

Lattice Problems and Norm Embeddings

Oded Regev * Ricky Rosen †

November 4, 2005

Abstract

We present reductions from lattice problems in the ℓ_2 norm to the corresponding problems in other norms such as ℓ_1 , ℓ_∞ (and in fact in any other ℓ_p norm where $1 \leq p \leq \infty$). We consider lattice problems such as the Shortest Vector Problem, Shortest Independent Vector Problem, Closest Vector Problem and the Closest Vector Problem with Preprocessing. The reductions are based on *embeddings of normed spaces*. Among other things, our reductions imply that the Shortest Vector Problem in the ℓ_1 norm and the Closest Vector Problem with Preprocessing in the ℓ_∞ norm are hard to approximate to within any constant (and beyond). Previously, the former problem was known to be hard to approximate to within $2 - \epsilon$, while no hardness result was known for the latter problem.

1 Introduction

1.1 Lattices

Given n -linearly independent vectors $b_1, \dots, b_n \in \mathbb{R}^m$, the lattice generated by them is the set of vectors

$$\mathcal{L}(b_1, \dots, b_n) := \left\{ \sum_{i=1}^n x_i b_i \mid x_i \in \mathbb{Z} \right\} \subseteq \mathbb{R}^m$$

We refer to n as the *rank* of the lattice and to m as the *dimension* of the lattice. We also say that the vectors b_1, \dots, b_n form a *basis* of the lattice.

There are several natural computational problems involving lattices. These problems play a fundamental role in various areas, such as computational number theory and cryptography. Two of the main lattice problems are the Shortest Vector Problem (SVP) and the Closest Vector Problem (CVP). In SVP, the goal is to find the shortest nonzero vector in a lattice given some basis of the lattice. In CVP, in addition to a lattice, we are given some target vector and the goal is to find the lattice vector that is closest to the target vector. In both problems, one needs to specify the norm by which the length of the shortest vector, or the distance to the closest vector is measured. A common choice is to take the ℓ_p norm for some $1 \leq p \leq \infty$, where for $1 \leq p < \infty$ we define $\|x\|_p = (\sum_i |x_i|^p)^{1/p}$ and $\|x\|_\infty = \max_i |x_i|$.

In this paper, we are mainly interested in approximation versions of these problems. For example, in the case of SVP, the goal is to find a lattice vector that is longer than the shortest lattice

*Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by an Alon Fellowship, by the Binational Science Foundation, and by the Israel Science Foundation.

†Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel.

vector by at most some given *approximation factor*. Similarly, in the case of CVP, the goal is to find a lattice vector whose distance to the target point is larger than that of the nearest lattice vector by at most the approximation factor.

The complexity of lattice problems has been investigated extensively in the last two decades, both in terms of approximation algorithms (e.g., [LLL82]), and computational hardness results (e.g., [DKRS03, Kho04]). Connections to average-case hardness have been obtained in the seminal work of Ajtai [Ajt96], which later led to lattice-based cryptosystems (e.g., [AD97, Reg03b]). For a more thorough introduction to lattices see [MG02a].

Algorithms: The earliest known lattice algorithm is attributed to Gauss [Gau66] who, in 1801, described an algorithm that solves 2-dimensional SVP. The first algorithm that works for any dimension was given almost two centuries later by Lenstra, Lenstra and Lovász (LLL) [LLL82]. In this celebrated result, they describe an efficient algorithm that approximates SVP to within $2^{n/2}$ where n is the rank of the lattice. Although this might seem like a very poor approximation factor, their algorithm had a significant impact on many important problems such as factoring polynomials over the rationals and simultaneous diophantine approximation [LLL82], integer programming [Len83], solving low-density subset-sum problems, breaking knapsack-based cryptosystems [LO85], and approximation algorithms for CVP [Bab86]. Moreover, it seems very difficult to obtain significantly better factors. The best known deterministic algorithm was given by Schnorr [Sch85] and achieves an approximation factor of $2^{O(n(\log \log n)^2 / \log n)}$. The best known algorithm is a recent randomized algorithm that achieves an approximation factor of $2^{O(n \log \log n / \log n)}$ [AKS01].

As for an *exact* solution to SVP, all known algorithms have exponential running time. The best known deterministic algorithm is the one by Kannan [Kan87], and has running time is $2^{O(n \log n)}$. The best known randomized algorithm has running time $2^{O(n)}$ and is given in a recent paper by Ajtai et al. [AKS01].

Hardness: There are many known hardness results for lattice problems. We will now review some of them in detail. As we will see, all known results give hardness of approximation for sub-polynomial factors. This is in stark contrast to the almost exponential approximation factors achieved by known algorithms. One might therefore suspect that the known hardness results can be substantially improved. However, it is a somewhat surprising fact that most lattice problems are *unlikely* to be NP-hard to approximate to within factors such as $O(\sqrt{n})$. Such results can be found in, e.g., [GG00, AR04]. For example, [AR04] shows that approximating SVP in the ℓ_2 norm to within a factor of $O(\sqrt{n})$ is in $\text{NP} \cap \text{coNP}$ and is therefore not NP-hard unless the polynomial hierarchy collapses.

Hardness of SVP: Already in 1981, van Emde Boas [vEB81] showed that SVP in the ℓ_∞ norm is NP-hard. Extending this result to other norms turned out to be very difficult. This was finally achieved by Ajtai in 1998 [Ajt98]. He showed that for any $p \geq 1$, SVP in the ℓ_p norm is hard assuming $\text{NP} \not\subseteq \text{RP}$. This was later improved, under the same complexity assumption, by Cai and Nerrurkar [CN99] who showed that the problem is hard to approximate to within $(1 + \frac{1}{n^\epsilon})$. The first constant factor hardness result was obtained by Micciancio [Mic01] who proved that for any $p < \infty, \epsilon > 0$, SVP in the ℓ_p norm is hard to approximate to within $\sqrt[p]{2} - \epsilon$ under the assumption that $\text{NP} \not\subseteq \text{RP}$. Dinur [Din02] boosted the hardness of SVP in the ℓ_∞ norm to $n^{O(1/\log \log n)}$. Khot [Kho03] showed a hardness result of the form $p^{0.99}$ for large enough p . Finally, in a recent

breakthrough result, Khot [Kho04] shows that for any p strictly greater than 1, SVP in the ℓ_p norm is hard to approximate to within any constant unless $\text{NP} \subseteq \text{RP}$. Under the stronger assumption that $\text{NP} \not\subseteq \text{RTIME}(2^{\text{poly}(\log n)})$, the hardness factor becomes $2^{(\log n)^{1/2-\epsilon}}$ for any $\epsilon > 0$.

Summarizing, the best known hardness result is $n^{O(1/\log \log n)}$ in ℓ_∞ [Din02], any constant (and beyond) in ℓ_p for $1 < p < \infty$ [Kho04], and $2 - \epsilon$ in ℓ_1 [Mic01].

