated by the current state of affairs for the ℓ_2 norm, where CVP is known to be in coNP only for some $\gamma(n) = O(\sqrt{n})$, but is known to be in coAM for some $\gamma(n) = O(\sqrt{n/\log n})$. In ℓ_p norms, however, our techniques do not seem to yield such an improvement. We explain below.

Recall that the main tool in [AR05] is a function f indicating whether a given point is close to, or far from, some given lattice. In the ℓ_2 norm, the measure inequality from [Ban93] gives an *exponentially small* bound on f in the "far" case. That is, for points $\mathbf{x}$ at a distance $\geq \sqrt{n}$ from the lattice, we have $f(\mathbf{x}) < 2^{-n}$. As pointed out in [AR05], the full strength of this bound is not needed for the coNP proof system; a small enough constant bound suffices. In contrast, the coAM protocol of [GG00] has soundness error as large as $1 - \frac{1}{\mathrm{poly}(n)}$, and therefore needs an inverse polynomial bound on $f(\mathbf{x})$ for completeness.

For ℓ_p norms, the measure inequalities of [Ban95] provide only a *constant* upper bound on $f(\mathbf{x})$ when $\mathbf{x}$ is at distance $\geq n^{1/p}$ from the lattice. As we have already explaiend, this constant bound is sufficient for the coNP proof system. However, it does not appear strong enough to yield a coAM protocol for $O(\sqrt{n/\log n})$ factors, due to the large soundness error.

We suspect that the way to resolve this issue is by improving the coNP proof system for the ℓ_2 norm by a $\sqrt{\log n}$ factor, which was left as an open problem in [AR05].[2] The coNP proof systems would then essentially subsume the known coAM protocols, for all values of p.

Equivalence among ℓ_p norms? A final challenging question is whether there are efficient approximation-preserving reductions from lattice problems in the ℓ_p norm to corresponding problems in the ℓ_2 norm (perhaps only for $p \geq 2$). Together with [RR06], this would imply that all ℓ_p norms are polynomially equivalent for *any* factor $\gamma(n)$.

1.4 Reader's Guide and Warning

Section 2 contains basic notation and concepts that are needed throughout the rest of the paper, and may be safely skimmed by the reader who is familiar with the recent literature. The remainder of the paper is conceptually divided into two independent parts, each containing one section devoted to analysis and one section devoted to complexity-theoretic applications.

Section 3 introduces some *measure inequalities* of Banaszczyk and their immediate implications. These facts provide a basic starting point for understanding how Gaussian measures relate to ℓ_p norms, and we encourage the reader to start here. Section 4 then applies the inequalities to the framework of [AR05], resulting in coNP proof systems for several lattice problems, and some other results relating to problems with preprocessing.

Section 5 develops a new analysis of discrete Gaussian distributions. It is completely self-contained and assumes no prior knowledge, though we still recommend absorbing Section 3 beforehand for intuition. Section 6 then applies this analysis by extending prior worst-case/average-case reductions to ℓ_p norms. This section is quite technical and requires familiarity with prior works.

Finally, a few words of warning: this work extends and generalizes results from a number of recent papers [AR05, LLM06, MR07, Reg05]. Wherever possible, we have attempted to present enough context to make the presentation as self-contained as possible, and detailed knowledge of the prior works unnecessary. In some cases, however, understanding and verifying our claims requires more familiarity with the details of the prior works. Moreover, at times our proofs need to rely upon facts that are established only implicitly in those works (for example, within lengthy proofs of other claims). In these cases, we have made our best attempt to encapsulate the main facts we need, and to provide some guidance to the reader who is interested in verifying the claims against the contents of the original papers.

2 Preliminaries

2.1 Notation

The real numbers are denoted by $\mathbb{R}$ and the integers by $\mathbb{Z}$. For any positive integer n, $[n]$ denotes $\{1, \ldots, n\}$. The function log is always taken to be the natural logarithm. We extend any function $f(\cdot)$ to a countable set A in the following way: $f(A) = \sum_{x \in A} f(x)$. We write $\mathrm{poly}(n)$ for some unspecified polynomial function in n. We write $\omega(f(n))$ to denote the set of functions (or a particular function in that set) growing faster than $c \cdot f(n)$ for any $c > 0$. A positive function $\epsilon(\cdot)$ is *negligible* in its parameter if it decreases faster than the inverse of any polynomial, i.e., if $\epsilon(n) = n^{-\omega(1)}$.

[2] For $p = \infty$, such a result would presumably also improve our coNP system to work for $\gamma(n) = O(\sqrt{n})$, but would have no further effect for finite p.

Vectors are written using bold lower-case letters, e.g., $\mathbf{x}$. For a vector $\mathbf{x}$, the ith component of $\mathbf{x}$ is denoted x_i. Matrices are written using bold capital letters, e.g., $\mathbf{X}$. The ith column vector of $\mathbf{X}$ is denoted $\mathbf{x}_i$. We often use matrix notation to denote a set of vectors, i.e., $\mathbf{S}$ also represents the set of its column vectors. We write $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \ldots)$ to denote the linear space spanned by its arguments. For a set $S \subseteq \mathbb{R}^n$, $\mathbf{v} \in \mathbb{R}^n$, and $c \in \mathbb{R}$, we let $S + \mathbf{x} = \{\mathbf{y} + \mathbf{x} \ : \ \mathbf{y} \in S\}$ and $cS = \{c\mathbf{y} \ : \ \mathbf{y} \in S\}$.

A norm $\|\cdot\|$ is a nonnegative real-valued function on $\mathbb{R}^n$ that satisfies the following axioms: $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{0}$, $\|c\mathbf{x}\| = |c| \, \|\mathbf{x}\|$ for any $c \in \mathbb{R}$, and $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ (the triangle inequality). The associated metric is $\text{dist}(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$. For $p \in [1, \infty)$, the ℓ_p norm on $\mathbb{R}^n$, denoted $\|\cdot\|_p$, is defined as

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

For $p = \infty$, the ℓ_∞ norm $\|\cdot\|_\infty$ is defined as $\|\mathbf{x}\|_\infty = \max_{i \in [n]} |x_i|$. For ease of notation, when $p \in [1, \infty]$ represents an ℓ_p norm we say that $1/p = 0$ for $p = \infty$ and $1/p = 0$ for $p = \infty$. As a special case of Hölder's inequality, for any $\mathbf{x} \in \mathbb{R}^n$ and any $p \in [2, \infty]$, we have

$$n^{1/p - 1/2} \|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_p \leq \|\mathbf{x}\|_2 ,$$

whereas for any $p \in [1, 2]$, we have

$$\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_p \leq n^{1/p - 1/2} \|\mathbf{x}\|_2 .$$

In the following, fix some arbitrary norm $\|\cdot\|$ on $\mathbb{R}^n$. By convention, we say that the norm of a matrix is the norm of its longest column: $\|\mathbf{X}\| = \max_i \|\mathbf{x}_i\|$. For any $\mathbf{t} \in \mathbb{R}^n$ and set $V \subseteq \mathbb{R}^n$, the distance from $\mathbf{t}$ to V is $\text{dist}(\mathbf{t}, V) = \inf_{\mathbf{v} \in V} \text{dist}(\mathbf{t}, \mathbf{v})$. The *open* unit ball is denoted $\mathcal{B}_n = \{\mathbf{x} \in \mathbb{R}^n \ : \ \|\mathbf{x}\| < 1\}$, and the *closed* unit ball by $\mathcal{C}_n = \overline{\mathcal{B}_n} = \{\mathbf{x} \in \mathbb{R}^n \ : \ \|\mathbf{x}\| \leq 1\}$. We often affix a superscript $p \in [1, \infty]$ to these expressions to indicate that the quantity is defined using the ℓ_p norm, e.g., dist^p or $\mathcal{B}_n^\infty$.

The Euler Gamma function for real $z > 0$ is defined as

$$\Gamma(z) = 2 \int_{r=0}^{\infty} r^{2z-1} e^{-r^2} \, dr.$$

The Gamma function satisfies the recursive formula $\Gamma(z + 1) = z\Gamma(z)$, and $\Gamma(z + 1) = z!$ for all nonnegative integers z.

2.2 Lattices

For a matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$ whose columns $\mathbf{b}_1, \ldots, \mathbf{b}_n$ are linearly independent, the n-dimensional *lattice*[3] Λ generated by the *basis* $\mathbf{B}$ is

$$\Lambda = \mathcal{L}(\mathbf{B}) = \left\{ \mathbf{Bc} = \sum\nolimits_{i \in [n]} c_i \cdot \mathbf{b}_i \ : \ \mathbf{c} \in \mathbb{Z}^n \right\}.$$

The *fundamental parallelepiped* of $\mathbf{B}$ is the half-open set $\mathcal{P}(\mathbf{B}) = \{\sum_i c_i \mathbf{b}_i \ : \ 0 \leq c_i < 1, i \in [n]\}$. The *dual lattice* of Λ, denoted Λ^*, is defined to be $\Lambda^* = \{\mathbf{x} \in \mathbb{R}^n \ : \ \forall \, \mathbf{v} \in \Lambda, \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbb{Z}\}$. From the symmetry of this definition, it is easy to see that $(\Lambda^*)^* = \Lambda$.

[3] Technically, this is the definition of a *full-rank* lattice, which is all we are concerned with in this work.

In the following definitions, fix some norm $\|\cdot\|$ on $\mathbb{R}^n$, and recall that $\mathcal{C}_n \subset \mathbb{R}^n$ is the closed unit ball under that norm. The *minimum distance* of a lattice Λ, denoted $\lambda_1(\Lambda)$, is the length of its shortest nonzero element: $\lambda_1(\Lambda) = \min_{0 \neq \mathbf{x} \in \Lambda} \|\mathbf{x}\|$. More generally, the *ith successive minimum* $\lambda_i(\Lambda)$ is the smallest radius r such that the (closed) ball $r\mathcal{C}_n$ contains i linearly independent vectors in Λ; formally, $\lambda_i(\Lambda) = \min\{r \in \mathbb{R} : \dim \operatorname{span}(\Lambda \cap r\mathcal{C}_n) \geq i\}$. The *covering radius* of Λ, denoted $\mu(\Lambda)$, is the smallest radius r such that closed balls $r\mathcal{C}_n$ centered at every point of Λ cover all of $\mathbb{R}^n$; formally, $\mu(\Lambda) = \max_{\mathbf{x} \in \mathbb{R}^n} \operatorname{dist}(\mathbf{x}, \Lambda)$. As above, for each of these quantities we often affix a superscript (p) to indicate that the quantity is measured in the ℓ_p norm, e.g., $\lambda_i^{(2)}$ and $\mu^{(p)}$.

2.3 Problems on Lattices

Here we define some standard worst-case problems on lattices. See [MG02, MR07] for further motivation and discussion of these problems. We define approximation problems parameterized by a positive function $\gamma = \gamma(n)$ of the lattice dimension n. The problems are defined relative to some (implicit) norm $\|\cdot\|$ on $\mathbb{R}^n$. We attach a superscript $p \in [1, \infty]$ to the name of the problem to indicate that this norm is the ℓ_p norm, e.g., GapSVP_γ^p or $\mathsf{GDD}_\gamma^{p,\phi}$.

We first define some promise (or "gap") problems, where the goal is to decide whether the input belongs to the YES set or the NO set (these two sets are disjoint, but not necessarily exhaustive; when the input belongs to neither set, any output is acceptable). In the complement of a promise problem, the YES and NO sets are merely swapped.

Definition 2.1 (Shortest Vector Problem)**.** An input to GapSVP_γ is a basis $\mathbf{B}$ of an n-dimensional lattice. It is a YES instance if $\lambda_1(\mathcal{L}(\mathbf{B})) \leq 1$, and is a NO instance if $\lambda_1(\mathcal{L}(\mathbf{B})) > \gamma(n)$.

Definition 2.2 (Closest Vector Problem)**.** An input to GapCVP_γ is a pair $(\mathbf{B}, \mathbf{v})$ where $\mathbf{B}$ is a basis of an n-dimensional lattice and $\mathbf{v} \in \mathbb{R}^n$. It is a YES instance if $\operatorname{dist}(\mathbf{v}, \mathcal{L}(\mathbf{B})) \leq 1$, and is a NO instance if $\operatorname{dist}(\mathbf{v}, \mathcal{L}(\mathbf{B})) > \gamma(n)$.

Formally (and as in [AR05]), we assume that the target vector $\mathbf{v}$ is specified relative to the basis $\mathbf{B}$ using some $\ell(n) = \operatorname{poly}(n)$ bits of precision, i.e., $\mathbf{v}$ is specified as $\mathbf{v} = \sum_{i \in [n]} a_i \mathbf{b}_i$, where the coefficients a_i are specified using at most $\ell(n)$ bits each. In particular, $\mathbf{v} \in \mathcal{L}(\mathbf{B})/2^{\ell(n)}$.

Definition 2.3 (Covering Radius Problem)**.** An input to GapCRP_γ is a basis $\mathbf{B}$ of an n-dimensional lattice. It is a YES instance if $\mu(\mathcal{L}(\mathbf{B})) \leq 1$ and is a NO instance if $\mu(\mathcal{L}(\mathbf{B})) > \gamma(n)$.

Note that the choice of the quantities 1 and γ in the above problems is arbitrary; by scaling the input instance, they can be replaced by β and $\beta \cdot \gamma$ (respectively) for any $\beta > 0$ without changing the problem.

We now define some lattice problems in their search versions. In these problems, the parameter ϕ represents any desired function from n-dimensional lattices to the positive reals, e.g., the nth successive minimum λ_n or the covering radius μ.

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

In the promise variant $\mathsf{GapGIVP}_\gamma^\phi$, the input is a YES instance if $\phi(\mathcal{L}(\mathbf{B})) \leq 1$ and a NO instance if $\phi(\mathcal{L}(\mathbf{B})) > \gamma$.

The special case $\phi = \lambda_n$ defines the shortest independent vectors problem SIVP_γ and its promise variant $\mathsf{GapSIVP}_\gamma$.

Definition 2.5 (Guaranteed Distance Decoding Problem)**.** An input to $\mathsf{GDD}^{p,\phi}_{\gamma}$ is a pair $(\mathbf{B}, \mathbf{t})$ where $\mathbf{B}$ is a basis of an n-dimensional lattice and $\mathbf{t} \in \mathbb{R}^n$. The goal is to output a lattice point $\mathbf{x} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{t} - \mathbf{x}\| \leq \gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$.

Various other problems on lattices are defined within the paper, closer to where they are needed.

2.4 Gaussian Measures

Our review of Gaussian measures over lattices follows the development by prior works [Reg04b, AR05, MR07]. For any $s > 0$ define the Gaussian function centered at $\mathbf{c}$ with parameter s as:

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

The subscripts s and $\mathbf{c}$ are taken to be 1 and $\mathbf{0}$ (respectively) when omitted.

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

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

(As above, we may omit the parameters s or $\mathbf{c}$.) Note that the denominator in the above expression is always finite (see, e.g., [AR05, Claim 2.4]), so the probability distribution is well-defined. Intuitively, $D_{\Lambda,s,\mathbf{c}}$ can be viewed as a "conditional" distribution, resulting from sampling $\mathbf{x} \in \mathbb{R}^n$ from a Gaussian centered at $\mathbf{c}$ with parameter s, and conditioning on the event $\mathbf{x} \in \Lambda$.