Hardness of other lattice problems: Already in 1981, van Emde Boas [vEB81] established that CVP is NP-hard in any ℓ_p norm. A series of hardness results for CVP were obtained in the following years. The best known hardness result shows that for any $1 \leq p \leq \infty$, CVP in the ℓ_p norm is NP-hard to approximate to within $n^{O(1/\log \log n)}$. For $p < \infty$, this result is due to Dinur, Kindler, Raz, and Safra [DKRS03]; Dinur [Din02] used similar techniques to obtain the result for $p = \infty$.

Another problem considered in this paper is a variant of CVP known as the Closest Vector Problem with Preprocessing (CVPP). The motivation for this problem comes from applications in cryptography and coding (see [MG02b] for more details). Informally, our goal in CVPP is to solve the CVP for *any fixed* lattice (so the input of the algorithm consists only of a target point). Since the lattice is fixed, our algorithm can be hardwired with some information on the lattice that is hard to compute, such as a short basis. Indeed, using such ideas, [AR04] present an algorithm that approximates CVPP to within $O(n/\sqrt{\log n})$ in the ℓ_p norm for $1 \leq p \leq \infty$. Feige and Micciancio [FM04] showed that for any $p < \infty$, CVPP is NP-hard to approximate to within $\sqrt[3]{5/3} - \epsilon$. This was later improved to $\sqrt[3]{3} - \epsilon$ by Regev [Reg03a]. More recently, Alekhnovich, Khot, Kindler and Vishnoi [AKKV05] improved this result by showing that CVPP is NP-hard to approximate to within any constant. They also showed hardness of $(\log n)^{1/p-\epsilon}$ for any $1 \leq p < \infty$ under the assumption that $\text{NP} \not\subseteq \text{DTIME}(2^{\text{poly} \log(n)})$. Note that no hardness result was previously known for $p = \infty$.

Two more lattice problems considered here are the *Shortest Independent Vector Problem* (SIVP) and the *Shortest Basis Problem* (SBP). Part of the interest in these problems comes from Ajtai's work [Ajt96] where they play an important role. In the former problem, the goal is to find n linearly independent vectors in a given lattice while minimizing the length of the longest vector in the set. In the latter problem, the goal is to find a basis of a given lattice while minimizing the length of the longest vector in the basis. Both problems can be defined in their approximation version. Blömer and Seifert [BS99] showed that in the ℓ_2 norm, both problems are NP-hard to approximate to within any constant. Under the stronger assumption $\text{NP} \not\subseteq \text{DTIME}(2^{\text{poly} \log(n)})$, they prove a hardness factor of $2^{\log^{1-\epsilon} n}$ for any $\epsilon > 0$.

Behavior as a function of p ? Considering the above results, we see many different behaviors as a function of p : one result works only for large enough p and improves as p increases [Kho03], some results only hold for $p = \infty$ [vEB81, Din02], most results weaken as p increases and do not work for $p = \infty$ [Mic01, FM04, Reg03a, AKKV05], and one result does not work for $p = 1$ [Kho04]. Is there any unifying pattern here? This is the question we set out to answer in this paper.

Our Results

Informally, our main result shows that for lattice problems, the ℓ_2 norm is the easiest. A more precise description is given in the following theorem.

Theorem 1.1 *For any $\epsilon > 0$, $\gamma \geq 1$ and $1 \leq p \leq \infty$ there exists a randomized polynomial time reduction from approximating SVP in the ℓ_2 norm to within $(1 + \epsilon)\gamma$ to approximating SVP in the ℓ_p norm to within γ . A similar result holds for CVP, SIVP, SBP, and CVPP.*

It is crucial for our applications that these reductions almost preserve the approximation factor. As we mention later, if one is willing to lose a factor of \sqrt{n} in the approximation factor, trivial reductions exist. For the case $p = 1$, we also present a deterministic reduction, but it loses a factor of $\sqrt{3}$ in the approximation ratio (as opposed to $1 + \epsilon$). This turns out to be sufficient for our applications. Furthermore, all our reductions in the case of CVPP are deterministic. Finally, our reductions hold also for the decision variants of these lattice problems.

The theorem implies that it is enough to show hardness of approximation for problems such as SVP, SIVP, and CVP in the ℓ_2 norm. Any such hardness result automatically implies hardness of approximation within essentially the same factor in all other ℓ_p norms for $1 \leq p \leq \infty$. By combining our theorem with known results, we obtain the following new hardness of approximation results.

- Using [Kho04], we obtain that SVP in the ℓ_1 norm is hard to approximate
 - to within any constant unless $\text{NP} \subseteq \text{RP}$, and
 - to within $2^{(\log n)^{1/2-\epsilon}}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{RTIME}(2^{\text{poly}(\log n)})$.

Previously, this problem was known to be hard to approximate to within $2 - \epsilon$ unless $\text{NP} \subseteq \text{RP}$ [Mic01].

- Using [AKKV05], we obtain that for any $2 < p < \infty$, CVPP in the ℓ_p norm is hard to approximate
 - to within $(\log n)^{1/2-\epsilon}$, for any $\epsilon > 0$, unless NP has quasi-polynomial sized circuits.

Previously, this problem was known to be hard to approximate to within $(\log n)^{1/p-\epsilon}$ under the same assumption. For the case $p = \infty$, we show that CVPP is hard to approximate

- to within any constant unless $\text{NP} \subseteq \text{P/poly}$, and
- to within $(\log n)^{1/2-\epsilon}$ for any $\epsilon > 0$, unless NP has quasi-polynomial sized circuits.

Previously, no hardness of approximation was known in the ℓ_∞ norm.

- Using [BS99], we obtain that SIVP and SBP in the ℓ_p norm for $1 \leq p \leq \infty$ is hard to approximate
 - to within any constant unless $\text{NP} \subseteq \text{BPP}$, and
 - to within $2^{(\log n)^{1-\epsilon}}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{BPTIME}(2^{\text{poly} \log(n)})$.

To the best of our knowledge, all previous results hold only for the ℓ_2 norm.

It also follows from the theorem that it is enough to find approximation algorithms for, say, SVP or CVP in, say, the ℓ_1 or ℓ_∞ norm. Any such algorithm yields an algorithm in the ℓ_2 norm with essentially the same approximation factor. Obtaining algorithms in the ℓ_1 and ℓ_∞ norm might be conceptually easier. For instance, [AKS01] first present an algorithm for SVP in the ℓ_∞ norm and then extend it to the ℓ_2 norm.

However, the usefulness of Theorem 1.1 in deriving algorithmic results in the case $p > 2$ is quite limited. The reason for this is that for $p > 2$, the reductions incur a polynomial blowup in the dimension. To see why this might be a problem, assume one comes up with an algorithm that solves SVP in the ℓ_∞ norm in time $2^{O(\sqrt{n})}$ where n is the dimension. Because of the polynomial blowup in the dimension, the running time of the resulting algorithm in the ℓ_2 norm is $2^{O(n^c)}$ for some possibly large constant c and is thus not better than known algorithms. To this end, we prove the following theorem. It shows a reduction from ℓ_2 to ℓ_p for any $2 < p \leq \infty$ that preserves the dimension. This comes at the expense of a small loss in the approximation factor.