The smoothing parameter. Micciancio and Regev [MR07] proposed a lattice quantity called the *smoothing parameter*:

Definition 2.6 ([MR07])**.** 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 due to the following (informally stated) 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 over the entire space. (Further discussion along with the precise statement of this fact can be found in [MR07, Section 3].) The smoothing parameter of any n-dimensional lattice is closely related to its nth successive minimum:

Lemma 2.7 ([MR07, Lemma 3.3])**.** *For any $p \in [2, \infty]$, any n-dimensional lattice Λ, and any $\epsilon > 0$,*

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

In particular, for any $\omega(\sqrt{\log n})$ function, there is a negligible function $\epsilon(n)$ for which

$$\eta_\epsilon(\Lambda) \quad \leq \quad \lambda^{(2)}_n(\Lambda) \cdot \omega(\sqrt{\log n}) \quad \leq \quad \lambda^{(p)}_n(\Lambda) \cdot n^{1/2-1/p} \cdot \omega(\sqrt{\log n}).$$

(We note that the inequalities for the ℓ_p norm in Lemma 2.7 follow immediately from standard relations between the ℓ_2 and ℓ_p norms.)

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

Lemma 2.8. *For any $\epsilon \in (0,1)$, $s \geq \eta_\epsilon(\Lambda)$, and $\mathbf{c} \in \mathbb{R}^n$, we have*

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

Proof. The first inequality is implicit in the proof of Lemma 4.4, and the second inequality is Lemma 2.9, of [MR07]. □

For our GapSVP reductions we also need the following lemma, which is implicit in the proofs of Lemma 4.5 and Corollary 4.6 of [MR07].

Lemma 2.9. *Let Λ be any n-dimensional lattice, let $\mathbf{w}, \mathbf{c} \in \mathbb{R}^n$, and let $s \geq \eta_\epsilon(\Lambda)$ for some $\epsilon \in (0,1)$. Then for any $\mathbf{v} \in \mathbb{R}^n$, we have*

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

3 Measure Inequalities for ℓ_p Norms

In this section we review some inequalities developed by Banaszczyk [Ban95] 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 proved to be a crucial quantity in obtaining transference theorems for lattices [Ban93, Ban95], and in the study of the computational complexity of lattice problems (see, e.g., [AR03, AR05, MR07]).

In a prior work of Banaszczyk [Ban93], it was shown that for $p = 2$ and radius $r = \sqrt{n}$, the ratio described above is *exponentially small* in n. The results below generalize this result to arbitrary ℓ_p norms, showing that the ratio is small for $r \sim n^{1/p}$. The ratio is *not*, generally speaking, exponentially small, but for our applications we only need it to be a small constant or negligibly small in n.

Lemma 3.1 ([Ban95, Lemma 2.9])**.** *For any n-dimensional lattice Λ, $p \in [1,\infty)$, $\mathbf{v} \in \mathbb{R}^n$, and $r > 0$,*

$$\frac{\rho((\Lambda - \mathbf{v}) \setminus r\mathcal{B}_n^p)}{\rho(\Lambda)} < 2n \cdot \Gamma\left(\frac{p}{2} + 1\right) \cdot (r\sqrt{\pi})^{-p}.$$

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 $\mathbf{v} \in \mathbb{R}^n$,*

$$\frac{\rho((\Lambda - \mathbf{v}) \setminus c_p n^{1/p} \cdot \mathcal{B}_n^p)}{\rho(\Lambda)} < 1/4.$$

(The $1/4$ bound is arbitrary, and may be replaced by any other constant $\delta > 0$ for some suitable constant c_p.)

Proof. Follows immediately from Lemma 3.1 by setting

$$r = \left(8n \cdot \Gamma\left(\frac{p}{2} + 1\right)\right)^{1/p} / \sqrt{\pi} = c_p \cdot n^{1/p}.$$ □

Lemma 3.3 ([Ban95, Lemma 2.10])**.** *For any n-dimensional lattice Λ, $\mathbf{v} \in \mathbb{R}^n$, and real $r > 0$,*

$$\frac{\rho((\Lambda - \mathbf{v}) \setminus r\mathcal{B}_n^\infty)}{\rho(\Lambda)} < 2n \cdot \exp(-\pi r^2).$$

Corollary 3.4. *Let $n \geq 3$. For any n-dimensional lattice Λ and $\mathbf{v} \in \mathbb{R}^n$,*

$$\frac{\rho((\Lambda - \mathbf{v}) \setminus \sqrt{\log n} \cdot \mathcal{B}_n^\infty)}{\rho(\Lambda)} < 1/4.$$

Proof. Follows immediately from Lemma 3.3 by setting $r = \sqrt{\log(8n)/\pi} \leq \sqrt{\log n}$ for $n \geq 3$. □

3.1 Smoothing Parameter

The measure inequalities also yield bounds on the smoothing parameter relative to the dual minimum distance in various ℓ_p norms. We use these bounds in Section 6.1.2 for the worst-case to average-case reductions for GapSVP in ℓ_p norms, in Sections 6.1.2 and 6.2.

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

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

In particular, for any $\omega(\sqrt{\log n})$ function, there is a negligible function $\epsilon(n)$ such that

$$\eta_\epsilon(\Lambda) \quad \leq \quad \omega(\sqrt{\log n})/\lambda_1^{(\infty)}(\Lambda^*) \quad \leq \quad n^{1/p} \cdot \omega(\sqrt{\log n})/\lambda_1^{(p)}(\Lambda^*).$$

Proof. The inequalities for the ℓ_p norm follow from the fact that $\lambda_1^{(\infty)}(\Lambda^*) \geq \lambda_1^{(p)}(\Lambda^*)/n^{1/p}$, by standard relations between the ℓ_∞ and ℓ_p norms.

For ℓ_∞ norm, let $L = \lambda_1^{(\infty)}(\Lambda^*)$. We have

$$\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) = \rho(s\Lambda^* \setminus \{\mathbf{0}\}) = \rho(s\Lambda^* \setminus sL \cdot \mathcal{B}_n^\infty).$$

Now let $s = \sqrt{\log(2n(1+1/\epsilon))/\pi}/L$. Applying Lemma 3.3,

$$\rho(s\Lambda^* \setminus \{\mathbf{0}\}) < \frac{\epsilon}{1+\epsilon} \cdot \rho(s\Lambda^*) = \frac{\epsilon}{1+\epsilon} \cdot (1 + \rho(s\Lambda^* \setminus \{\mathbf{0}\})).$$

Rearranging terms, we get $\rho(s\Lambda^* \setminus \{\mathbf{0}\}) < \epsilon$, as desired. □

We remark that for $p \leq 2$, the bound in Lemma 3.5 is slightly looser than the known bound of

$$\eta_\epsilon(\Lambda) \leq n^{1/p}/\lambda_1^{(p)}(\Lambda^*)$$

for $\epsilon = 2^{-n}$, which follows from [MR07, Lemma 3.2]. The extra factor in the numerator of our bound (which is identical to the extra factor in Lemma 2.7) arises from needing to deal with arbitrarily small $\epsilon > 0$, rather than a fixed constant like $1/4$. The best dependence on ϵ seems to come by bounding the minimum distance of Λ^* in the ℓ_∞ norm and applying Lemma 3.3 (which introduces the extra factor), rather than by applying Lemma 3.1 directly with the minimum distance of Λ^* in the ℓ_p norm.

4 Problems in coNP

In this section, we show that for $p \geq 2$ and certain $\gamma(n) = \tilde{O}(\sqrt{n})$ approximation factors, the following promise problems in ℓ_p norm are in coNP: the closest vector problem GapCVP^p_γ, the shortest vector problem GapSVP^p_γ, the covering radius problem GapCRP^p_γ, and the shortest independent vectors problem $\mathsf{GapSIVP}^p_\gamma$. This implies that these problems are not NP-hard unless the polynomial-time hierarchy collapses (see [GG00, Gol, Cai98] for a discussion of some subtleties concerning promise problems and the polynomial-time hierarchy). For similar approximation factors, we also show that the closest vector with preprocessing GapCVPP and bounded distance decoding with preprocessing BDD in the ℓ_p norm are *easy* (i.e., in P).

The results in this section follow from an application of the measure inequalities from Section 3 to prior work by Aharonov and Regev [AR05], who developed the main techniques for the ℓ_2 norm.

4.1 Closest Vector Problem

The main result we need is $\mathsf{GapCVP}^p_\gamma \in \mathsf{coNP}$ for $p \in [2, \infty)$ and appropriate $\gamma(n) = O(\sqrt{n})$ (and, for $p = \infty$, some $\gamma(n) = O(\sqrt{n \log n})$). Other problems are then be placed in coNP via known reductions to GapCVP, which work for arbitrary norms and approximation factors.

We start with an informal overview of the main proof technique of Aharonov and Regev [AR05]. They show that for any n-dimensional lattice Λ, there is a positive function $f : \mathbb{R}^n \to [0, 1]$ that indicates whether an arbitrary point $\mathbf{v} \in \mathbb{R}^n$ is close to, or far from, the lattice (in ℓ_2 norm). The function f is in fact the (normalized) sum of Gaussians centered at every lattice point, i.e.,

$$f(\mathbf{v}) = \rho(\Lambda - \mathbf{v})/\rho(\Lambda).$$

When $\mathbf{v}$ is within distance (say) $1/100$ of Λ, it is relatively straightforward to show that $f(\mathbf{v}) \geq 1/2$. On the other hand, when $\mathbf{v}$ is far away from Λ, the measure inequalities of [Ban93, Ban95] imply that $f(\mathbf{v})$ is quite small. (For example, when $\mathrm{dist}^2(\mathbf{v}, \Lambda) > \sqrt{n}$, we have $f(\mathbf{v}) < 2^{-n}$.)

A main technical result of [AR05] is that f can be *succinctly approximated* by an efficiently-computable function $f_{\mathbf{W}}$, where $\mathbf{W} \in \mathbb{R}^{n \times N}$ is a matrix made up of $N = \mathrm{poly}(n)$ vectors $\mathbf{w}_i$ from the dual lattice Λ^*. The vectors $\mathbf{w}_i$ are chosen independently at random from the Fourier spectrum of f, which just so happens to be the discrete Gaussian distribution D_{Λ^*}. With good probability over these random choices, $f_{\mathbf{W}}$ is a very good (pointwise) approximation to f.

Putting together all these facts results in an NP proof system for the *complement* of GapCVP_γ, for $\gamma(n) = 100\sqrt{n}$. The witness that a point $\mathbf{v}$ is far from Λ is simply a suitable matrix $\mathbf{W}$ defining the function $f_{\mathbf{W}} \approx f$. The verifier accepts if $f_{\mathbf{W}}(\mathbf{v})$ is small (say, less than $1/2$), and if $f_{\mathbf{W}}$ is a good enough approximation to f (this is more technical, but can also be checked efficiently via a spectral test on $\mathbf{W}$). When $\mathrm{dist}^2(\mathbf{v}, \Lambda) > \sqrt{n}$, then $f_{\mathbf{W}}(\mathbf{v}) \approx f(\mathbf{v}) < 2^{-n}$ is very small, and the verifier accepts. On the other hand, when $\mathrm{dist}(\mathbf{v}, \Lambda) \leq 1/100$, then $f_{\mathbf{W}}(\mathbf{v}) \geq 1/2$ for any acceptable $\mathbf{W}$, causing the verifier to reject.

Overview of analysis for ℓ_p norms. Now consider the ℓ_p norm for $p \geq 2$. It turns out that we can use *exactly the same* witness and verifier; 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 basic relations among norms. In such a case, we already are guaranteed that the verifier 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 $f_{\mathbf{W}}(\mathbf{v}) \approx f(\mathbf{v}) < 1/4$, so the verifier accepts. The resulting gap factor is therefore $O(n^{1/p}/n^{1/p-1/2}) = O(\sqrt{n})$.

Unfortunately, when $1 \leq p < 2$, the above analysis breaks down. Using the measure inequalities, we can show that the verifier still accepts when $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c_p n^{1/p}$ for an appropriate constant c_p. However, soundness is compromised: if $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p-1/2}$, it may also be the case that $\mathrm{dist}^2(\mathbf{v}, \Lambda) = n^{1/p-1/2} \gg 1$, and $f(\mathbf{v})$ may be small enough to fool the verifier. In order to guarantee that $\mathrm{dist}^2(\mathbf{v}, \Lambda) \leq 1/100$, we must also require that $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq 1/100$. This yields a gap factor of $O(n^{1/p})$, which already follows trivially from the original analysis in the ℓ_2 norm and its relationship to ℓ_p norms. We do not know if there is an alternate proof system that improves upon this factor.

We now proceed more formally.

Theorem 4.1. *For any $p \in [2, \infty)$, there is a constant $c_p \approx \sqrt{p}$ such that $\mathsf{GapCVP}^p_{100c_p\sqrt{n}}$ belongs to $\mathsf{NP} \cap \mathsf{coNP}$.*
For $p = \infty$, $\mathsf{GapCVP}^\infty_{100\sqrt{n \log n}}$ belongs to $\mathsf{NP} \cap \mathsf{coNP}$.

In order to prove the theorem, we need to recall a few tools from [AR05]. First we recall the verifier algorithm $\mathcal{V}$ for the complement of GapCVP: the input is a basis $\mathbf{B}$ for an n-dimensional lattice $\Lambda = \mathcal{L}(\mathbf{B})$ and a target point $\mathbf{v} \in \mathbb{R}^n$, and the witness is a matrix $\mathbf{W} \in \mathbb{R}^{n \times N}$, for some large enough $N = \mathrm{poly}(n)$. The verifier checks the following conditions, accepting if all three hold true and rejecting otherwise:

1. Check that $f_{\mathbf{W}}(\mathbf{v}) < 1/2$, where $f_{\mathbf{W}}$ is the function $f_{\mathbf{W}}(\mathbf{v}) = \frac{1}{N} \sum_{i \in [N]} \cos(2\pi \langle \mathbf{v}, \mathbf{w}_i \rangle)$.
2. Check that $\mathbf{w}_i \in \Lambda^*$ for all $i \in [N]$, i.e., that each column of $\mathbf{W}$ is in the dual lattice.
3. Check that the largest eigenvalue of the positive semidefinite matrix $\mathbf{W}\mathbf{W}^T \in \mathbb{R}^{n \times n}$ is at most $3N$.

Using standard algorithms, $\mathcal{V}$ can be implemented in polynomial time.

The first two facts we need are concerned with completeness. In the following two lemmas, suppose we choose the N columns $\mathbf{w}_i$ of an $n \times N$ matrix $\mathbf{W}$ independently according to D_{Λ^*}, i.e., the discrete Gaussian distribution over the dual lattice Λ^*. The first fact says that with good probability, the eigenvalues of $\mathbf{W}\mathbf{W}^T$ are not too large, as required by Condition 3 of the verifier. The second fact says that with good probability, $f_{\mathbf{W}}$ is a very good approximation to f essentially everywhere (more precisely, on a fine grid of any desired precision). The proofs of these facts are somewhat lengthy, so we omit them and direct the interested reader to [AR05].

Lemma 4.2 ([AR05, Lemma 6.3])**.** *The probability that $\mathbf{W}$ satisfies Condition 3 of the verifier algorithm $\mathcal{V}$ is at least $3/4$.*

Lemma 4.3 ([AR05, Lemma 1.3])**.** *Let $\ell(n) = \mathrm{poly}(n)$ be a precision parameter, let $c > 0$ be any constant, and let $N = \ell(n) \cdot n^{2c+2} = \mathrm{poly}(n)$. Then with probability at least $3/4$ over the choice of $\mathbf{W}$, we have*

$$|f_{\mathbf{W}}(\mathbf{v}) - f(\mathbf{v})| \leq \frac{1}{n^c}$$

for every $\mathbf{v} \in \Lambda/2^{\ell(n)}$*, where $f_{\mathbf{W}}$ is defined as in Condition 1 of the verifier algorithm $\mathcal{V}$.*

The final fact that we need is concerned with the soundness of the verifier algorithm $\mathcal{V}$.