Theorem 1.2 *For any $\gamma \geq 1$ and $2 \leq p \leq \infty$ there exists a randomized polynomial time dimension-preserving reduction from approximating SVP in the ℓ_2 norm to within γ' to approximating SVP in the ℓ_p norm to within γ where in the case $p < \infty$, $\gamma' = c(p)\gamma$ for some constant $c(p)$ and in the case $p = \infty$, $\gamma' = O(\sqrt{\log n \gamma})$. A similar result holds for CVP, SIVP, SBP, and CVPP.*

1.2 Embeddings

It is a well-known fact (and not too difficult to verify) that for any vector $x \in \mathbb{R}^n$ and any p , $n^{1/p-1/2}\|x\|_2 \leq \|x\|_p \leq \|x\|_2$ if $p > 2$ and $\|x\|_2 \leq \|x\|_p \leq n^{1/p-1/2}\|x\|_2$ if $1 \leq p < 2$. This fact alone already allows one to construct (trivial) reductions among lattice problems in different norms. These reductions, however, lose a polynomial factor (in n) in the approximation factor, and are therefore useless when dealing with sub-polynomial approximation factors (as is the case in virtually all hardness of approximation results of lattice problems). Our goal in this paper is to present reductions that (almost) preserve the approximation factor. Our reductions are based on the notion of an *embedding*, described next.

Let X be the space \mathbb{R}^n equipped with the ℓ_q norm and let Y be the space \mathbb{R}^m equipped with the ℓ_p norm. By an *embedding* from X to Y we mean a linear function $f : X \rightarrow Y$.¹ If an embedding f satisfies that for any two points $x, x' \in X$,

$$(1 - \epsilon)\|x - x'\|_q \leq \|f(x) - f(x')\|_p \leq (1 + \epsilon)\|x - x'\|_q,$$

then we say that the *distortion* of f is $\frac{1+\epsilon}{1-\epsilon}$. Since f is linear, this is equivalent to the property that for all $x \in X$

$$(1 - \epsilon)\|x\|_q \leq \|f(x)\|_p \leq (1 + \epsilon)\|x\|_q. \tag{1}$$

Moreover, it is equivalent to require this only for x such that $\|x\|_q = 1$ (again by linearity). Having a low distortion is a desirable property since it means that distances are almost preserved. We note that another notion of embedding, namely embedding of finite metrics, has recently attracted much attention (see, e.g., [ALN05]). In contrast, our focus in this paper is on embeddings of normed spaces (containing an infinite number of points).

Embeddings of normed spaces have fascinated mathematicians for many decades, and recently also found many applications in computer science (see e.g., [AMS99, Das99, Ind00, Ind03] and the survey by Indyk [Ind01]). Broadly speaking, such embeddings can be categorized into two types. In the first type, the dimension of the host space m is (much) smaller than that of the original space n (a good example is the Johnson-Lindenstrauss lemma). Such embeddings are known as dimension reduction embeddings and can be used, for example, to design more efficient algorithms. The

¹We remark that one can easily extend this definition (and most of our results) to arbitrary finite dimensional Banach spaces, but for simplicity we restrict our attention to ℓ_p spaces.

second type of embedding consists of embeddings between different norms. As the name suggest, the main feature of such embeddings is that the norm of the host space is different from that in the original space. Typically in such embeddings, the dimension of the host space m is larger than that of the original space n . Such embeddings can be used to extend a solution for a problem in one norm to a solution in another norm. It is this latter type that we use in this paper.

The celebrated theorem of Dvoretzky [Dvo61] was the first to establish that embeddings between norms exist. More specifically, it shows that if X is the space \mathbb{R}^n equipped with the ℓ_2 norm and Y is the space \mathbb{R}^m equipped with some ℓ_p norm, then for large enough m , a random orthogonal mapping from X to Y is with high probability an embedding with very low distortion, namely, $1 + \epsilon$. (An equivalent, and more common formulation of Dvoretzky's theorem says that most low-dimensional slices of an ℓ_p space are close to Euclidean; yet another equivalent formulation says that most low-dimensional slices of the ℓ_p ball are close to an ℓ_2 ball.) This result was later improved by Figiel, Lindenstrauss and Milman [FLM77] who gave a tight estimate on how large m should be. Specifically, they show that for any $p < \infty$, it is enough to take $m = \text{poly}(n)$.

However, for the case $p = \infty$, things are not that good: [FLM77] showed that in order to achieve any constant distortion, m must be *exponential* in n . A way to bypass this problem was suggested by Indyk in [Ind03]. His idea was to relax the distortion requirement and to instead settle for the requirement that for any fixed $x \in X$, each of the inequalities in (1) holds with high probability over the choice of f . He then shows that such an embedding exists for $m = \text{poly}(n)$. We will discuss this in more detail later.

All embedding results mentioned so far are probabilistic. There are a few known explicit (deterministic) embeddings of ℓ_2 into ℓ_1 . One such embedding has distortion $\sqrt{3}$ [Ber97, LLR95]; another has distortion only $1 + \epsilon$ but requires m to be super-polynomial in n [Ind00]. Another relevant result is that of Johnson and Schechtman [JS82] who extended the result of [FLM77] by showing that for any $1 \leq q < p \leq 2$, ℓ_p^n can be embedded into ℓ_q^m with low distortion.

In most of our applications we use these embedding results in a black-box fashion. Nevertheless, it is instructive to review how such embedding results are established. Specifically, we concentrate on the elegant proof of [FLM77] and consider the case $p = 1$. Recall that our goal is to prove that with high probability over the choice of f , Equation (1) (with $p = 1$, $q = 2$) holds for any $x \in \mathbb{R}^n$. Moreover, it is enough to show this for any x satisfying $\|x\|_2 = 1$. The main idea in the proof is what is known as the *concentration of measure phenomenon*. In our case, it implies that a y chosen uniformly from the unit ℓ_2 sphere $S^{m-1} = \{y \in \mathbb{R}^m \mid \|y\|_2 = 1\}$ satisfies that $\|y\|_1$ is in the range $(1 \pm \epsilon)M$ with probability $1 - 2^{-cm}$ for some constant c and some M that depends only on m . Let f be a random orthogonal mapping from \mathbb{R}^n to \mathbb{R}^m normalized by $1/M$. Fix some $x \in \mathbb{R}^n$ with $\|x\|_2 = 1$. Using the fact that a random orthogonal mapping maps x to a uniform element of S^{m-1} , we see that $\|f(x)\|_1 \in (1 \pm \epsilon)$ holds with probability $1 - 2^{-cm}$. By an application of a union bound, we obtain that for any fixed set $K \subseteq \mathbb{R}^n$ of cardinality at most $2^{cm/2}$, Equation (1) is satisfied for all $x \in K$ with probability at least $1 - 2^{-cm/2}$. To complete the proof, [FLM77] show that it is enough if (1) holds for all x in some 'dense enough' set of points on the unit ℓ_2 sphere, and that there exist such a set of cardinality 2^{Cn} for some constant C . Hence, by choosing $m > 2Cn/c = O(n)$ we obtain the required result. For more details, see the friendly survey by Ball [Bal97, Sections 8 and 9].

1.3 Our techniques

For the proof of our main theorem we use three embedding techniques. The first is the one by Figiel et al. [FLM77], and is used in all our reductions to ℓ_p norms for $p < \infty$. The second is an embedding into ℓ_∞ . We construct such an embedding with the property that it does not contract *any* point in \mathbb{R}^n (i.e., the left inequality in (1) holds for any x) and that for any fixed point in \mathbb{R}^n there is a polynomially small probability of expansion (i.e., for any fixed x , the right inequality in (1) holds with high probability). This embedding is based on the work of Indyk [Ind03] and could be useful in other applications. We use this embedding in almost all our reductions to the ℓ_∞ norm. The third embedding technique is used in our CVPP reduction for the ℓ_∞ case. There, we need to find some sort of embedding that is good for all $x \in \mathbb{R}^n$ *simultaneously*. For this purpose, instead of considering just one embedding, we take a sample of embeddings into ℓ_∞ . We show that with high probability, such a sample has the following property. For any point $x \in \mathbb{R}^n$, there is an embedding in the sample that does not expand it by ‘too much’ and moreover, no embedding in the sample contracts it by ‘too much’.

Finally, for the proof of Theorem 1.2 we analyze a remarkably simple embedding that consists of nothing more than a random rotation. Somewhat surprisingly, we show that this simple embedding is already good enough to be used for lattice reductions.

(Oded: outline?)

2 Preliminaries

The space \mathbb{R}^n equipped with the ℓ_p norm is denoted by ℓ_p^n . We denote by S^{n-1} the set of points $\{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$. The ℓ_p distance of a point t from a set of points \mathcal{L} is denoted by $\text{dist}_p(t, \mathcal{L}) = \inf_{v \in \mathcal{L}} \|v - t\|_p$.

Definition 2.1 (successive minima) For a lattice \mathcal{L} , we define its i th successive minimum in the ℓ_p norm as

$$\lambda_i^p(\mathcal{L}) = \inf \{r \mid B_p(r) \text{ contains } i \text{ linearly independent vectors from } \mathcal{L}\}$$

where $B_p(r)$ is defined as $\{x \mid \|x\|_p \leq r\}$. In particular, $\lambda_1^p(\mathcal{L})$ is the length of the shortest nonzero vector in the ℓ_p norm.

Fact 2.2 For $m \geq n$, a full rank linear mapping $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ preserves linear dependencies. In particular, if \mathcal{L} is a lattice in \mathbb{R}^n then $f(\mathcal{L})$ is a lattice in \mathbb{R}^m .

Definition 2.3 (θ -net) A θ -net on the sphere $S^{n-1} = \{x \mid \|x\|_2 = 1\}$ is a set of points $K_0 \subseteq S^{n-1}$ such that for any $x \in S^{n-1}$ there is a point $y \in K_0$ for which $\text{dist}_2(x, y) \leq \theta$.

In this paper we consider the following lattice problems. In all problems, $\gamma \geq 1$ is some arbitrary approximation factor, $1 \leq p \leq \infty$ is the norm involved, and the input lattice is given in the form of some arbitrary basis. For each problem we describe both its search variant and its decision variant.

Definition 2.4 (SVP $_\gamma^p$) In SVP $_\gamma^p$, we are given a lattice $\mathcal{L} \subseteq \mathbb{R}^n$. Our goal is to find a nonzero vector $v \in \mathcal{L}$ such that $\|v\|_p \leq \gamma \cdot \lambda_1^p(\mathcal{L})$.

Definition 2.5 (GapSVP $^p_\gamma$) In GapSVP $^p_\gamma$, we are given a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and a parameter $d > 0$. In YES instances, $\lambda_1^p(\mathcal{L}) \leq d$ whereas in NO instances $\lambda_1^p(\mathcal{L}) > \gamma \cdot d$.

Definition 2.6 (CVP $^p_\gamma$) In CVP $^p_\gamma$, we are given a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and a target vector $t \in \mathbb{R}^n$. Our goal is to find a vector $v \in \mathcal{L}$ such that $\|v - t\|_p \leq \gamma \cdot \text{dist}_p(t, \mathcal{L})$.

Definition 2.7 (GapCVP $^p_\gamma$) In GapCVP $^p_\gamma$, we are given a lattice $\mathcal{L} \subseteq \mathbb{R}^n$, a target vector $t \in \mathbb{R}^n$, and a parameter $d > 0$. In YES instances $\text{dist}_p(v, \mathcal{L}) \leq d$. In NO instances $\text{dist}_p(v, \mathcal{L}) > \gamma d$.

We now consider the preprocessing variant of CVP. Recall the intuition that in CVPP we try to solve CVP instances for any *fixed* lattice. In other words, we ask whether for any fixed lattice \mathcal{L} there exists a circuit that solves CVP instances of the form (\mathcal{L}, t) . We make this precise in the following definitions by introducing a computationally unlimited ‘preprocessing’ function.

Definition 2.8 (CVPP $^p_\gamma$) A solution to CVPP $^p_\gamma$ is a (computationally unlimited) function P that maps any given lattice $\mathcal{L} \subseteq \mathbb{R}^n$ to a circuit of size polynomial in the description of \mathcal{L} . This circuit, given any target vector $t \in \mathbb{R}^n$, solves the CVP $^p_\gamma$ instance given by (\mathcal{L}, t) .

Definition 2.9 (GapCVPP $^p_\gamma$) A solution to GapCVPP $^p_\gamma$ is a (computationally unlimited) function P that maps any given lattice $\mathcal{L} \subseteq \mathbb{R}^n$ to a circuit of size polynomial in the description of \mathcal{L} . This circuit, given any target vector $t \in \mathbb{R}^n$, solves the GapCVP $^p_\gamma$ instance given by (\mathcal{L}, t) .

Notice that the above definitions are non-uniform. A uniform formulation also exists. Here, given a uniform sequence of bases $\{B_n\}_{n=1}^\infty$ of increasing dimension (i.e., an algorithm that given n generates the basis B_n), P generates an algorithm that given n , and any point t , solves CVP $^p_\gamma(B_n, t)$. In this paper we focus on the non-uniform case.

We now review the Shortest Independent Vectors Problem (SIVP). In SIVP p , we are given a lattice \mathcal{L} and our goal is to find a set of n linearly independent vectors $\{v_i\}_{i=1}^n$ in \mathcal{L} such that $\max_i \|v_i\|_p$ is minimized.