Lemma 4.4 ([AR05, Section 6.1])**.** *If* $\mathrm{dist}^2(\mathbf{v}, \mathcal{L}(\mathbf{B})) \leq 1/100$*, then* $\mathcal{V}$ *always rejects on* $(\mathbf{B}, \mathbf{v})$ *and any* $\mathbf{W}$*.*

Proof. We repeat the short proof from [AR05] for self-containment. Suppose that Conditions 2 and 3 are passed; we show that Condition 1 must fail. Because Condition 2 is passed, the function $f_{\mathbf{W}}$ is periodic over $\mathcal{L}(\mathbf{B})$. It therefore suffices to prove that $f_{\mathbf{W}}(\mathbf{v}) \geq 1/2$ for any $\mathbf{v}$ such that $\|\mathbf{v}\| \leq 1/100$. Now because Condition 3 is passed, we have

$$\frac{1}{N} \sum_{i \in [N]} \langle \mathbf{v}, \mathbf{w}_i \rangle^2 = \frac{1}{N} \mathbf{v}^T \mathbf{W} \mathbf{W}^T \mathbf{v} \leq \frac{1}{N} \cdot \frac{3N}{10000} = \frac{3}{10000},$$

where the inequality follows by expressing $\mathbf{v}$ in the (orthonormal) eigenvector basis of $\mathbf{W}\mathbf{W}^T$. Using the inequality $\cos\theta \geq 1 - \theta^2/2$, we then get

$$f_{\mathbf{W}}(\mathbf{v}) = \frac{1}{N} \sum_{i \in [N]} \cos(2\pi \langle \mathbf{v}, \mathbf{w}_i \rangle) \geq 1 - \frac{4\pi^2}{2N} \sum_{i \in [N]} \langle \mathbf{v}, \mathbf{w}_j \rangle^2 \geq 1 - \frac{6\pi^2}{10000} > \frac{1}{2}. \qquad \square$$

We are now ready to prove the main theorem.

Proof of Theorem 4.1. Membership in NP is trivial, as are the cases $n = 1, 2$. Thus it suffices to give an NP verifier for the complement of GapCVP, assuming $n \geq 3$. Consider an instance $(\mathbf{B}, \mathbf{v})$ to the *complement* of GapCVP, and let $\Lambda = \mathcal{L}(\mathbf{B})$. Without loss of generality, we can assume by scaling the input that for $p \in [2, \infty)$, NO instances are such that $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p - 1/2}/100$, while YES instances are such that $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c_p n^{1/p}$. Likewise, for $p = \infty$, NO instances are such that $\mathrm{dist}(\mathbf{v}, \Lambda) \leq n^{-1/2}/100$, while YES instances are such that $\mathrm{dist}(\mathbf{v}, \Lambda) > \sqrt{\log n}$.

First we show soundness: suppose that $(\mathbf{B}, \mathbf{v})$ is a NO instance, that is, $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p - 1/2}/100$. Then by the properties of ℓ_p norms, we have $\mathrm{dist}^2(\mathbf{v}, \Lambda) \leq 1/100$. By Lemma 4.4, $\mathcal{V}$ always rejects on any witness $\mathbf{W}$, as desired.

We now show completeness: suppose that $p \in [2, \infty)$ and $(\mathbf{B}, \mathbf{v})$ is a YES instance, that is, $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c_p n^{1/p}$ for appropriate constant c_p. Then Corollary 3.2 implies that

$$f(\mathbf{v}) = \frac{\rho(\Lambda - \mathbf{v})}{\rho(\Lambda)} = \frac{\rho((\Lambda - \mathbf{v}) \backslash c_p n^{1/p} \cdot \mathcal{B}_n^p)}{\rho(\Lambda)} < 1/4.$$

Now suppose we choose the witness $\mathbf{W}$ as above, by sampling each column $\mathbf{w}_i$ independently according to D_{Λ^*}. By this choice, clearly $\mathcal{V}$'s Condition 2 holds true, and by Lemma 4.2, Condition 3 holds true except with probability at most $1/4$. Furthermore, by setting $c = 2$ in Lemma 4.3, we have that $f_{\mathbf{W}}(\mathbf{v}) \leq f(\mathbf{v}) + 1/n^2 < 1/2$ except with probability at most $1/4$. We conclude that all conditions hold true with at least $1/2$ probability over the choice of the witness $\mathbf{W}$, implying that there exists some $\mathbf{W}$ that causes $\mathcal{V}$ to accept.

For $p = \infty$, if $(\mathbf{B}, \mathbf{v})$ is a YES instance we have $\mathrm{dist}^\infty(\mathbf{v}, \Lambda) > \sqrt{\log n}$. Repeating the above argument using Corollary 3.4, the proof is complete. $\square$

4.2 Other Problems in coNP

Theorem 4.5. *For any* $p \in [2, \infty)$*, there is a constant* $c_p \approx \sqrt{p}$ *such that the problems* GapSVP^p_γ*,* $\mathsf{GapCRP}^p_{2\gamma}$*,* $\mathsf{GapSIVP}^p_\gamma$ *belong to* coNP*, for* $\gamma(n) = 100 c_p \sqrt{n}$*.*

For $p = \infty$*, all of the above problems belong to* coNP *for* $\gamma(n) = 100\sqrt{n \log n}$*.*

Proof. The claims follow via known approximation- and norm-preserving reductions from the various problems to GapCVP.

For the shortest vector problem (in any norm), there is a deterministic non-adaptive Cook reduction from GapSVP_γ to GapCVP_γ (in any norm, for any $\gamma \geq 1$) due to Goldreich *et al* [GMSS99, Theorem 6]. As shown in [AR05, Lemma A.1], it follows that if GapCVP_γ is in coNP, then so is GapSVP_γ. Essentially, a YES instance of GapSVP maps to at least one YES instance of GapCVP (which lacks a coNP witness), whereas a NO instance of GapSVP maps to several NO instances of GapCVP. Therefore, it suffices to give coNP witnesses for all of the GapCVP instances produced by the reduction.

For the covering radius problem, the proof of Theorem 4.5 of Guruswami *et al* [GMR05] implicitly describes a nondeterministic Karp reduction from the complement of GapCRP_γ to the complement of GapCVP_γ (in any norm, for any $\gamma(n) \geq 1$). Using this reduction, their proof establishes that if GapCVP_γ is in coNP, then so is $\mathsf{GapCRP}_{2\gamma}$ (the reduction loses a factor of two in the approximation).

For the shortest independent vectors problem, Theorem 4.9 and Corollary 4.10 of [GMR05] establish (via a nondeterministic reduction) that if GapCVP_γ is in coNP, then so is $\mathsf{GapSIVP}_\gamma$ (for any $\gamma(n) \geq 1$). Their proofs are independent of the choice of norm, except in the proof of Theorem 4.9 where specific properties of the ℓ_2 norm are used to argue that a certain quantity M can be made polynomial in the input size. Because the ℓ_p norm is always within a $\sqrt{n}$ factor of the ℓ_2 norm, the same proof establishes that there is a suitable polynomial M for any ℓ_p norm. □

4.3 Easy Problems with Preprocessing

Lattice problems with *preprocessing* model situations in which a lattice is fixed long before an actual instance of the problem is generated. For example, a fixed lattice $\Lambda = \mathcal{L}(\mathbf{B})$ may be used as a kind of error-correcting code, and a (noisy) received message would be represented by a target point $\mathbf{v}$. The goal would be to decode $\mathbf{v}$ to a nearby element of Λ, given some suitable (short) advice about Λ that assists the decoding process. Because the lattice is fixed far in advance, the advice can be viewed as the output of a preprocessing phase, whose running time does not count toward the complexity of the decoding algorithm. (For additional motivation and discussion of preprocessing, see [FM04].)

Here we define two variants of the closest vector problem with preprocessing. The first, defined by Feige and Micciancio [FM04], is a decision version whose goal is to efficiently distinguish points that are close to the lattice from those than are very far from the lattice. The second, due to Liu *et al* [LLM06], is a search version whose goal is to efficiently decode a target point to its (unique) closest lattice point, under the promise that the target is within a certain fraction (say, 1/10) of the minimum distance of the lattice. In the following problems, the norm $\|\cdot\|$ is implicit; we attach a superscript $p \in [1, \infty]$ to the problem name to indicate the ℓ_p norm.

Definition 4.6. A solution to the *closest vector problem with preprocessing* $\mathsf{GapCVPP}_\gamma$ (respectively, *bounded distance decoding problem with preprocessing* BDD_γ) is given by a preprocessing function P (which may be very hard to compute) and a decision algorithm (respectively, decoding algorithm) D having the following properties:

- On input an n-dimensional basis $\mathbf{B}$ for lattice $\Lambda = \mathcal{L}(\mathbf{B})$, P returns a succinct advice string A, i.e., the length of A is at most a fixed polynomial in the length of $\mathbf{B}$.
- For GapCVPP: given $A = P(\mathbf{B})$ and a target vector $\mathbf{v} \in \mathbb{R}^n$, D accepts if $\mathrm{dist}(\mathbf{v}, \Lambda) \leq 1$ and rejects if $\mathrm{dist}(\mathbf{v}, \Lambda) > \gamma(n)$ (otherwise, any output is acceptable).

- For BDD: given $A = P(\mathbf{B})$ and a target $\mathbf{v} \in \mathbb{R}^n$ such that $\mathrm{dist}(\mathbf{v}, \Lambda) \leq \lambda_1(\Lambda)/\gamma(n)$, D outputs the unique $\mathbf{x} \in \Lambda$ closest to $\mathbf{v}$ (in order to guarantee that $\mathbf{x}$ is unique, we require $\gamma > 2$).

The complexity of the solution is measured by the running time of D alone.

(Formally, as with GapCVP we fix some precision parameter $\ell(n) = \mathrm{poly}(n)$ in advance and represent the target vector $\mathbf{v}$ relative to the basis $\mathbf{B}$ using coefficients having at most $\ell(n)$ bits of precision, so that $\mathbf{v} \in \Lambda/2^{\ell(n)}$.)

Aharonov and Regev showed that $\mathsf{GapCVPP}$ in the ℓ_2 norm is easy for any $\gamma(n) = O(\sqrt{n/\log n})$ factor [AR05]. Their result extends to ℓ_p norms using a similar analysis as above.

Theorem 4.7. *For any $p \in [2, \infty]$, $\mathsf{GapCVPP}^p_{10\sqrt{n}}$ can be solved in deterministic polynomial time.*

Proof. Let $\Lambda = \mathcal{L}(\mathbf{B})$ where $\mathbf{B}$ is the input basis. By scaling, we can assume without loss of generality that YES instances are such that $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p-1/2} \cdot \sqrt{\log n}$, while NO instances are such that $\mathrm{dist}^p(\mathbf{v}, \Lambda) > 10n^{1/p} \cdot \sqrt{\log n}$.

By Lemma 4.3, there is some $N = \mathrm{poly}(n)$ and a matrix $\mathbf{W} \in \mathbb{R}^{n \times N}$ such that the function $f_{\mathbf{W}}$ approximates the function f at any point $\mathbf{v} \in \mathcal{L}(B)/2^{\ell(n)}$ to within n^{-10}. Given $\mathbf{B}$, the preprocessing function outputs such $\mathbf{W}$ as the advice. Given this advice and a target point $\mathbf{v}$, the decision algorithm D outputs YES if $f_{\mathbf{W}}(\mathbf{v}) \geq n^{-4}$, otherwise it outputs NO.

Suppose $\mathbf{v}$ is a YES instance. Let $d = \mathrm{dist}^2(\mathbf{v}, \Lambda) \leq \sqrt{\log n}$. Lemma 3.2 of [AR05] establishes that $f(\mathbf{v}) \geq \exp(-\pi d^2) \geq n^{-\pi}$, therefore $f_{\mathbf{W}}(\mathbf{v}) \geq n^{-\pi} - n^{-10} \geq n^{-4}$ and D accepts. Now suppose $\mathbf{v}$ is a NO instance, so that $\mathrm{dist}^\infty(\mathbf{v}, \Lambda) > 10\sqrt{\log n}$. By Lemma 3.3, we have $f(\mathbf{v}) < n^{-100}$, so $f_{\mathbf{W}}(\mathbf{v}) < n^{-100} + n^{-10} < n^{-4}$ and D rejects, as desired. □

Building upon the techniques of [AR05], Liu *et al* [LLM06] showed that the search problem BDD_γ (with preprocessing) in the ℓ_2 norm is also easy for certain $\gamma(n) = O(\sqrt{n/\log n})$ factors. Their decoding algorithm follows a "hill-climbing" approach using the function $f_{\mathbf{W}} \approx f$ to move closer and closer to the lattice Λ, until it becomes close enough that the nearest lattice vector can be found by other means. The hill-climbing process works as long as f can be used to compute (to a high degree of precision) the distance $\mathrm{dist}(\mathbf{w}, \Lambda)$ for all $\mathbf{w}$ that are suitably close to Λ. The main claim can be abstracted as follows.

Lemma 4.8 ([LLM06, Theorem 2])**.** *There is an preprocessing function P and a deterministic polynomial-time algorithm D having the following properties:*

1. *Given a basis $\mathbf{B}$ of an n-dimensional lattice $\Lambda = \mathcal{L}(\mathbf{B})$, P outputs a succinct advice string A.*
2. *Suppose that*[4]
$$f(\mathbf{w}) \leq \rho(\mathbf{w}) + n^{-100}$$
for all $\mathbf{w} \in \mathbb{R}^n$ of length $\|\mathbf{w}\|_2 \leq \sqrt{\log n}$. Then given advice string A and any $\mathbf{v}$ such that $\mathrm{dist}^2(\mathbf{v}, \Lambda) \leq \frac{1}{2}\sqrt{\log n}$, D outputs the vector $\mathbf{x} \in \Lambda$ closest to $\mathbf{v}$ (in the ℓ_2 norm).

We now extend this result to solve BDD in ℓ_p norms.

Theorem 4.9. *For any $p \in [2, \infty]$, $\mathsf{BDD}^p_{20\sqrt{n}}$ can be solved in deterministic polynomial time.*

[4]Recall that $f(\mathbf{w}) \geq \rho(\mathbf{w})$ for all lattices Λ and all $\mathbf{w} \in \mathbb{R}^n$, by [AR05, Lemma 3.2]. The choice of n^{-100} is rather arbitrary and can be replaced with the inverse of any suitably large polynomial in n.

Proof. Let $\Lambda = \mathcal{L}(\mathbf{B})$ where $\mathbf{B}$ is the input basis. By scaling, we can assume without loss of generality that $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p-1/2} \cdot \sqrt{\log n}$ and $\lambda_1^{(p)}(\Lambda) \geq 10n^{1/p} \cdot \sqrt{\log n}$. This implies that

$$\mathrm{dist}^2(\mathbf{v}, \Lambda) \leq \tfrac{1}{2}\sqrt{\log n} \qquad \text{and} \qquad \lambda_1^{(p)}(\Lambda) \geq \lambda_1^{(\infty)}(\Lambda) \geq 10\sqrt{\log n}.$$

In particular, there is a unique $\mathbf{x} \in \Lambda$ that is closest to $\mathbf{v}$ in both the ℓ_2 and ℓ_p norms. Therefore it suffices to use the algorithm guaranteed by Lemma 4.8, which finds $\mathbf{x}$.