Definition 2.10 (SIVP $^p_\gamma$) In SIVP $^p_\gamma$, we are given a rank- n lattice $\mathcal{L} \subseteq \mathbb{R}^m$ and our goal is to find a set of n linearly independent vectors $\{v_i\}_{i=1}^n$ in \mathcal{L} such that $\max_i \|v_i\|_p \leq \gamma \cdot \lambda_n^p(\mathcal{L})$.

Definition 2.11 (GapSIVP $^p_\gamma$) In GapSIVP $^p_\gamma$, we are given a rank- n lattice $\mathcal{L} \subseteq \mathbb{R}^m$ and a parameter $d > 0$. In YES instances, $\lambda_n^p(\mathcal{L}) \leq d$ whereas in NO instances $\lambda_n^p(\mathcal{L}) > \gamma d$.

A closely related problem is the Shortest Basis Problem (SBP). In SBP p , given a lattice \mathcal{L} , our goal is to find a basis $\{v_i\}_{i=1}^n$ of \mathcal{L} such that $\max_i \|v_i\|_p$ is minimized. For the following definitions it is useful to define $\kappa^p(\mathcal{L})$ as the minimum of $\max_i \|v_i\|_p$ over all bases $\{v_i\}_{i=1}^n$ of \mathcal{L} .

Definition 2.12 (SBP $^p_\gamma$) In SBP $^p_\gamma$, we are given a rank- n lattice $\mathcal{L} \subseteq \mathbb{R}^m$ and the goal is to find a basis $\{v_i\}_{i=1}^n$ of \mathcal{L} , such that $\max_i \|v_i\|_p \leq \gamma \cdot \kappa^p(\mathcal{L})$.

Definition 2.13 (GapSBP $^p_\gamma$) In GapSBP $^p_\gamma$, we are given a rank- n lattice $\mathcal{L} \subseteq \mathbb{R}^m$ and a parameter $d > 0$. In YES instances, $\kappa^p(\mathcal{L}) \leq d$ and in NO instances $\kappa^p(\mathcal{L}) > \gamma \cdot d$.

Most of our reductions are randomized. Recall that a *randomized Karp reduction* from a decision problem A to a decision problem B is a polynomial time randomized machine R mapping instances of A to instances of B , such that for every $x \in A$, $R(x) \in B$ with probability at least $2/3$ and similarly, for every $x \notin A$, $R(x) \notin B$ with probability at least $2/3$. In addition, a *randomized Cook reduction* R from A to B is a randomized polynomial time machine R that solves A by using an oracle to B . We require that for any instance x of A , $R(x)$ is a correct solution with probability at least $2/3$.

3 Embedding

In this section we introduce several families of embeddings that will be used later to prove the main theorem. In Section 3.1 we state the result of [FLM77] and indicate how to apply it in our setting. Then, in Section 3.2, we discuss the result of Indyk [Ind03] for embedding into ℓ_∞ . We then present a generalization of it that will be used in most of our reductions into ℓ_∞ . In our CVPP reduction into ℓ_∞ , we combine several embeddings together. This method is described and analyzed in Section 3.3.

3.1 Embedding into ℓ_p for $p < \infty$

Definition 3.1 (Embedding family for $p < \infty$) For any real $1 \leq p < \infty$ and integers $m \geq n$ we define a distribution $\mathcal{F}(p, n, m)$ over embedding functions $f: \ell_2^n \rightarrow \ell_p^m$. Choosing a function f from $\mathcal{F}(p, n, m)$ is done as follows. First choose n orthonormal vectors in S^{m-1} uniformly at random² and let A be the $m \times n$ matrix whose columns are the orthonormal vectors. Then define $f(x) = \nu_p \cdot Ax$ where $\nu_p = \nu_p(n, m)$ is some normalization factor.

Clearly, any embedding function $f \sim \mathcal{F}(p, n, m)$ is linear and moreover has full column rank. In particular, by Fact 2.2, for any lattice \mathcal{L} in ℓ_2^n , $f(\mathcal{L})$ is a lattice in ℓ_p^m .

The following result shows that for large enough m , $\mathcal{F}(p, n, m)$ is a low-distortion family of embeddings from ℓ_2^n into ℓ_p^m .

Theorem 3.2 [FLM77] There exists $\nu_p(n, m)$ such that for all $p < \infty$, $\epsilon > 0$ and n there is m such that with probability at least $1 - 2^{-\Omega(n)}$, a randomly chosen embedding function $f \sim \mathcal{F}(p, n, m)$ satisfies that for all $x \in \mathbb{R}^n$

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p \leq (1 + \epsilon) \|x\|_2. \quad (2)$$

A choice of m that satisfies this is

$$m = \begin{cases} \epsilon^{-2}n & 1 \leq p < 2 \\ \epsilon^{-p}p^{-p/2}n^{p/2} & 2 \leq p < \infty \end{cases}.$$

For our applications, we will need to efficiently sample a function from $\mathcal{F}(p, n, m)$. Let us briefly describe how this can be done. In order to sample n orthonormal vectors, first choose n vectors independently from the uniform distribution on S^{m-1} . With probability 1, the vectors are linearly independent and we can therefore apply the Gram-Schmidt orthogonalization process to them.

²More precisely, from the unique rotationally invariant probability measure on the set of all n -tuples of orthogonal vectors in S^{m-1} .

By normalizing the resulting vectors, we obtain n orthogonal vectors with the correct distribution. Another issue to consider is the computation of the constant ν_p . It is possible to calculate it directly from its definition as given in [FLM77]. A simpler solution, which is sufficient for our purposes, is the following. Choose a random matrix A as in Definition 3.1. Then, for some fixed vector x , compute $\|x\|_2/\|Ax\|_p$. As can be seen from Equation (2), with probability $1 - 2^{-\Omega(n)}$, this value is in the range $(1 \pm \epsilon)\nu_p$ and we can therefore use it instead of ν_p with only a small increase in the distortion.

3.2 Embedding into ℓ_∞

Definition 3.3 (Embedding family for $p = \infty$) For any $m \geq n$, we define a distribution $\mathcal{F}(\infty, n, m)$ over embedding functions $f : \ell_2^n \rightarrow \ell_\infty^m$. Choosing a function f from $\mathcal{F}(\infty, n, m)$ is done as follows. Let z_1, z_2, \dots, z_m be m random vectors in \mathbb{R}^n chosen independently from the n -dimensional standard normal distribution (i.e., each coordinate of z_i is chosen independently according to $N(0, 1)$). We then define

$$f(x) = \nu_\infty \cdot (\langle z_1, x \rangle, \langle z_2, x \rangle, \dots, \langle z_m, x \rangle)$$

where $\nu_\infty = (2 \ln(m/\sqrt{\ln m}))^{-1/2}$ is a normalization factor.

As before $f \sim \mathcal{F}(\infty, n, m)$ is a linear transformation, moreover, it is easy to see that with probability 1, it has full rank (for $m \geq n$).

Definition 3.4 (Probability of expansion and contraction) For a family \mathcal{F} of embedding functions from ℓ_q^n to ℓ_p^m and a vector $v \in \ell_q^n$, we define the probability of expansion by more than $(1 + \epsilon)$ as

$$\Pr_{f \sim \mathcal{F}} \left[\|f(v)\|_p \geq (1 + \epsilon) \|v\|_q \right].$$

Similarly, we define the probability of contraction by more than $(1 - \epsilon)$ as

$$\Pr_{f \sim \mathcal{F}} \left[\|f(v)\|_p \leq (1 - \epsilon) \|v\|_q \right].$$

Theorem 3.5 [Ind03] For any small enough positive reals $\epsilon, \delta, \eta > 0$, the family $\mathcal{F}(\infty, n, m)$ for $m = (\delta^{-1} + \log \eta^{-1} + 1/\epsilon)^{O(1/\epsilon)}$ satisfies that its probability of contraction by more than $(1 - \epsilon)$ is at most η and its probability of expansion by more than $(1 + \epsilon)$ is at most δ .

For completeness, we include a proof in Appendix A.1. Notice that in order to keep $m \leq \text{poly}(n)$ we can take, for any constant $\epsilon > 0$, η to be exponentially small in n and δ only polynomially small. Partly for this reason, this embedding is sometimes known as an asymmetric embedding.

Notice also that, unlike Theorem 3.2, this theorem does not imply any bound on the distortion of the embedding. Indeed, it only tells us that for a randomly chosen f , ‘most points’ are not distorted by ‘too much’. However, for our reductions to work, we need the embedding to satisfy a stronger property: a randomly chosen f should not contract *any* $x \in \mathbb{R}^n$ by ‘too much’. This stronger property is shown in the following theorem.

Theorem 3.6 For any $\epsilon > 0$, any large enough n , and any $\delta > 0$, the family $\mathcal{F}(\infty, n, m)$ for $m = (n \log n + \delta^{-1} + 1/\epsilon)^{O(1/\epsilon)}$ satisfies the following two properties. Its probability of expansion by more than $1 + \epsilon$ is at most δ and with probability $1 - m2^{-\Omega(n)}$ over the choice of $f \sim \mathcal{F}(\infty, n, m)$, we have that for any $x \in \mathbb{R}^n$,

$$\|f(x)\|_\infty \geq (1 - \epsilon) \|x\|_2.$$

Proof: The proof starts by taking a fine enough net on S^{n-1} . We then choose η to be small enough, and show that with high probability on the choice of f , none of the points on the net contracts under f . Finally, we complete the proof of the theorem in Lemma 3.10 by showing that if f does not contract any point on the net, then it also does not contract any point in \mathbb{R}^n (in fact, this lemma requires some rough bound on the maximum expansion of f ; such a bound is given in Lemma 3.9). We start with two standard claims.

Claim 3.7 *For any $\theta > 0$, the unit sphere S^{n-1} has a θ -net of cardinality at most $(1 + 2/\theta)^n$.*

Proof: We follow [FLM77]. Let K_0 be a maximal set of points in S^{n-1} with pairwise distances greater than θ . This set is a θ -net. The ℓ_2 -balls of radius $\frac{\theta}{2}$ centered at each $y \in K_0$, (i.e., $y + B_2(\frac{\theta}{2})$) are disjoint and contained in $B_2(1 + \theta/2)$. By a volume argument it follows that $|K_0| \leq (1 + \theta/2)^n / (\theta/2)^n = (1 + 2/\theta)^n$. ■

Claim 3.8 *For a vector z chosen from the standard n -dimensional normal distribution,*

$$\Pr[\|z\|_2 \geq 2\sqrt{n}] \leq 2^{-n}.$$

Proof: By a change of variable, we see that

$$\int_{\mathbb{R}^n} e^{-\|z\|_2^2/8} dz = 2^n \int_{\mathbb{R}^n} e^{-\|z\|_2^2/2} dz.$$

On the other hand,

$$\int_{\mathbb{R}^n} e^{-\|z\|_2^2/8} dz \geq \int_{\|z\|_2 \geq 2\sqrt{n}} e^{-\|z\|_2^2/8} dz \geq 4^n \int_{\|z\|_2 \geq 2\sqrt{n}} e^{-\|z\|_2^2/2} dz$$

where we use that $e^{-\|z\|_2^2/8} / e^{-\|z\|_2^2/2} = e^{3\|z\|_2^2/8} \geq 4^n$ whenever $\|z\|_2 \geq 2\sqrt{n}$. The claim follows by combining the two equations. ■

The following lemma gives a rough bound on the maximum expansion of f .

Lemma 3.9 *Let f be chosen randomly from $\mathcal{F}(\infty, n, m)$. Then, with probability $1 - m2^{-\Omega(n)}$, for any $x \in \mathbb{R}^n$,*

$$\|f(x)\|_\infty \leq 2\sqrt{n} \|x\|_2.$$

Proof: We have

$$\|f(x)\|_\infty = \nu_\infty \max_i |\langle z_i, x \rangle| < \max_i \|z_i\|_2 \|x\|_2$$

where $\nu_\infty < 1$ is the normalization factor from Definition 3.3. By a union bound and Claim 3.8,

$$\Pr \left[\max_i \|z_i\|_2 \geq 2\sqrt{n} \right] \leq m \Pr [\|z_i\|_2 \geq 2\sqrt{n}] \leq m \cdot 2^{-n}.$$

■

Lemma 3.10 *Let K_0 be some θ -net on the sphere S^{n-1} , and let f be an embedding function that satisfies*

$$\forall y \in K_0 \quad \|f(y)\|_\infty \geq (1 - \epsilon) \|y\|_2 \quad \text{and} \quad \forall x \in \mathbb{R}^n \quad \|f(x)\|_\infty \leq 2\sqrt{n} \|x\|_2.$$

Then for any $x \in \mathbb{R}^n$

$$\|f(x)\|_\infty \geq (1 - (\epsilon + 2\sqrt{n}\theta)) \|x\|_2.$$

Proof: By linearity of f , it is enough to prove the theorem for any $x \in S^{n-1}$. So let $x \in S^{n-1}$ and let $y \in K_0$ be such that $\|x - y\|_2 \leq \theta$. By the triangle inequality and the linearity of f ,

$$\begin{aligned} \|f(x)\|_\infty &\geq \|f(y)\|_\infty - \|f(x) - f(y)\|_\infty \\ &= \|f(y)\|_\infty - \|f(x - y)\|_\infty \\ &\geq (1 - \epsilon) \|y\|_2 - 2\sqrt{n}\theta \\ &= (1 - \epsilon) - 2\sqrt{n}\theta. \end{aligned}$$

■

We now return to the proof of the theorem. Set $\eta = 2^{-10n \log n}$ and $m = (\delta^{-1} + n \log n + 1/\epsilon)^{O(1/\epsilon)}$ as in Theorem 3.5. The bound on the expansion probability of $\mathcal{F}(\infty, n, m)$ follows immediately from Theorem 3.5. In the rest of the proof we prove the second claim.