It remains to prove the hypothesis of Lemma 4.8. Suppose $\|\mathbf{w}\|_2 \leq \sqrt{\log n}$, which implies $\|\mathbf{w}\|_\infty \leq \sqrt{\log n}$. Then because $\lambda_1^{(\infty)}(\Lambda) \geq 10\sqrt{\log n}$, it is the case that

$$(\Lambda - \mathbf{w}) \setminus \{\mathbf{w}\} = (\Lambda - \mathbf{w}) \setminus (9\sqrt{\log n} \cdot \mathcal{B}_n^\infty).$$

Therefore we have

$$f(\mathbf{w}) = \frac{\rho(\Lambda - \mathbf{w})}{\rho(\Lambda)} = \frac{\rho(\mathbf{w})}{\rho(\Lambda)} + \frac{\rho((\Lambda - \mathbf{w}) \setminus (9\sqrt{\log n} \cdot \mathcal{B}_n^\infty))}{\rho(\Lambda)} \leq \rho(\mathbf{w}) + n^{-100},$$

where the inequality follows from the fact that $\rho(\Lambda) \geq 1$ and by Lemma 3.3. □

We note that the approximation factors in Theorems 4.7 and 4.9 are an $O(\sqrt{\log n})$ factor looser than those in [AR05, Theorem 1.4] and [LLM06, Theorem 2]. This stems from a corresponding looseness in the measure inequalities for ℓ_p norms. In the ℓ_2 norm, we have $f(\mathbf{v}) < 2^{-n}$ when $\mathrm{dist}(\mathbf{v}, \Lambda) > \sqrt{n}$ [Ban93]. In contrast, the inequalities for ℓ_p norms from [Ban95] require $\mathrm{dist}(\mathbf{v}, \Lambda) \geq \Omega(n^{1/p} \cdot \sqrt{\log n})$ to obtain even an inverse polynomial (say, $1/n$) upper bound on $f(\mathbf{v})$. While it is possible to bound $f(\mathbf{v})$ by a constant (say, $1/4$) when $\mathrm{dist}(\mathbf{v}, \Lambda) \geq c_p n^{1/p}$, this merely leads to the loss of a $\sqrt{\log n}$ factor on the YES instances, to no overall benefit.

5 Analysis of Discrete Gaussians

In this section, we develop new tools for analyzing worst-case to average-case reductions that use Gaussian measures. Our main result is a general bound on the *moments* of discrete Gaussian distributions over lattices. These moments are nearly identical to those of *continuous* Gaussian distributions. As a consequence, many essential facts about continuous Gaussians also carry over to the discrete case. These include, for example, exponential tail bounds and "nice" behavior of sums of independent samples.

Our analysis seems to be a natural continuation to prior study of discrete Gaussians. Using ideas of Banaszczyk [Ban93], Micciancio and Regev [MR07] analyzed the low-order moments of discrete Gaussians. Lyubashevsky and Micciancio [LM06] extended this analysis to higher moments. Unfortunately, even at this stage the analysis becomes quite cumbersome. Our analysis is more general but actually more modular and arguably simpler, due to the techniques from [Ban95].

5.1 Overview of Techniques

Here we give a simplified overview of the techniques for analyzing a single discrete Gaussian over a lattice, centered at the origin. Our main result are general enough to apply to *sums* of several discrete Gaussians (even over different lattices, having different centers, etc.).

Let Λ be a sufficiently dense lattice in $\mathbb{R}^n$, and suppose that $\mathbf{x} \in \Lambda$ is a random variable with distribution D_Λ.[5] 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 $\mathrm{E}[|x_i|^p]$, i.e., the pth moment of $|x_i|$. The crucial tool we use is an exponential tail inequality on x_i:

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

$$\Pr_{\mathbf{x}\sim D_\Lambda}[|x_i| \geq r] \leq \exp(-\Theta(r^2)).$$

This inequality is stated precisely and in full generality as Lemma 5.1 below. We remark that for *continuous* Gaussians, proving this inequality is straightforward using direct integration. However, for *discrete* Gaussians the path is not so straightforward.

To prove the tail inequality and complete the analysis, we draw upon techniques of [Ban95]. First, consider $\mathbf{x}$'s probability distribution as a positive function $D = D_\Lambda : \Lambda \to \mathbb{R}^+$. Then our goal is to bound the total measure assigned by D to a subset $\Lambda^- = \{\mathbf{x} \in \Lambda \ : \ |x_i| \geq r\}$ of the lattice. The general strategy is to find some positive function $g : \Lambda \to \mathbb{R}^+$ satisfying two conditions:

1. The measure $(D \cdot g)(\Lambda)$ on the entire lattice exceeds the measure $D(\Lambda) = 1$ by only at most a "small" factor c.
2. The measure $(D \cdot g)(\Lambda^-)$ on the subset exceeds the measure $D(\Lambda^-)$ by at least a "very large" factor C.

Because D and g are positive, we then 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 very small quantity. It turns out that a good choice for the function g that satisfies the two requirements is $g(\mathbf{x}) = \cosh(2\pi r x_i)$, where $\cosh(x) = \frac{1}{2}(e^x + e^{-x})$ is the hyperbolic cosine function. Intuitively, this choice works because g is relatively small when $|x_i|$ is small (which is where most of the measure of D lies), but becomes very large when $|x_i|$ is large.

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

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

The final integral is the definition of the Γ function, and evaluates to approximately $(\sqrt{p})^p$. When plugged into Equation (1), this yields

$$\mathrm{E}[\|\mathbf{x}\|_p] \leq \sqrt{p} \cdot n^{1/p}.$$

[5] In the general case, $\mathbf{x}$ may be drawn from $D_{\Lambda,s,\mathbf{c}}$ for any parameter s and arbitrary centers $\mathbf{c}$. In order to illuminate the key ideas, we focus on the simpler case in this overview.

That is, for any fixed ℓ_p norm (for finite p), a sample from an n-dimensional discrete Gaussian has expected norm $\sim n^{1/p}$, just as is the case for a continuous Gaussian. For the ℓ_∞ norm, a similar argument shows that the norm is bounded by $\sim \sqrt{\log n}$ with good probability.

We devote the remainder of this section to the full statement of the result and its proof.

5.2 Main Results

Here we state our central claims concerning discrete Gaussians, deferring their proofs to Section 5.3.

Let $\mathbf{U} = \{\mathbf{u}_1, \ldots, \mathbf{u}_d\}$ be a set of $d \geq 1$ orthonormal vectors in $\mathbb{R}^n$. Define the "$\mathbf{U}$ norm" as

$$\|\mathbf{x}\|_{\mathbf{U}} = \sum\nolimits_{i \in [d]} |\langle \mathbf{x}, \mathbf{u}_i \rangle|$$

for any $\mathbf{x} \in \mathbb{R}^n$. For example, consider the case $d = 1$. Then $\|\mathbf{x}\|_{\mathbf{U}}$ is simply the magnitude of the component of $\mathbf{x}$ parallel to the unit vector $\mathbf{u}_1$. More generally, the $\mathbf{U}$ norm is akin to the ℓ_1 norm within the subspace spanned by $\mathbf{U}$.

When a $\mathbf{U}$ norm is clear from context, it is convenient to use the following notation: for any $r \geq 0$ and any $\mathbf{c} \in \mathbb{R}^n$, define the open "$\mathbf{U}$ cylinder" $Q_{r,\mathbf{c}}$ of radius r centered at $\mathbf{c}$ as

$$Q_{r,\mathbf{c}} = \{\mathbf{x} \in \mathbb{R}^n \ : \ \|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}} < r\}.$$

The following tail inequality is the central tool for proving our main results.

Lemma 5.1 (Tail Inequality)**.** *For any n-dimensional lattice Λ, any $\mathbf{c} \in \mathbb{R}^n$, and any $r \geq 0$,*

$$\rho_{\mathbf{c}}(\Lambda \setminus Q_{r,\mathbf{c}}) \quad \leq \quad 2^d \cdot \exp(-\pi r^2/d) \cdot \rho(\Lambda).$$

(Recall that $d = |\mathbf{U}|$; the 2^d term is a side effect of the proof techniques when working with the $\mathbf{U}$ norm. Fortunately, we have $d = 1$ in all of our applications; the work of [PR07] also requires $d = 2$, but no more.)

Our main theorem concerns the moments of discrete Gaussian distributions about their centers. Of course, moments are defined for distributions over $\mathbb{R}$, whereas a discrete Gaussian is distributed over $\mathbb{R}^n$. To be completely precise, we bound the moments of $\|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}}$, i.e., the distance (in the $\mathbf{U}$ norm) of a discrete Gaussian sample $\mathbf{x}$ from its center $\mathbf{c}$.

Theorem 5.2 (Main Theorem: Moments of discrete Gaussians)**.** *For any n-dimensional lattice Λ, real $p \in [1, \infty)$, $\mathbf{c} \in \mathbb{R}^n$, and $\mathbf{U}$ as above,*

$$\mathop{\mathrm{E}}_{\mathbf{x} \sim D_{\Lambda,\mathbf{c}}} \left[\|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}}^p\right] \leq 2^d \cdot \left(\frac{d}{\pi}\right)^{p/2} \cdot \Gamma\left(\frac{p}{2} + 1\right) \cdot \frac{\rho(\Lambda)}{\rho_{\mathbf{c}}(\Lambda)}.$$

The bound given by the above theorem is quite precise, but is cumbersome to work with directly. We now present a corollary that is more suitable for our applications, in which we need to bound the ℓ_p norm of the *weighted sum* of several independent samples from a discrete Gaussian over some fixed lattice.

Corollary 5.3 (Weighted sums of discrete Gaussians)**.** *Let m be a positive integer, $\epsilon \leq 1/(2m+1)$ be a positive real, and let $\mathbf{z} \in \mathbb{R}^m$ be a vector of m weights. Let Λ be an n-dimensional lattice, let $s \geq \eta_\epsilon(\Lambda)$, and let $\mathbf{C} \in \mathbb{R}^{n \times m}$.*

For any $p \in [1, \infty)$, there is a constant $c_p \approx \sqrt{p}$ (depending only on p) such that

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

where the expectation is taken over independent samples $\mathbf{x}_i \sim D_{\Lambda,s,\mathbf{c}_i}$ for each $i \in [m]$.
For $p = \infty$ and any $r \geq 0$,

$$\Pr_{\mathbf{x}_i \sim D_{\Lambda,s,\mathbf{c}_i}} [\|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_\infty \geq s \|\mathbf{z}\|_2 \cdot r] \leq 2en \cdot \exp(-\pi r^2).$$

In particular, for $r = \sqrt{\log n}$ and $n \geq 3$ the above probability is at most $2/3$, and for $r = r(n) = \omega(\sqrt{\log n})$ the above probability is a negligible function in n.

We remark that Corollary 5.3 can be generalized further, e.g., to deal with samples $\mathbf{x}_i \sim D_{\Lambda_i,s_i,\mathbf{c}_i}$ for arbitrary n-dimensional lattices Λ_i and Gaussian parameters $s_i \geq \eta_\epsilon(\Lambda_i)$. We omit such generalizations for simplicity, and because the current formulation is expressive enough for our applications.

5.3 Proofs of Claims

We now prove the claims of the previous subsection. The proofs are technical in places; the reader who is interested only in the complexity-theoretic applications may wish to skip to Section 6. We first prove the main theorem on the moments of discrete Gaussians (Theorem 5.2) and its corollary, assuming the tail inequality (Lemma 5.1). We conclude the section with the proof of the tail inequality.

5.3.1 Proofs of Main Theorem and Corollary

Proof of Theorem 5.2. The proof closely follows the structure of Equation (2) from our overview, but generalized to the $\mathbf{U}$ norm. We have

$$\begin{aligned}
\mathop{\mathrm{E}}_{\mathbf{x} \sim D_{\Lambda,\mathbf{c}}} \left[\|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}}^p \right] &= \sum_{\mathbf{x} \in \Lambda} \|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}}^p \cdot D_{\Lambda,\mathbf{c}}(\mathbf{x}) && \text{(def. of E)} \\
&= p \int_{r=0}^{\infty} r^{p-1} \cdot \sum_{\mathbf{x} \in \Lambda \setminus Q_{r,\mathbf{c}}} D_{\Lambda,\mathbf{c}}(\mathbf{x})\, dr && \text{(calculus; see (2))} \\
&\leq 2^d \cdot p \int_{r=0}^{\infty} r^{p-1} \exp(-\pi r^2/d) \cdot \frac{\rho(\Lambda)}{\rho_{\mathbf{c}}(\Lambda)}\, dr && \text{(Lemma 5.1)} \\
&= 2^d \cdot (d/\pi)^{p/2} \cdot \Gamma\left(\tfrac{p}{2} + 1\right) \cdot \frac{\rho(\Lambda)}{\rho_{\mathbf{c}}(\Lambda)}. && \text{(def. of } \Gamma\text{)}
\end{aligned}$$

□

Proof of Corollary 5.3. The main idea behind the proof is to combine the m Gaussian samples from the n-dimensional lattice Λ into one "super-sample" from a Gaussian over an nm-dimensional lattice, at which point we can apply Theorem 5.2.

First, we may assume without loss of generality that $\mathbf{z} \neq \mathbf{0}$ (otherwise, the result is trivially true). Furthermore, we can assume that $s = 1$ by replacing Λ with Λ/s, $\mathbf{C}$ by $\mathbf{C}/s$, and $\mathbf{X}$ by $\mathbf{X}/s$.

Now define a new lattice $\Lambda' = \Lambda \times \cdots \times \Lambda$ as the Cartesian product of m copies of Λ (equivalently, viewing Λ as an additive subgroup of $\mathbb{R}^n$, Λ' is the direct sum of m copies of Λ). Then Λ' is an nm-dimensional lattice in $\mathbb{R}^{nm}$. Likewise, define $\mathbf{c}' = (\mathbf{c}_1, \ldots, \mathbf{c}_m) \in \mathbb{R}^{nm}$ and $\mathbf{x}' = (\mathbf{x}'_1, \ldots, \mathbf{x}'_m) \in \mathbb{R}^{nm}$.

By a routine calculation using the definition of ρ, for any countable sets $A_i \subset \mathbb{R}^n$ and any $\mathbf{a}_i \in \mathbb{R}^n$ for $i \in [m]$, we have

$$\rho_{(\mathbf{a}_1,\ldots,\mathbf{a}_m)}(A_1 \times \cdots \times A_m) = \prod\nolimits_{i\in[m]} \rho_{\mathbf{a}_i}(A_i).$$

It follows that $\mathbf{x}'$ is distributed according to $D_{\Lambda',\mathbf{c}'}$. Furthermore, by Lemma 2.8 and by hypothesis on ϵ,

$$\frac{\rho(\Lambda')}{\rho_{\mathbf{c}'}(\Lambda')} \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)^m \leq \left(1+\frac{1}{m}\right)^m \leq \exp(1) = e. \tag{3}$$

We are now ready to analyze the ℓ_p norm of the weighted sum of Gaussians. For $j \in [n]$, let $\mathbf{e}_j$ denote the jth standard basis element of $\mathbb{R}^n$. Then the jth coordinate of $(\mathbf{X}-\mathbf{C})\mathbf{z}$ is

$$\sum_{i\in[m]} \langle z_i \cdot (\mathbf{x}_i - \mathbf{c}_i), \mathbf{e}_j\rangle = \left\langle \mathbf{x}'-\mathbf{c}', (z_1\mathbf{e}_j,\ldots,z_m\mathbf{e}_j)\right\rangle = \|\mathbf{z}\|_2 \cdot \left\langle \mathbf{x}'-\mathbf{c}', \mathbf{w}_j\right\rangle, \tag{4}$$

where $\mathbf{w}_j \in \mathbb{R}^{nm}$ is the unit vector parallel to $(z_1\mathbf{e}_j,\ldots,z_m\mathbf{e}_j) \in \mathbb{R}^{nm}$ (this is where we use the fact that $\mathbf{z} \neq \mathbf{0}$).

Suppose $p \in [1,\infty)$. We then have

$$\begin{aligned}
&\mathrm{E}\left[\|(\mathbf{X}-\mathbf{C})\mathbf{z}\|_p\right] \\
&\quad= \|\mathbf{z}\|_2 \cdot \mathrm{E}\left[\left(\sum\nolimits_{j\in[n]} \left|\left\langle \mathbf{x}'-\mathbf{c}', \mathbf{w}_j\right\rangle\right|^p\right)^{1/p}\right] && \text{(Equation (4))} \\
&\quad\leq \|\mathbf{z}\|_2 \cdot \left(\sum\nolimits_{j\in[n]} \mathrm{E}\left[\left|\left\langle \mathbf{x}'-\mathbf{c}', \mathbf{w}_j\right\rangle\right|^p\right]\right)^{1/p} && \text{(Jensen's Inequality, linearity of E)} \\
&\quad\leq \|\mathbf{z}\|_2 \cdot \sqrt{1/\pi}\cdot(2en\cdot\Gamma(p/2+1))^{1/p} && \text{(Theorem 5.2, Inequality (3))} \\
&\quad\leq \|\mathbf{z}\|_2 \cdot c_p \cdot n^{1/p}. && \text{(constant } c_p \approx \sqrt{p})
\end{aligned}$$

Now suppose $p = \infty$. Here we use the tail inequality (Lemma 5.1) directly. By Equation (4),

$$\|(\mathbf{X}-\mathbf{C})\mathbf{z}\|_\infty = \|\mathbf{z}\|_2 \cdot \max_{j\in[n]} \left|\left\langle \mathbf{x}'-\mathbf{c}', \mathbf{w}_j\right\rangle\right|.$$

By Lemma 5.1 and Inequality (3), for every $j \in [n]$ we have

$$\Pr_{\mathbf{x}'\sim D_{\Lambda',\mathbf{c}'}}\left[\left|\left\langle \mathbf{x}'-\mathbf{c}', \mathbf{w}_j\right\rangle\right| \geq r\right] \leq 2e\cdot\exp(-\pi r^2).$$

The claim follows by applying the union bound over all $j \in [n]$. □

5.3.2 Proof of Tail Inequality

Our proof of the tail inequality mirrors the discussion from the overview in Section 5.1, where the main goal is to find a suitable function g that satisfies the two conditions.

Proof of Lemma 5.1. First, for any $r \geq 0$ define the positive function $g_r : \Lambda \to \mathbb{R}^+$ as:

$$g_r(\mathbf{x}) = \prod_{k\in[d]} \cosh(2\pi r \left\langle \mathbf{x}-\mathbf{c}, \mathbf{u}_k\right\rangle / d).$$

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

Claim 5.4. *For any $r \geq 0$,*

$$\sum_{\mathbf{x} \in \Lambda} \rho_{\mathbf{c}}(\mathbf{x}) \cdot g_r(\mathbf{x}) \quad \leq \quad \exp(\pi r^2/d) \cdot \rho(\Lambda).$$

Claim 5.5. *For any $r \geq 0$,*

$$\sum_{\mathbf{x} \in \Lambda \setminus Q_{r,\mathbf{c}}} \rho_{\mathbf{c}}(\mathbf{x}) \cdot g_r(\mathbf{x}) \quad \geq \quad \frac{\exp(2\pi r^2/d)}{2^d} \cdot \rho_{\mathbf{c}}(\Lambda \setminus Q_{r,\mathbf{c}}).$$

Then we see that

$$\begin{aligned}
\frac{\exp(2\pi r^2/d)}{2^d} \cdot \rho_{\mathbf{c}}(\Lambda \setminus Q_{r,\mathbf{c}}) \quad &\leq \quad \sum_{\mathbf{x} \in \Lambda \setminus Q_{r,\mathbf{c}}} \rho_{\mathbf{c}}(\mathbf{x}) \cdot g_r(\mathbf{x}) && \text{(Claim 5.5)} \\
&\leq \quad \sum_{\mathbf{x} \in \Lambda} \rho_{\mathbf{c}}(\mathbf{x}) \cdot g_r(\mathbf{x}) && (\rho,\ g_r \text{ positive}) \\
&\leq \quad \exp(\pi r^2/d) \cdot \rho(\Lambda). && \text{(Claim 5.4)}
\end{aligned}$$

Clearing the coefficient on the left completes the proof of Lemma 5.1. □

To conclude, it remains to justify the two claims.

Proof of Claim 5.4. We start by analyzing terms of the following form, which appear when we expand $g_r(\mathbf{x})$ according to its definition:

$$\begin{aligned}
&\rho_{\mathbf{c}}(\mathbf{x}) \cdot \exp\left(\textstyle\sum_{k \in [d]} 2\pi r \left\langle \mathbf{x} - \mathbf{c}, \pm \mathbf{u}_k \right\rangle / d\right) \\
&\quad = \rho_{\mathbf{c}}(\mathbf{x}) \cdot \exp\Big(2\pi \Big\langle \mathbf{x} - \mathbf{c}, \underbrace{\textstyle\sum_{k \in [d]} \pm \mathbf{u}_k r/d}_{\mathbf{c}'} \Big\rangle\Big) \\
&\quad = \exp\Big(-\pi\Big((\mathbf{x} - \mathbf{c})^2 - 2\left\langle \mathbf{x} - \mathbf{c}, \mathbf{c}' \right\rangle\Big)\Big) \\
&\quad = \exp\Big(-\pi\Big(\mathbf{x} - \underbrace{(\mathbf{c} + \mathbf{c}')}_{\mathbf{c}''}\Big)^2 + \pi(\mathbf{c}')^2\Big) && (5) \\
&\quad = \exp\left(\pi r^2/d\right) \cdot \exp\left(-\pi(\mathbf{x} - \mathbf{c}'')^2\right) && (6) \\
&\quad = \exp(\pi r^2/d) \cdot \rho_{\mathbf{c}''}(\mathbf{x})
\end{aligned}$$

Equation (5) is by completing the square. Equation (6) is by $(\mathbf{c}')^2 = \|\mathbf{c}'\|_2^2 = r^2/d$, regardless of the pattern of $\pm$'s, by orthonormality of $\{\mathbf{u}_k\}$.

We now analyze the expression that appears in the statement of Claim 5.4. Expanding the definition of g_r using $\cosh(x) = \frac{1}{2}(e^x + e^{-x})$, we see that the expression $\rho_{\mathbf{c}}(\mathbf{x}) \cdot g_r(\mathbf{x})$ contains 2^d terms of the form

$$\frac{1}{2^d} \cdot \rho_{\mathbf{c}}(\mathbf{x}) \cdot \prod_{k \in [d]} \exp\left(\pm 2\pi r \left\langle \mathbf{x} - \mathbf{c}, \mathbf{u}_k \right\rangle / d\right) = \frac{1}{2^d} \cdot \rho_{\mathbf{c}}(\mathbf{x}) \cdot \exp\left(\textstyle\sum_{k \in [d]} 2\pi r \left\langle \mathbf{x} - \mathbf{c}, \pm \mathbf{u}_k \right\rangle / d\right),$$

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

$$\frac{\exp(\pi r^2/d)}{2^d} \cdot \rho_{\mathbf{c}''}(\Lambda) \leq \frac{\exp(\pi r^2/d)}{2^d} \cdot \rho(\Lambda),$$

where the inequality is due to Lemma 2.8. Combining all 2^d terms, the claim follows. □

Proof of Claim 5.5. By the definition of g_r and the inequality $\cosh(x) \geq \frac{1}{2}\exp(|x|)$, we have

$$g_r(\mathbf{x}) \geq \frac{1}{2^d} \prod_{k \in [d]} \exp\left(2\pi r \left|\langle \mathbf{x} - \mathbf{c}, \mathbf{u}_k \rangle / d\right|\right) = \frac{1}{2^d} \cdot \exp\left(2\pi r \left\|\mathbf{x} - \mathbf{c}\right\|_{\mathbf{U}} / d\right).$$

Then because $\|\mathbf{x} - \mathbf{c}\|_{\mathbf{U}} \geq r$ for any $\mathbf{x} \in \Lambda \setminus Q_{r,\mathbf{c}}$, and by positivity ρ, the claim follows. □

6 Worst-Case to Average-Case Reductions

Here we use the results from the previous section (namely, Theorem 5.2 and Corollary 5.3) to provide an analysis in ℓ_p norms of two prior worst-case to average-case reductions that rely on Gaussian measures. The first, due to Micciancio and Regev [MR07], shows that finding "small" nonzero solutions to *random homogeneous linear systems* over $\mathbb{Z}_q$ (for an appropriate choice of modulus q) is as hard as approximating several worst-case lattice problems in the ℓ_2 norm to within $\tilde{O}(n)$ factors. We extend this result to all ℓ_p norms, $p \in [2, \infty]$, maintaining essentially the same approximation factors.

The second reduction, due to Regev [Reg05], shows that solving the "*learning with errors*" (LWE) problem on the average (under a Gaussian-like error distribution) is as hard as approximating worst-case lattice problems in the ℓ_2 norm to within factors as small as $\tilde{O}(n)$ *for quantum algorithms*. We also extend this result to all ℓ_p norms, $p \in [2, \infty]$, for essentially the same approximation factors.

We remark that our results were also used to analyze a worst-case to average-case reduction based on so-called *ideal* lattices having special algebraic structure [PR07], and that it appears possible to do the same for related reductions and the cryptographic schemes based upon them [Mic07, PR06, LM06, LM08].

6.1 Random Homogeneous Linear Systems

Our exposition in this section closely follows the organization of Section 5 in [MR07]. Our goal is to reduce worst-case lattice problems in ℓ_p norms to the average-case problem of finding "small" nonzero solutions to random homogeneous linear systems over $\mathbb{Z}_q$, i.e., modulo q. This problem goes all the way back to Ajtai's seminal work [Ajt04]; we use the following definition from [MR07]:

Definition 6.1. The *small integer solution* problem (in the ℓ_2 norm), denoted SIS, is the following: for an integer q, matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and real $\beta > 0$, find a nonzero integer vector $\mathbf{z} \in \mathbb{Z}^m \setminus \{\mathbf{0}\}$ such that $\mathbf{A}\mathbf{z} = \mathbf{0} \bmod q$ and $\|\mathbf{z}\|_2 \leq \beta$.

For functions $q(n)$, $m(n)$, and $\beta(n)$, the average-case problem $\mathsf{SIS}_{q,m,\beta}$ is defined to be 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 $\beta \geq \sqrt{m} \cdot q^{n/m}$, a simple pigeonhole argument implies that any instance of $\mathsf{SIS}_{q,m,\beta}$ always has a nonzero solution [MR07, Lemma 5.2]. We implicitly take $\beta = \sqrt{m} \cdot q^{n/m}$ when it is otherwise left unspecified.

The SIS problem is related to the shortest vector problem SVP on a suitably defined family of random lattices, and its average-case hardness immediately implies a collection of one-way and collision-resistant cryptographic hash functions (see, e.g., [MR07, Section 5.1]). More recently, SIS has been used as the foundation for a "direct" construction of lattice-based trapdoor functions and "hash-and-sign" signature schemes [GPV08].

Instead of reducing directly from, say, SIVP in the worst-case to SIS on the average, prior works have introduced an intermediate (worst-case) lattice problem. The intermediate problem is at least as hard as other well-studied lattice problems (such as SIVP), and yet also admits a reduction to the relevant average-case problem. This serves two purposes: first, the reduction from the intermediate problem better reflects the essence of the main technique and admits a simpler probabilistic analysis. Second, the main results are more modular, because there are elementary (worst-case to worst-case) reductions from several standard lattice problems to the intermediate problem. In [MR07], the intermediate problem is called *incremental guaranteed distance decoding* (a variation of the problem first introduced in [Mic07]). As with other lattice problems, it can be defined relative to any norm $\|\cdot\|$, and we affix a superscript p to indicate the ℓ_p norm.

Definition 6.2 (Incremental Guaranteed Distance Decoding)**.** An input to $\mathsf{IncGDD}^{\phi}_{\gamma,g}$ is a tuple $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ where $\mathbf{B}$ is a basis of an n-dimensional lattice, $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$ is a set of n linearly independent lattice vectors, $\mathbf{t} \in \mathbb{R}^n$ is a target vector, and $r > \gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$ is a real number. The goal is to output a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{s} - \mathbf{t}\| \leq \|\mathbf{S}\| / g + r$.

Typically, g is some small constant (say, 8), and ϕ is some parameter of the lattice (say, the smoothing parameter η_ϵ or the nth successive minimum λ_n). Informally, the goal of IncGDD is to decode an arbitrary target vector $\mathbf{t}$ to a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ that is within a distance not much larger than $\|\mathbf{S}\| / g$. Intuitively, this task appears difficult because the polynomial-time nearest plane algorithm [Bab86] (which is the standard algorithm for decoding on lattices) is only guaranteed to produce a lattice vector within distance $(\sqrt{n}/2) \|\mathbf{S}\|_2$ (in the ℓ_2 norm) of the target.

In [MR07], the core worst-case to average-case reduction is from IncGDD to SIS, and is described in Theorem 5.9 of that work. An examination of its proof (specifically, the analysis of the reduction's success event) reveals the following specific details about the reduction.

Proposition 6.3. *Let $\|\cdot\|$ denote any norm on $\mathbb{R}^n$, let $g(n), \gamma(n) > 0$, let $\epsilon(n)$ be negligible in n, and let $m(n), \beta(n) = \mathrm{poly}(n)$, $q(n) \geq n \cdot g(n) \cdot \beta(n) \cdot \sqrt{m(n)}$. Let $\mathcal{F}$ be an oracle that solves $\mathsf{SIS}_{q,m,\beta}$ (in the ℓ_2 norm) on the average with non-negligible probability.*

There is a probabilistic polynomial time oracle algorithm (making a single call to $\mathcal{F}$) that, given an instance $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ of $\mathsf{IncGDD}^{\eta_\epsilon}_{\gamma,g}$, and conditioned on an event that occurs with non-negligible probability, outputs a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ of the form $\mathbf{s} = \mathbf{t} + \mathbf{u} - (\mathbf{X} - \mathbf{C})\mathbf{z}$, where:[6]

1. $\|\mathbf{u}\| \leq \|\mathbf{S}\| / g(n)$,
2. $\mathbf{z} \in \mathbb{Z}^m$ *and* $\|\mathbf{z}\|_2 \leq \beta$,
3. $\mathbf{X}, \mathbf{C} \in \mathbb{R}^{n \times m}$ *and the columns* $\mathbf{x}_i$ *are distributed independently according to* $D_{\mathcal{L}(\mathbf{B}),s,\mathbf{c}_i}$, *for* $s = 2r/\gamma(n) \geq 2\eta_\epsilon(\mathcal{L}(\mathbf{B}))$.