For $\theta = \frac{1}{n}$ let K_0 be a θ -net of cardinality $|K_0| \leq (1 + 2/\theta)^n$. Such a net exists by Claim 3.7. By Theorem 3.5, the probability that a randomly chosen $f \sim \mathcal{F}(\infty, n, m)$ contracts a point $y \in K_0$ by more than $(1 - \epsilon)$ is at most η . By the union bound we obtain that with probability at least $1 - \eta \cdot (1 + 2/\theta)^n \geq 1 - 2^{-\Omega(n)}$ a random f satisfies that for every $y \in K_0$

$$\|f(y)\|_\infty \geq (1 - \epsilon) \|y\|_2.$$

Moreover, by Lemma 3.9, we know that with probability $1 - m2^{-\Omega(n)}$, a random f satisfies that for all $x \in \mathbb{R}^n$, $\|f(x)\|_\infty \leq 2\sqrt{n} \|x\|_2$. We can now apply Lemma 3.10 and obtain that for all $x \in \mathbb{R}^n$

$$\|f(x)\|_\infty \geq \left(1 - \left(\epsilon + \frac{2}{\sqrt{n}}\right)\right) \|x\|_2 \geq (1 - 2\epsilon) \|x\|_2$$

for large enough n . Since ϵ is arbitrary, this completes the proof of the theorem. ■

3.3 Combining Several Embeddings in ℓ_∞

In the previous subsection, we showed that an embedding function chosen from the family $\mathcal{F}(\infty, n, m)$ does not contract any $x \in \mathbb{R}^n$ by more than $1 - \epsilon$. On the other hand, the bound we have on expansion is much weaker and only says that the probability of expansion is some small polynomial. As shown in [FLM77], this bound is in fact close to being tight, so there is no hope to improve the analysis substantially. For almost all of our applications, this small probability of expansion causes no difficulty and we can use the family $\mathcal{F}(\infty, n, m)$ directly.

The only exception is our application to CVPP. The problem is that in CVPP, we must fix an embedding of the lattice already in the preprocessing function. This embedding should be good for *any* query point $x \in \mathbb{R}^n$. We need an embedding that will not expand nor contract any $x \in \mathbb{R}^n$. Hence, it might seem at first that we have to use a low-distortion embedding. This is in fact what we do for $p < \infty$; however, by the result of [FLM77], such an embedding does not exist for $p = \infty$ unless m is exponential in n .

We solve this issue by combining several embedding functions. More precisely, we sample $k = O(n)$ embedding functions f_1, \dots, f_k from $\mathcal{F}(\infty, n, m)$, and show that with high probability, for *any* $x \in \mathbb{R}^n$ there is at least one f_i that does not expand x and none of f_1, \dots, f_k contracts x . In particular, this shows that for any $x \in \mathbb{R}^n$, $\min_i \|f_i(x)\|_\infty$ gives a good estimate of $\|x\|_2$. This estimate is the one we use in our application to CVPP.

Theorem 3.11 For any $\epsilon > 0$ and any large enough n the following holds. Let $k = 10n$ and $m = (n \log n)^{O(1/\epsilon)}$. Then for f_1, \dots, f_k chosen independently from $\mathcal{F}(\infty, n, m)$ we have that with high probability, for any $x \in \mathbb{R}^n$,

$$\exists i, \|f_i(x)\|_\infty \leq (1 + \epsilon) \|x\|_2 \quad \text{and} \quad (3)$$

$$\forall j, \|f_j(x)\|_\infty \geq (1 - \epsilon) \|x\|_2. \quad (4)$$

Proof: Choose $m = (n \log n)^{O(1/\epsilon)}$ as given by Theorem 3.6 with ϵ and $\delta = 1/n$. For any fixed j , Theorem 3.6 says that with probability $1 - 2^{-\Omega(n)}$, for any x we have $\|f_j(x)\|_\infty \geq (1 - \epsilon) \|x\|_2$. By the union bound, we obtain that (4) holds with probability $1 - k2^{-\Omega(n)} = 1 - 2^{-\Omega(n)}$.

In the remainder of the proof we prove (3). We first show that it holds for all y in some θ -net and then show how to extend it to any $x \in \mathbb{R}^n$. Let $\theta = 1/n$ and let K_0 be a θ -net on S^{n-1} with $|K_0| \leq (1 + 2/\theta)^n$, as given by Claim 3.7. Fix some $y \in K_0$. By Theorem 3.6, we know that the probability that a randomly chosen $f \sim \mathcal{F}(\infty, n, m)$ expands y by more than $(1 + \epsilon)$ is at most δ . Hence, the probability that all of f_1, \dots, f_k expand y by more than $(1 + \epsilon)$ is at most $\delta^k = 2^{-10n \log n}$. By the union bound it follows that the probability that there exists a $y \in K_0$ which every f_i expands by more than $(1 + \epsilon)$ is at most $(1 + 2/\theta)^n \cdot 2^{-10n \log n} \leq 2^{-8n \log n}$. Hence with probability $1 - 2^{-\Omega(n)}$, (3) holds for all $y \in K_0$. We complete the proof of the theorem in the following lemma. To apply it, notice that by Lemma 3.9, we have that with probability $1 - 2^{-\Omega(n)}$ no f_i expands any point by more than $2\sqrt{n}$.

Lemma 3.12 Let f_1, f_2, \dots, f_k be such that

$$\forall y \in K_0 \quad \exists i, \|f_i(y)\|_\infty \leq (1 + \epsilon) \|y\|_2 \quad \text{and} \quad \forall x \in \mathbb{R}^n \quad \forall j, \|f_j(x)\|_\infty \leq 2\sqrt{n} \|x\|_2$$

hold where K_0 is some θ -net on S^{n-1} . Then

$$\forall x \in \mathbb{R}^n, \exists i, \|f_i(x)\|_\infty \leq (1 + \epsilon + 2\sqrt{n}\theta) \|x\|_2.$$

Proof: By linearity of f_i , it is enough to prove the lemma for $x \in S^{n-1}$. So for any $x \in S^{n-1}$, let $y \in K_0$ be such that $\|x - y\|_2 \leq \theta$ and let i be such that $\|f_i(y)\|_\infty \leq (1 + \epsilon) \|y\|_2$. Using the triangle inequality and the linearity of f_i ,

$$\|f_i(x)\|_\infty \leq \|f_i(y)\|_\infty + \|f_i(x - y)\|_\infty \leq (1 + \epsilon) \|y\|_2 + 2\sqrt{n}\theta = 1 + \epsilon + 2\sqrt{n}\theta.$$

■

■

4 Reductions for Lattice Problems

In this section we prove Theorem 1.1. We start with a sequence of reductions for the ‘standard’ lattice problems, namely SVP, CVP, SIVP, and SBP and for each problem we consider both the search variant and the decision variant. Although these reductions are not exactly the same, they follow an almost identical pattern. The reader might therefore wish to read the first reduction and then skip directly to Section 4.4 where we describe the CVPP reduction (which is somewhat more subtle).