[6] For the reader interested in checking these claims against the proof of [MR07, Theorem 5.9], we have condensed some notation, using $\mathbf{u}$ in place of $\mathbf{x} - \mathbf{C}\mathbf{z}$, and $\mathbf{C}$ in place of $\mathbf{C} + \mathbf{T}$.

We remark that Condition 1 is a consequence of Lemma 5.8 in [MR07], which is stated only in terms of the ℓ_2 norm. However, the relevant part of its proof relies only on the basic norm axioms, and therefore applies for any norm. We also remark that a recent work [GPV08] has given a simplified and slightly tighter form of Proposition 6.3, in which the $\mathbf{u}$ term is removed and the modulus $q(n)$ is smaller.

Using Proposition 6.3, we can now show that solving SIS on the average (in the ℓ_2 norm) is as hard as solving IncGDD in the ℓ_p norm.

Theorem 6.4. *For any $p \in [1, \infty)$, there is a constant c_p such that the following holds: for any $g(n) > 0$, negligible $\epsilon(n)$, polynomially-bounded $m(n), \beta(n) = \mathrm{poly}(n)$, and $q(n) \geq n \cdot g(n) \cdot \beta(n)\sqrt{m(n)}$, there is a probabilistic polynomial time reduction from solving $\mathsf{IncGDD}^{p,\eta_\epsilon}_{\gamma,g}$ in the worst case for $\gamma(n) = 4c_p n^{1/p} \cdot \beta(n)$ to solving $\mathsf{SIS}_{q,m,\beta}$ (in the ℓ_2 norm) on the average with non-negligible probability.*

For $p = \infty$, the same holds for $\gamma(n) = 2\beta(n)\sqrt{\log n}$.

Proof. The reduction is exactly the algorithm described in Proposition 6.3; all that remains is the analysis. Let $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ be the input IncGDD instance. We show that with non-negligible probability (over all the randomness of the reduction and its oracle), the reduction outputs an $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{s} - \mathbf{t}\|_p \leq \|\mathbf{S}\|_p / g(n) + r$ (by a standard repetition argument, the probability of success can then be made negligibly close to 1). Proposition 6.3 states that, conditioned on an event E that occurs with non-negligible probability, the output is some $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ of the form $\mathbf{s} = \mathbf{t} + \mathbf{u} - (\mathbf{X} - \mathbf{C})\mathbf{z}$ satisfying the three conditions. In particular, by the triangle inequality,

$$\|\mathbf{s} - \mathbf{t}\|_p \leq \|\mathbf{u}\|_p + \|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p \leq \|\mathbf{S}\|_p / g(n) + \|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p ,$$

where the inequality is due to Condition 1. Therefore it suffices to show that, conditioned on any fixed values of $\mathbf{C}$ and $\mathbf{z}$ that may arise when event E occurs, $\|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p \leq r$ with some non-negligible probability (say, at least $1/2$).

Recall that by Condition 2, $\|\mathbf{z}\|_2 \leq \beta(n)$, and by Condition 3, the $\mathbf{x}_i \sim D_{\mathcal{L}(\mathbf{B}),s,\mathbf{c}_i}$ are independent for $s = 2r/\gamma(n) \geq 2\eta_\epsilon(\mathcal{L}(\mathbf{B}))$.

Suppose $p \in [1, \infty)$, and recall that $\gamma(n) = 4c_p n^{1/p} \cdot \beta(n)$. Applying Corollary 5.3, we have

$$\mathrm{E}\left[\|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p\right] \leq (2r/\gamma(n)) \cdot \beta(n) \cdot c_p n^{1/p} = r/2.$$

Then by Markov's inequality, it follows that $\|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p \leq r$ except with probability at most $1/2$, as desired.

Now suppose $p = \infty$, and recall that $\gamma(n) = 2\beta(n) \cdot \sqrt{\log n}$. Applying Corollary 5.3, we have

$$\Pr\left[\|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p \geq r\right] \leq \Pr\left[\|(\mathbf{X} - \mathbf{C})\mathbf{z}\|_p \geq s \|\mathbf{z}\|_2 \cdot \sqrt{\log n}\right] \leq 2e/n^{\pi - 1},$$

which is at most $1/2$ for all sufficiently large n. □

6.1.1 Connection to Other Worst-Case Problems

As shown in Lemmas 5.10 through 5.12 of [MR07] (using some reductions from [Mic07]), the IncGDD problem is at least as hard as several other lattice problems, via straightforward worst-case to worst-case reductions. However, in some cases their analysis is specific to the ℓ_2 norm. (Specifically,

they rely on the fact that a vector of ℓ_2 norm r orthogonal to a subspace $H \subset \mathbb{R}^n$ is at ℓ_2 distance r from H, which is not true in all norms.) In order to generalize the reductions to arbitrary norms, we use the following lemma. Technically, the second claim of the lemma assumes that the norm satisfies a mild boundedness condition and is efficiently computable (in particular, all ℓ_p norm satisfy these requirements). For the remainder of the paper, we implicitly restrict our attention to such norms.

Lemma 6.5. *For any norm $\|\cdot\|$ on $\mathbb{R}^n$ and any proper subspace $H \subset \mathbb{R}^n$, there is an $\mathbf{x}_H \in \mathbb{R}^n$ such that $\|\mathbf{x}_H\| = 1$ and $\operatorname{dist}(\mathbf{x}_H, H) = 1$.*

Furthermore, there is a polynomial-time algorithm for finding such an $\mathbf{x}_H$ given a basis for H.

Proof. Let $\mathbf{v} \in \mathbb{R}^n$ be any point such that $\operatorname{dist}(\mathbf{v}, H) = 1$. (Such a point always exists because for any $\mathbf{w} \notin H$ we have $\operatorname{dist}(\mathbf{w}, H) > 0$, so it suffices to let $\mathbf{v} = \mathbf{w}/\operatorname{dist}(\mathbf{w}, H)$.) Let $\mathbf{u} \in H$ be such that $\|\mathbf{v} - \mathbf{u}\| = 1$. We claim that $\mathbf{x}_H = \mathbf{v} - \mathbf{u}$ serves to prove the claim. We have already established that $\|\mathbf{x}_H\| = 1$. Now if $\operatorname{dist}(\mathbf{x}_H, H) < 1$, then there exists some $\mathbf{t} \in H$ such that $\|\mathbf{x}_H - \mathbf{t}\| = \|\mathbf{v} - (\mathbf{u} + \mathbf{t})\| < 1$, and since $\mathbf{u} + \mathbf{t} \in H$, we have $\operatorname{dist}(\mathbf{v}, H) < 1$, a contradiction.

For the second claim, note that the above description of $\mathbf{x}_H$ is constructive in the following sense: if, given a basis for H and any $\mathbf{v} \in \mathbb{R}^n$, we can efficiently find some $\mathbf{u} \in H$ such that $\|\mathbf{v} - \mathbf{u}\| = \operatorname{dist}(\mathbf{v}, H)$, then we can compute the desired $\mathbf{x}_H$. Finding such $\mathbf{u}$ reduces to minimizing the convex function $f(\mathbf{u}) = \|\mathbf{v} - \mathbf{u}\|$ subject to the constraint $\mathbf{u} \in H' \subset H$, where H' is a sufficiently large convex bounded subset of H. For this purpose we may use any polynomial-time algorithm for convex optimization, e.g., the ellipsoid or interior point methods. See, e.g., [NN94] for details on the interior point method, including a formal treatment of the computability and boundedness conditions. □

We also need the following lower bound on the covering radius in an arbitrary norm. Our proof is an adaptation of the proof (for the ℓ_2 norm) of Theorem 7.9 in [MG02].

Lemma 6.6. *In any norm and for any n-dimensional lattice Λ, $\mu(\Lambda) \geq \lambda_n(\Lambda)/2$.*

Proof. Write $\mu = \mu(\Lambda)$ and $\lambda_n = \lambda_n(\Lambda)$. Suppose for contradiction that $\lambda_n > 2\mu$, and let $\epsilon > 0$ be such that $\epsilon < \lambda_n - 2\mu$. To obtain a contradiction, we iteratively construct a set of linearly independent lattice vectors $\mathbf{s}_1, \ldots, \mathbf{s}_n \in \Lambda$ such that $\|\mathbf{s}_i\| < \lambda_n$ for all $i \in [n]$.

For any $i \in [n]$, let $H_{i-1} = \operatorname{span}(\mathbf{s}_1, \ldots, \mathbf{s}_{i-1})$. Let $\mathbf{t}_i = (\mu + \epsilon) \cdot \mathbf{x}_{H_{i-1}}$, where $\mathbf{x}_{H_{i-1}}$ is the point guaranteed by Lemma 6.5 for subspace H_{i-1}. Therefore we have $\|\mathbf{t}_i\| = (\mu + \epsilon)$ and $\operatorname{dist}(\mathbf{t}_i, H_{i-1}) = (\mu + \epsilon)$. Let $\mathbf{s}_i \in \Lambda$ be such that $\|\mathbf{s}_i - \mathbf{t}_i\| \leq \mu$. Then $\mathbf{s}_i \notin H_{i-1}$, because by the triangle inequality, the distance from $\mathbf{s}_i$ to H_{i-1} is

$$\operatorname{dist}(\mathbf{s}_i, H_{i-1}) \geq \operatorname{dist}(\mathbf{t}_i, H) - \|\mathbf{s}_i - \mathbf{t}_i\| = (\mu + \epsilon) - \|\mathbf{s}_i - \mathbf{t}_i\| \geq \epsilon > 0.$$

Therefore, $\mathbf{s}_i$ is linearly independent from $\mathbf{s}_1, \ldots, \mathbf{s}_{i-1}$. Furthermore, by the triangle inequality, $\|\mathbf{s}_i\| \leq \|\mathbf{t}_i\| + \|\mathbf{s}_i - \mathbf{t}_i\| \leq (\mu + \epsilon) + \mu < \lambda_n$. By induction on i, the proof is complete. □

We can now generalize Lemmas 5.10 through 5.12 from [MR07] to hold for any norm. For completeness, we include sketches of their proofs, which are very similar (or identical) to the originals.

Lemma 6.7. *For any $\gamma(n) \geq 1$, any ϕ, and any norm $\|\cdot\|$, there is a deterministic polynomial-time reduction from $\mathsf{GIVP}^{\phi}_{8\gamma}$ to $\mathsf{IncGDD}^{\phi}_{\gamma,8}$.*

Proof. Given a basis $\mathbf{B}$, we use an iterative process to produce n linearly independent vectors $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{S}\| \leq 8\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$. Let $\mathbf{S} = \mathbf{B}$ initially. At each iteration, choose some longest vector $\mathbf{s}_i \in \mathbf{S}$. Compute a vector $\mathbf{t} = \mathbf{x}_H \cdot \|\mathbf{S}\|/2$, where $H = \mathrm{span}(\mathbf{s}_1, \ldots, \mathbf{s}_{i-1}, \mathbf{s}_{i+1}, \ldots, \mathbf{s}_n)$ and $\mathbf{x}_H$ is the point guaranteed by Lemma 6.5. Then apply the IncGDD oracle on the instance $(\mathbf{B}, \mathbf{S}, \mathbf{t}, \|\mathbf{S}\|/8)$. If it fails, abort and output $\mathbf{S}$. Otherwise, we obtain a lattice vector $\mathbf{u} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{t} - \mathbf{u}\| \leq \|\mathbf{S}\|/8 + \|\mathbf{S}\|/8 = \|\mathbf{S}\|/4$, which we substitute for $\mathbf{s}_i$ and continue the process.

Note that by the triangle inequality, we have $\mathrm{dist}(\mathbf{u}, H) \geq \mathrm{dist}(\mathbf{t}, H) - \|\mathbf{t} - \mathbf{u}\| \geq \|\mathbf{S}\|/2 - \|\mathbf{S}\|/4 > 0$, so $\mathbf{u} \notin H$ and therefore $\mathbf{S}$ always contains n linearly independent vectors. When the oracle fails, by definition of IncGDD it must be the case that $\|\mathbf{S}\| \leq 8\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$, as desired. Finally, to analyze the runtime we observe that $\|\mathbf{u}\| \leq \|\mathbf{t}\| + \|\mathbf{S}\|/8 \leq 3\|\mathbf{S}\|/4$, so $\sum_i \log \|\mathbf{s}_i\|$ decreases by a constant with every iteration. Because $\|\cdot\|$ is efficiently computable, the initial value of $\sum_i \log \|\mathbf{s}_i\|$ is polynomial in the input size. □

Lemma 6.8. *For any $\gamma(n) \geq 1$, any ϕ, and any norm $\|\cdot\|$, there is a deterministic polynomial-time reduction from $\mathsf{GDD}^{\phi}_{3\gamma}$ to $\mathsf{IncGDD}^{\phi}_{\gamma,8}$.*

Proof. The proof is identical to the proof of Lemma 5.11 in [MR07]. Given a basis $\mathbf{B}$ and a vector $\mathbf{t}$, the goal is to find a vector $\mathbf{u} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{t} - \mathbf{u}\| \leq 3\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$. First, we apply the reduction claimed by Lemma 6.7 to obtain a set of n linearly independent lattice vectors $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{S}\| \leq 8\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$. We then apply a binary search to find a value r for which an oracle call on $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r/2)$ fails, but an oracle call on $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ succeeds. Because the former call fails, it must be the case that $r \leq 2\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$. The latter call yields a lattice vector $\mathbf{u} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{t} - \mathbf{u}\| \leq \|\mathbf{S}\|/8 + r \leq \gamma(n)\phi(\mathcal{L}(\mathbf{B})) + 2\gamma(n)\phi(\mathcal{L}(\mathbf{B})) = 3\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$, as desired. □

Lemma 6.9. *For any $\gamma(n) \geq 1$ and any norm $\|\cdot\|$, there is a randomized polynomial-time reduction from GapCRP_{γ} to $\mathsf{GDD}^{\lambda_n}_{\gamma/4}$.*

Proof. The proof is identical to the proof of Lemma 5.12 in [MR07]. Given a basis $\mathbf{B}$, the goal is to decide whether $\mu(\mathcal{L}(\mathbf{B})) \leq 1$ or $\mu(\mathcal{L}(\mathbf{B})) > \gamma$. The reduction chooses a point $\mathbf{t} \in \mathcal{P}(\mathbf{B})$ uniformly at random, then calls the GDD oracle on the instance $(\mathbf{B}, \mathbf{t})$ to obtain a lattice vector $\mathbf{x} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{t} - \mathbf{x}\| \leq (\gamma/4)\lambda_n(\mathcal{L}(\mathbf{B}))$. If $\|\mathbf{t} - \mathbf{x}\| \leq \gamma/2$ the reduction accepts, otherwise it rejects.

Suppose the input is a YES instance, i.e., $\mu(\mathcal{L}(\mathbf{B})) \leq 1$. Then $\|\mathbf{t} - \mathbf{x}\| \leq \gamma\lambda_n(\mathcal{L}(\mathbf{B}))/4 \leq \gamma\mu(\mathcal{L}(\mathbf{B}))/2 \leq \gamma/2$, where we have used Lemma 6.6 for the second inequality. Therefore, YES instances are always accepted by the reduction. Now suppose the input is a NO instance, i.e., $\mu(\mathcal{L}(\mathbf{B})) > \gamma$. By Lemma 4.1 of [GMR05] (which holds for any norm), a random $\mathbf{t}$ chosen as above satisfies $\mathrm{dist}(\mathbf{t}, \mathcal{L}(\mathbf{B})) \geq \mu(\mathcal{L}(\mathbf{B}))/2$ with probability at least $1/2$. Therefore, $\mathrm{dist}(\mathbf{t}, \mathcal{L}(\mathbf{B})) > \gamma/2$ with probability at least $1/2$, and the instance is rejected. □

From now on, we consider certain "good" choices of the parameters $m(n)$ and $q(n)$, and fix $\beta(n) = \sqrt{m(n)} \cdot q(n)^{n/m(n)}$ to guarantee the existence of solutions to $\mathsf{SIS}_{q,m,\beta}$.

Corollary 6.10. *For any $p \in [1, \infty)$ and any $m(n) = \Theta(n \log n)$, there exist some $q(n) = O(n^2 \log n)$ and $\gamma(n) = O(n^{1/2+1/p} \cdot \sqrt{\log n})$ such that for any negligible function $\epsilon(n)$, solving $\mathsf{SIS}_{q,m}$ on the average with non-negligible probability is at least as hard as solving both $\mathsf{GIVP}^{p,\eta_\epsilon}_{\gamma}$ and $\mathsf{GDD}^{p,\eta_\epsilon}_{\gamma}$ in the worst case.*

For $p = \infty$, the same holds for some $\gamma(n) = O(n^{1/2} \log n)$.

Proof. Let $g(n) = 8$; then there is a $q(n) = O(n^2 \log n)$ that satisfies the conditions of our Theorem 6.4. The claim follows by Lemmas 6.7 and 6.8. □

Finally, by connecting the smoothing parameter η_ϵ to the nth successive minimum λ_n in the ℓ_p norm for $p \in [2, \infty]$, we obtain the following end result.

Theorem 6.11. *For any $p \in [2, \infty)$ and any $m(n) = \Theta(n \log n)$, there exists some $q(n) = O(n^2 \log n)$ such that for any function $\gamma(n) = \omega(n \log n)$, solving $\mathsf{SIS}_{q,m}$ on the average with non-negligible probability is as hard as solving the following problems in the worst case: SIVP^p_γ, $\mathsf{GDD}^{p,\lambda_n}_\gamma$, and GapCRP^p_γ.*

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

Proof. Let $\gamma'(n) = O(n^{1/2+1/p} \cdot \sqrt{\log n})$ be the function guaranteed by Corollary 6.10 for which we can solve $\mathsf{GIVP}^{p,\eta_\epsilon}$ and $\mathsf{GDD}^{p,\eta_\epsilon}$ in the worst case to within $\gamma'(n)$ factors. Define

$$\alpha(n) = \frac{\gamma(n)}{\gamma'(n) \cdot n^{1/2-1/p}} = \frac{\gamma(n)}{O(n \cdot \sqrt{\log n})} = \omega(\sqrt{\log n}).$$

Then by Lemma 2.7, there is a negligible $\epsilon(n)$ such that $\eta_\epsilon(\Lambda) \leq \lambda^{(p)}_n(\Lambda) \cdot n^{1/2-1/p} \cdot \alpha(n)$.

Suppose $p \in [2, \infty)$. By the above bound on η_ϵ relative to $\lambda^{(p)}_n$, this implies that we can solve SIVP^p and $\mathsf{GDD}^{p,\lambda_n}$ to within $\gamma'(n) \cdot n^{1/2-1/p} \cdot \alpha(n) = \gamma(n)$ factors, as desired. (For $p = \infty$, an identical argument uses the $\gamma'(n) = O(n^{1/2+1/p} \cdot \log n)$ guaranteed by Corollary 6.10.)

The result for GapCRP follows by combining the result for GDD with Lemma 6.9. □

Remark on the case $1 \leq p < 2$. We point out that Theorem 6.4 and Corollary 6.10 apply for all ℓ_p norms, $1 \leq p \leq \infty$. The difficulty in obtaining $\tilde{O}(n)$ approximation factors for $p < 2$ arises from the relationship between the smoothing parameter η_ϵ and $\lambda^{(p)}_n$. Unlike the case $p \geq 2$, we cannot conclude that $\eta_\epsilon \leq \tilde{O}(n^{1/2-1/p}) \cdot \lambda^{(p)}_n$. Instead, the best bound we can obtain is $\eta_\epsilon \approx \lambda^{(p)}_n$, which yields an overall approximation factor of $\gamma(n) = \tilde{O}(n^{1/2+1/p}) = \tilde{O}(n^{3/2})$ for SIVP and the other problems above.

6.1.2 Connection to Shortest Vector Problem

The above results do not immediately imply a reduction that solves the shortest vector problem in the worst case (even its decision version). This is because the minimum distance λ_1 of a lattice may be significantly smaller than the smoothing parameter η_ϵ, but the reduction from Theorem 6.4 may "stop working" once the Gaussian parameter drops below η_ϵ. For the worst-case problem GapSVP, Theorem 5.23 of [MR07] applies the core worst-case to average-case reduction to the *dual* lattice. Using additional techniques from [AR05], the theorem demonstrates that solving SIS on the average is at least as hard as solving GapSVP_γ in the ℓ_2 norm (in the worst case) for some $\gamma(n) = O(n\sqrt{\log n})$ factor. A close examination of the reduction and its analysis yields the following more specific restatement of the theorem.

Proposition 6.12. *Let $\|\cdot\|$ denote any norm on $\mathbb{R}^n$, let $\gamma'(n) \geq 1$, let $q(n), m(n), \beta(n) = \mathrm{poly}(n)$ with* odd *$q(n) \geq 4\sqrt{m(n)} \cdot n^{1.5} \cdot \beta(n)$, and let $\mathcal{F}$ be an oracle that solves $\mathsf{SIS}_{q,m,\beta}$ on the average with non-negligible probability $\delta(n)$ (where the probability is taken over the choice of the input SIS instance and $\mathcal{F}$'s internal randomness).*

There is a probabilistic oracle algorithm running in time $\mathrm{poly}(n)/\delta$ *that, given access to* $\mathcal{F}$ *and an input basis* $\mathbf{B}$ *for an* n*-dimensional lattice* $\Lambda = \mathcal{L}(\mathbf{B})$*, has the following behavior:*

- *If* $\lambda_1^{(2)}(\Lambda) \leq 1$*, the algorithm always accepts. (Note that the minimum distance here is measured in the* ℓ_2 *norm.)*
- *If* $\lambda_1(\Lambda) > \gamma'$*, the algorithm rejects with overwhelming probability, provided that all the conditions on* Λ *stated below hold. (Note that the minimum distance here is measured in the given norm* $\|\cdot\|$*.)*

The conditions on Λ *are as follows:*

1. $s \geq 2\eta_\epsilon(\Lambda^*)$ *for some negligible function* $\epsilon(n)$*, where* s *is a Gaussian parameter used by the algorithm (whose value can be computed from the input as desired).*
2. $s \cdot \beta(n) \leq 1/8\pi$*.*
3. *For any* $\mathbf{w}, \mathbf{c} \in \mathbb{R}^n$ *and any* $\mathbf{v} \in \mathbb{R}^n$ *such that* $\mathrm{dist}(\mathbf{v}, \Lambda) > \gamma'$*, it is the case that*

$$\left| \mathop{\mathrm{E}}_{\mathbf{x} \sim D_{\Lambda^*, s, \mathbf{c}}} [\cos(2\pi \langle \mathbf{x} + \mathbf{w}, \mathbf{v} \rangle)] \right| \leq 1/3.$$

For the reader interested in checking this restatement against the original Theorem 5.23, we have made the following observations and changes (the uninterested reader can safely skip this paragraph). We have fixed the parameter $d = 1$ without loss of generality (by scaling the input instance) to match our definition of GapSVP. Our acceptance condition ($\lambda_1^{(2)}(\Lambda) \leq 1$) is exactly the same as in the original proof. The rejection condition generalizes the original analysis for NO instances, which was done in three parts. First, it is shown that "N calls to $\mathcal{W}$ succeed" assuming that $s \geq \eta_\epsilon(\Lambda^*)$ for $\epsilon = 2^{-n}$, but it is clear that any negligible $\epsilon(n)$ suffices (this is our Condition 1). Second, it is shown that "$\Pr[f_{\mathbf{W}}(\mathbf{t}) \geq 1/2]$ is small" by giving a 2^{-n+1} upper bound on the expression from our Condition 3 and applying Hoeffding's bound; it is clear that a $1/3$ bound suffices. Third, it is shown that "test (c) [the eigenvalue test] is satisfied with high probability," using Lemma 5.20 and the fact that $s\beta \leq 1/8\pi$ (our Condition 2). Proving the hypotheses of Lemma 5.20 is done using only the fact that $s \geq 2\eta_\epsilon(\Lambda^*)$ and other facts that are independent of the given norm $\|\cdot\|$.

Using Proposition 6.12, we obtain a reduction that solves GapSVP_γ in the ℓ_p norm for any $\gamma(n) = \omega(n \log n)$ factor. Note that this is an $\omega(\sqrt{\log n})$ factor looser than the reduction obtained in Theorem 5.23 of [MR07]. This is due entirely to the slightly looser bounds on the smoothing parameter relative to the dual minimum distance in the ℓ_p norm versus the ℓ_2 norm (see the discussion following Lemma 3.5).

Theorem 6.13. *For any* $p \in [2, \infty]$*, any* $q(n), m(n), \beta(n) = \mathrm{poly}(n)$ *with* odd $q(n) \geq 4\sqrt{m(n)} \cdot n^{1.5} \cdot \beta(n)$*, and any* $\gamma(n) = 16\pi\sqrt{n} \cdot \beta(n) \cdot \omega(\sqrt{\log n})$*, solving* $\mathsf{SIS}_{q,m,\beta}$ *on the average with non-negligible probability is as hard as solving* GapSVP^p_γ *in the worst case.*

In particular, for any $m(n) = \Theta(n \log n)$*, there exists an odd* $q(n) = O(n^{2.5} \log n)$ *such that for any function* $\gamma(n) = \omega(n \log n)$*, solving* $\mathsf{SIS}_{q,m}$ *on the average is as hard as solving* GapSVP^p_γ*.*

Proof. Let $\Lambda = \mathcal{L}(\mathbf{B})$, where basis $\mathbf{B}$ is the input instance of GapSVP. By scaling, without loss of generality we can assume that YES instances are such that $\lambda_1^{(p)}(\Lambda) \leq n^{1/p-1/2}$, and NO instances are such that $\lambda_1^{(p)}(\Lambda) > \gamma'(n) \stackrel{\text{def}}{=} \gamma(n) \cdot n^{1/p-1/2} = 16\pi n^{1/p} \cdot \beta(n) \cdot \alpha(n)$ where $\alpha(n) = \omega(\sqrt{\log n})$.

The reduction is the one claimed by Proposition 6.12, using a Gaussian parameter

$$s = \frac{2n^{1/p} \cdot \alpha(n)}{\gamma'(n)}.$$

If $\mathbf{B}$ is a YES instance, then by the properties of ℓ_p norms, we have $\lambda_1^{(2)}(\Lambda) \leq n^{1/2-1/p} \cdot \lambda_1^{(p)}(\Lambda) \leq 1$, so the reduction always accepts. If $\mathbf{B}$ is a NO instance, then $\lambda_1^{(p)}(\Lambda) > \gamma'(n)$, so the reduction rejects (with overwhelming probability) as long as the conditions in Proposition 6.12 are satisfied. We now show that this is the case.

For Condition 1, observe that by Lemma 3.5, there exists a negligible function $\epsilon(n)$ such that

$$\eta_\epsilon(\Lambda^*) \leq \frac{n^{1/p} \cdot \alpha(n)}{\lambda_1^{(p)}(\Lambda)} < \frac{n^{1/p} \cdot \alpha(n)}{\gamma'(n)} < s/2,$$

so $s > 2\eta_\epsilon(\Lambda^*)$, as desired. For Condition 2, by definition of s and $\gamma'(n)$ we immediately obtain $s \cdot \beta(n) = 1/(8\pi)$. For Condition 3, we apply Lemma 2.9. Note that by definition of s, we have $\gamma'(n) \geq \omega(n^{1/p}/s)$, and in particular, $\gamma'(n) \geq c_p n^{1/p}/s$ for any positive constant c_p and all large enough n. Now if $\text{dist}(\mathbf{v}, \Lambda) > \gamma' \geq c_p n^{1/p}/s$, then $\text{dist}(s\mathbf{v}, s\Lambda) > c_p n^{1/p}$, so by Corollary 3.2 we have

$$\frac{\rho_{1/s}(\Lambda - \mathbf{v})}{\rho_{1/s}(\Lambda)} = \frac{\rho(s\Lambda - s\mathbf{v})}{\rho(s\Lambda)} < 1/4.$$

Then by Lemma 2.9, we see that Condition 3 is satisfied. □

6.2 Learning With Errors

Regev defined the *learning with errors* (LWE) problem (a generalization to larger moduli of the "learning parity with noise" problem), and showed that LWE is hard on the average unless there are efficient *quantum* algorithms for approximating the worst-case problems SIVP and GapSVP in the ℓ_2 norm [Reg05]. (We remark that there are no known quantum algorithms that outperform the best known classical algorithms for worst-case lattice problems.) Regev used LWE to construct a public-key cryptosystem, and more recently, LWE has served as the foundation for several other cryptographic schemes, including chosen ciphertext-secure cryptosystems [PW08], oblivious transfer protocols [PVW07], and identity-based encryption [GPV08].

The essence of the worst-case to average-case reduction from [Reg05] is a *quantum* strategy that, given an oracle for solving LWE on the average, generates samples from the discrete Gaussian $D_{\Lambda,s}$. The process continues for iteratively smaller values of s, all the way down to a value of s that can be as small as $\Theta(\sqrt{n}) \cdot \eta_\epsilon(\Lambda)$.

We now give some background for defining the LWE problem and stating the quantum reduction. Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ be the additive group on the interval $[0, 1)$ with modulo 1 addition. For positive integers n and $q \geq 2$, a vector $\mathbf{s} \in \mathbb{Z}_q^n$, and a probability distribution χ on $\mathbb{T}$, define $A_{\mathbf{s},\chi}$ as the distribution on $\mathbb{Z}_q^n \times \mathbb{T}$ obtained by choosing a vector $\mathbf{a} \in \mathbb{Z}_q^n$ uniformly at random, choosing an error term $e \in \mathbb{T}$ according to χ, and outputting $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle / q + e)$, where the addition is performed in $\mathbb{T}$.

Definition 6.14. For an integer function $q = q(n)$ and an error distribution χ on $\mathbb{T}$, the goal of the (worst-case) $\mathsf{LWE}_{q,\chi}$ problem is to find $\mathbf{s}$, given access to samples from $A_{\mathbf{s},\chi}$ for some $\mathbf{s} \in \mathbb{Z}_q^n$.

As described in [Reg05, Section 4], there are a number of variants of LWE (such as an average-case version, and a version where the samples are from $\mathbb{Z}_q^n \times \mathbb{Z}_q$) that are at least as hard as the version of the problem from Definition 6.14.

We are primarily concerned with a version of LWE in which the error distribution χ over $\mathbb{T}$ is Gaussian-like. For any $\alpha > 0$, define the continuous one-dimensional Gaussian distribution D_α over $\mathbb{R}$ to have density function $\exp(-\pi x^2/\alpha^2)/\alpha$ for all $x \in \mathbb{R}$. Define Ψ_α to be the distribution on $\mathbb{T}$ obtained by taking a sample from D_α and reducing modulo 1.

As in [Reg05], we define the following intermediate worst-case problem to make the results more modular. The main theorem from that work follows.

Definition 6.15 (Discrete Gaussian Sampling Problem)**.** An input to DGS^ϕ is a basis $\mathbf{B}$ of an n-dimensional lattice $\Lambda = \mathcal{L}(\mathbf{B})$ and a parameter $s \geq \phi(\Lambda)$. The goal is to output a sample distributed according to $D_{\Lambda,s}$.

Proposition 6.16 ([Reg05, Theorem 3.1])**.** *Let $\epsilon(n)$ be any negligible function, let $q(n)$ be an integer, and let $\alpha(n) \in (0,1)$ be such that $\alpha(n) \cdot q(n) > 2\sqrt{n}$. Then there is a* quantum *polynomial-time reduction from solving DGS^ϕ for $\phi(\Lambda) = \sqrt{2n} \cdot \eta_\epsilon(\Lambda)/\alpha(n)$ to solving $\mathsf{LWE}_{q,\Psi_\alpha}$ given a* $\mathrm{poly}(n)$ *number of samples.*

In order to reduce SIVP to DGS, we need a few lemmas. The reduction follows in Corollary 6.20.

Lemma 6.17 ([Reg05, Claim 2.13])**.** *There is a constant $c > 0$ such that for any n-dimensional lattice and any $\epsilon \in (0, \frac{1}{2})$, we have $\eta_\epsilon(\Lambda) \geq c \cdot \lambda_n^{(2)}(\Lambda)/n$.*

Lemma 6.18 ([Reg05, Corollary 3.14])**.** *Let Λ be an n-dimensional lattice and let $s \geq \sqrt{2} \cdot \eta_\epsilon(\Lambda)$ for some $\epsilon \leq \frac{1}{10}$. Then a set of n^2 vectors, each chosen independently from $D_{\Lambda,s}$, contains n linearly independently vectors, except with probability exponentially small in n.*

Lemma 6.19. *Let $p \in [1,\infty]$, let Λ be an n-dimensional lattice, let $s \geq \eta_\epsilon(\Lambda)$ for some $\epsilon \leq 1/3$, and let $r \geq 0$. Then for a sample $\mathbf{x} \sim D_{\Lambda,s}$, we have $\|\mathbf{x}\|_p \leq s \cdot n^{1/p} \cdot r$ except with probability at most $2en \cdot \exp(-\pi r^2)$.*

Proof. Because $\|\mathbf{x}\|_p \leq n^{1/p} \cdot \|\mathbf{x}\|_\infty$, it suffices to consider $p = \infty$. The lemma follows immediately from Corollary 5.3, for the special case $m = 1$, $\mathbf{z} = (1)$, and $\mathbf{C} = \mathbf{0}$. □

Corollary 6.20. *Let $p \in [2,\infty]$ and $\epsilon(n) \leq 1/10$. For any function $\phi(\Lambda) \geq \sqrt{2} \cdot \eta_\epsilon(\Lambda)$ and any $\gamma(n) = n^{1/p} \cdot \omega(\sqrt{\log n})$, there is a polynomial time reduction from $\mathsf{GIVP}_\gamma^{p,\phi}$ to DGS^ϕ.*

In particular, for the parameters described in Proposition 6.16, solving $\mathsf{LWE}_{q,\Psi_\alpha}$ is as hard as solving $\mathsf{SIVP}_{\gamma'}^p$ with a quantum algorithm, for any $\gamma' = \omega((n \log n)/\alpha(n))$.

Proof. The second claim follows immediately from Proposition 6.16 and the second inequality of Lemma 2.7 (relating η_ϵ to $\lambda_n^{(p)}$).

The reduction and its analysis are almost exactly as in [Reg05, Lemma 3.15]. The main idea is simply to take n^2 samples from $D_{\Lambda,s}$ and choose an arbitrary subset of n linearly independent vectors from among them, which exist with high probability by Lemma 6.18. To be precise, the

function $s(\Lambda)$ might not be efficiently computable, so the actual reduction tries a polynomial number of different values for s.

In detail, the input is a lattice Λ represented by some basis $\mathbf{B}$. We aply the LLL algorithm [LLL82] to obtain a set $\mathbf{S} \subset \Lambda$ of n linearly independent lattice vectors, where $\|\mathbf{S}\|_2 \leq 2^n \cdot \lambda_n^{(2)}(\Lambda)$. For each $i \in \{0, \ldots, 2n\}$, call the DGS oracle n^2 times on the input $(\mathbf{B}, s_i)$ for $s_i = 2^{-i} \cdot \|\mathbf{S}\|_2$, and let $\mathbf{S}_i$ be the resulting set of vectors. When finished, look for a set of n linearly independent vectors in each of $S, S_0, \ldots, S_{2n}$, and output the shortest such set.

We now show that the reduction is correct. If $\phi(\Lambda) \geq \|\mathbf{S}\|_2$, then $\|\mathbf{S}\|_p \leq \|\mathbf{S}\|_2 \leq \gamma(n) \cdot \phi(\Lambda)$, and we are done. Otherwise, let $i \in \{0, \ldots, 2n\}$ be such that $\phi(\Lambda) \leq s_i \leq 2\phi(\Lambda)$; such an i must exist by Lemma 6.17. By Lemma 6.18, $\mathbf{S}_i$ contains n linearly independent vectors with overwhelming probability. Moreover, by Lemma 6.19, $\|\mathbf{S}_i\|_p \leq \gamma(n) \cdot \phi(\Lambda)$ with overwhelming probability. Hence the reduction succeeds. □

In order to reduce GapSVP to DGS, we need to introduce one more intermediate problem (which is defined relative to any implicit norm).

Definition 6.21. An input to GapCVP'_γ is a pair $(\mathbf{B}, \mathbf{v})$ where $\mathbf{B}$ is a basis of an n-dimensional lattice and $\mathbf{v} \in \mathbb{R}^n$. It is a YES instance if $\mathrm{dist}(\mathbf{v}, \mathcal{L}(\mathbf{B})) \leq 1$, and is a NO instance if both $\mathrm{dist}(\mathbf{v}, \mathcal{L}(\mathbf{B})) > \gamma$ and $\lambda_1(\mathcal{L}(\mathbf{B})) > \gamma$.

It is shown in [GMSS99] that for any $\gamma(n) \geq 1$, there is a norm- and approximation-preserving reduction from GapSVP to GapCVP'. Therefore, it suffices to show a reduction from GapCVP' to DGS in the ℓ_p norm. Essentially, the reduction just uses the DGS oracle to generate a witness for the coNP verifier of [AR05].

Corollary 6.22. *For any $p \in [2, \infty)$, there is a constant $c_p > 0$ such that the following holds. For any $\gamma'(n) \geq 1$, there is a polynomial time reduction from GapSVP^p_γ to DGS^ϕ, where $\gamma(n) = 100 c_p \sqrt{n} \cdot \gamma'(n)$ and $\phi(\Lambda) = c_p n^{1/p} \cdot \gamma'(n) / \lambda_1^{(p)}(\Lambda^*)$.*

For $p = \infty$, the same is true for $\gamma(n) = 100\sqrt{n \log n} \cdot \gamma'(n)$ and $\phi(\Lambda) = \sqrt{\log n} \cdot \gamma'(n) / \lambda_1^{\infty}(\Lambda^*)$.

In particular, for the parameters described in Proposition 6.16, solving $\mathsf{LWE}_{q,\Psi_\alpha}$ is as hard as solving GapSVP^p_γ with a quantum algorithm, for some $\gamma = \tilde{O}(n/\alpha(n))$.

Proof. The final part of the claim follows directly from Proposition 6.16 and Lemma 3.5.

First suppose $p \in [2, \infty)$. Recall the algorithm $\mathcal{V}$ from Section 4.1, which has the following properties. Its input is a basis $\mathbf{B}$ for a lattice $\Lambda = \mathcal{L}(\mathbf{B})$, a point $\mathbf{v} \in \mathbb{R}^n$, and a matrix $\mathbf{W} \subset \Lambda^*$ of N dual lattice vectors, for some $N = \mathrm{poly}(n)$. When $\mathrm{dist}^p(\mathbf{v}, \mathcal{L}(\mathbf{B})) \leq n^{1/p - 1/2}/100$, $\mathcal{V}$ always rejects. When $\mathrm{dist}^p(\mathbf{v}, \mathcal{L}(\mathbf{B})) \geq c_p n^{1/p}$ (for appropriate constant c_p), and $\mathbf{w}_i$ are chosen independently from D_{Λ^*}, then $\mathcal{V}$ accepts with overwhelming probability.

The reduction works as follows. The input is $(\mathbf{B}, \mathbf{v})$; assume without loss of generality (by scaling) that either $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p - 1/2}/100$, or both $\mathrm{dist}^p(\mathbf{v}, \Lambda)$ and $\lambda_1^{(p)}(\Lambda)$ exceed $c_p n^{1/p} \cdot \gamma'(n)$. Call the DGS oracle N times on the basis $\mathbf{B}^{-T}$ (which is a basis for Λ^*) and parameter $s = 1$ to generate the columns of $\mathbf{W}$. Call $\mathcal{V}(\mathbf{B}, \mathbf{v}, \mathbf{W})$, and accept if and only if $\mathcal{V}$ rejects.

We now show correctness. In the YES case, we have $\mathrm{dist}^p(\mathbf{v}, \Lambda) \leq n^{1/p - 1/2}/100$, so $\mathcal{V}$ always rejects (for any $\mathbf{W}$). In the NO case, we have $s = 1 > \phi(\Lambda^*) = c_p n^{1/p} \cdot \gamma'(n) / \lambda_1^{(p)}(\Lambda)$, so the $\mathbf{w}_i$ are actual samples from D_{Λ^*}. Moreover, $\mathrm{dist}^p(\mathbf{v}, \Lambda) > c_p n^{1/p} \cdot \gamma'(n) \geq c_p n^{1/p}$, so $\mathcal{V}$ accepts with overwhelming probability.

The proof is nearly identical for $p = \infty$ using the properties of $\mathcal{V}$ relative to the ℓ_∞ norm. □

7 Acknowledgments

I gratefully thank Vadim Lyubashevsky, Oded Regev, Alon Rosen, and the anonymous reviewers for many helpful and constructive comments.

References

[ABSS97] Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. Syst. Sci.*, 54(2):317–331, 1997. Preliminary version in FOCS 1993.

[Ajt98] Miklós Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19, 1998.

[Ajt04] Miklós Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.

[AKKV05] Mikhail Alekhnovich, Subhash Khot, Guy Kindler, and Nisheeth K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing. In *FOCS*, pages 216–225, 2005.

[AKS01] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610, 2001.

[AKS02] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *IEEE Conference on Computational Complexity*, pages 53–57, 2002.

[AR03] Dorit Aharonov and Oded Regev. A lattice problem in quantum NP. In *FOCS*, pages 210–219, 2003.

[AR05] Dorit Aharonov and Oded Regev. Lattice problems in NP $\cap$ coNP. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004.

[Bab86] László Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.

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

[Ban95] Wojciech Banaszczyk. Inequalites for convex bodies and polar reciprocal lattices in R^n. *Discrete & Computational Geometry*, 13:217–231, 1995.

[BDCGL92] Shai Ben-David, Benny Chor, Oded Goldreich, and Michael Luby. On the theory of average case complexity. *J. Comput. Syst. Sci.*, 44(2):193–219, 1992.

[BK02] Piotr Berman and Marek Karpinski. Approximating minimum unsatisfiability of linear equations. In *SODA*, pages 514–516, 2002.

[BN07] Johannes Blömer and Stefanie Naewe. Sampling methods for shortest vectors, closest vectors and successive minima. In *ICALP*, pages 65–77, 2007.

[Cai98] Jin-Yi Cai. A relation of primal-dual lattices and the complexity of shortest lattice vector problem. *Theor. Comput. Sci.*, 207(1):105–116, 1998.

[CN97] Jin-Yi Cai and Ajay Nerurkar. An improved worst-case to average-case connection for lattice problems. In *FOCS*, pages 468–477, 1997.

[CN99] Jin-Yi Cai and Ajay 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.

[Din02] Irit Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. *Theor. Comput. Sci.*, 285(1):55–71, 2002. Preliminary version in CIAC 2000.

[DKRS03] Irit Dinur, Guy Kindler, Ran Raz, and Shmuel Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003. Preliminary version in FOCS 1998.

[FM04] Uriel Feige and Daniele Micciancio. The inapproximability of lattice and coding problems with preprocessing. *J. Comput. Syst. Sci.*, 69(1):45–67, 2004. Preliminary version in CCC 2002.

[GG00] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000. Preliminary version in STOC 1998.

[GMR05] Venkatesan Guruswami, Daniele Micciancio, and Oded Regev. The complexity of the covering radius problem. *Computational Complexity*, 14:90–121, 2005. Preliminary version in CCC 2004.

[GMSS99] Oded Goldreich, Daniele Micciancio, Shmuel Safra, and Jean-Pierre Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Inf. Process. Lett.*, 71(2):55–61, 1999.

[Gol] Oded Goldreich. Note at `http://www.wisdom.weizmann.ac.il/~oded/p_lp.html`.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008. To appear. Full version available at `http://eprint.iacr.org/2007/432`.

[HR06] Ishay Haviv and Oded Regev. Hardness of the covering radius problem on lattices. In *IEEE Conference on Computational Complexity*, pages 145–158, 2006.

[HR07] Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *STOC*, pages 469–477, 2007.

[Kho05] Subhash Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005. Preliminary version in FOCS 2004.

[Lev86] Leonid A. Levin. Average case complete problems. *SIAM J. Comput.*, 15(1):285–286, 1986.

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

[LLM06] Yi-Kai Liu, Vadim Lyubashevsky, and Daniele Micciancio. On bounded distance decoding for general lattices. In *APPROX-RANDOM*, pages 450–461, 2006.

[LLS90] Jeffrey C. Lagarias, Hendrik W. Lenstra, Jr., and Claus-Peter Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.

[LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155, 2006. Full version in ECCC Report TR05-142.

[LM08] Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, pages 37–54, 2008.

[MG02] Daniele Micciancio and Shafi 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.

[Mic00] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035, 2000.

[Mic04] Daniele Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM J. Comput.*, 34(1):118–169, 2004. Preliminary version in STOC 2002.

[Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, December 2007. Preliminary version in FOCS 2002.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.

[NN94] Yurii Nesterov and Arkadii Nemirov. *Interior-Point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics (SIAM), 1994.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006. Full version in ECCC TR05-158.

[PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487, 2007. Full version in ECCC Report TR06-147.

[PVW07] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. Cryptology ePrint Archive, Report 2007/348, 2007. In submission. Full version available at `http://eprint.iacr.org/2007/348`.

[PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *STOC*, 2008. To appear. Full version available at `http://eprint.iacr.org/2007/279`.

[Reg04a] Oded Regev. Improved inapproximability of lattice and coding problems with preprocessing. *IEEE Transactions on Information Theory*, 50(9):2031–2037, 2004. Preliminary version in CCC 2003.

[Reg04b] Oded Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004. Preliminary version in STOC 2003.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005. Revised version available from author's web page.

[RR06] Oded Regev and Ricky Rosen. Lattice problems and norm embeddings. In *STOC*, pages 447–456, 2006.

[Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[vEB81] Peter van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, University of Amsterdam, 1981.