We now give an informal description of the idea in our reductions for standard lattice problems. Recall that the goal is to solve the problem in ℓ_2 using a solution in ℓ_p . In the case $p < \infty$, this follows almost immediately from the embedding of Theorem 3.2. Indeed, since an embedding function $f \sim \mathcal{F}(p, n, m)$ preserves distances almost exactly, we can apply it to our input lattice (and target vector, in the case of CVP). A solution to the resulting instance in ℓ_p is also a solution to the original question in ℓ_2 .

The case $p = \infty$ requires some more care. Here, we can no longer use Theorem 3.2, and instead we use Theorem 3.6 which gives only a weaker guarantee on the expansion. Luckily, essentially the same reduction works. The main observation is that all we care about, say in the case of SVP and CVP, is that one *fixed* vector does not expand by too much and that no other vector contracts by too much. For instance, in the case of SVP, we want the shortest vector not to expand. Other vectors can expand, but we must make sure that they do not contract. This is exactly what Theorem 3.6 guarantees.

In the $p < \infty$ case of CVPP we use the same technique as in the standard lattice problems, namely Theorem 3.2. For $p = \infty$, we can no longer use Theorem 3.6. The reason for this, as outlined in Section 3.3, is that the circuits created by the preprocessing function have to be hardwired with some sort of embedding function that is good for *all* possible target vectors. Theorem 3.11 provides us with such an embedding.

4.1 SVP reductions

Theorem 4.1 *For all $\epsilon > 0$ and for all $1 \leq p \leq \infty$, there is a randomized Cook reduction from SVP_γ^2 to $\text{SVP}_{\gamma'}^p$, where $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

Proof: The reduction is given an n -dimensional lattice \mathcal{L} . The reduction randomly chooses $f \sim \mathcal{F}(p, n, m)$ that is used to embed ℓ_2^n in ℓ_p^m where m is given by

$$m = \begin{cases} \epsilon^{-2}n & 1 \leq p < 2 & \text{by Theorem 3.2} \\ \epsilon^{-p}p^{-p/2}n^{p/2} & 2 \leq p < \infty & \text{by Theorem 3.2} \\ (n \log n)^{O(1/\epsilon)} & p = \infty & \text{by setting } \delta = 1/n \text{ in Theorem 3.6} \end{cases}. \quad (5)$$

The reduction then invokes the oracle of $\text{SVP}_{\gamma'}^p$ with $f(\mathcal{L})$ and obtains some $s' \in f(\mathcal{L})$. The output of the reduction is the lattice vector $f^{-1}(s') \in \mathcal{L}$ (where f^{-1} is defined on the image of f).

We now prove the correctness of the reduction. Let $w \in \mathcal{L}$ be a vector that achieves $\lambda_1^2(\mathcal{L})$. By Theorem 3.2 (for $p < \infty$) or Theorem 3.6 (for $p = \infty$) with probability at least $1 - 2^{-\Omega(n)}$ over $f \sim \mathcal{F}(p, n, m)$, for any $x \in \mathbb{R}^n$,

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p. \quad (6)$$

In addition, for the vector w , with probability at least $1 - 1/n$ (and in fact $1 - 2^{-\Omega(n)}$ for $p < \infty$),

$$\|f(w)\|_p \leq (1 + \epsilon) \|w\|_2. \quad (7)$$

So with probability at least, say, $1 - \frac{2}{n}$, both (6) and (7) hold. We complete the proof by showing that in such a case $\|f^{-1}(s')\|_2 \leq \gamma \lambda_1^2(\mathcal{L})$. Using $\|s'\|_p \leq \gamma' \cdot \lambda_1^p(f(\mathcal{L}))$,

$$\begin{aligned} \|f^{-1}(s')\|_2 &\stackrel{(6)}{\leq} \frac{1}{1-\epsilon} \|s'\|_p \leq \frac{1}{1-\epsilon} \gamma' \cdot \lambda_1^p(f(\mathcal{L})) \\ &\leq \frac{1}{1-\epsilon} \gamma' \|f(w)\|_p \stackrel{(7)}{\leq} \frac{1+\epsilon}{1-\epsilon} \gamma' \|w\|_2 = \gamma \cdot \lambda_1^2(\mathcal{L}) \end{aligned}$$

where the third inequality holds since $f(w) \in f(\mathcal{L})$. \blacksquare

Theorem 4.2 *For any $\epsilon > 0$, $1 \leq p \leq \infty$, $\gamma > 1$ there is a randomized Karp reduction from GapSVP_γ^2 to $\text{GapSVP}_{\gamma'}^p$ where $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

Proof: The input of the reduction is a GapSVP_γ^2 instance (\mathcal{L}, d) . The reduction randomly chooses $f \sim \mathcal{F}(p, n, m)$ where m is as in Equation (5). The output of the reduction is the $\text{GapSVP}_{\gamma'}^p$ instance $(f(\mathcal{L}), d')$ where $d' = (1 + \epsilon)d$.

We now prove the correctness of the reduction. First we consider a YES instance where there exists $u \in \mathcal{L}$ such that $\|u\|_2 \leq d$. From Theorem 3.6 (for $p = \infty$) or Theorem 3.2 (for $p < \infty$) we obtain that with probability at least $1 - 1/\text{poly}(n)$ (and in fact $1 - 2^{-\Omega(n)}$ for $p < \infty$) $\|f(u)\|_p \leq (1 + \epsilon)\|u\|_2 \leq (1 + \epsilon)d = d'$. Since $f(u) \in f(\mathcal{L})$, $\lambda_1^p(f(\mathcal{L})) \leq (1 + \epsilon)d = d'$ and hence the output is a YES instance for $\text{GapSVP}_{\gamma'}^p$.

We now consider a NO instance where $\lambda_1^2(\mathcal{L}) > \gamma d$. From Theorem 3.6 (for $p = \infty$) or Theorem 3.2 (for $p < \infty$), with probability $1 - 2^{-\Omega(n)}$ over the choice of $f \sim \mathcal{F}(\infty, n, m)$, f does not contract any $x \in \mathbb{R}^n$ by more than $(1 - \epsilon)$. In particular, f does not contract any $u \in \mathcal{L}$. Thus, if for all nonzero $u \in \mathcal{L}$, $\|u\|_2 > \gamma d$ then with probability at least $1 - 2^{-\Omega(n)}$,

$$\forall v \in f(\mathcal{L}), v \neq 0 \quad \|v\|_p \geq (1 - \epsilon) \|f^{-1}(v)\|_2 > (1 - \epsilon)\gamma d = \frac{1 - \epsilon}{1 + \epsilon}\gamma d = \gamma' d'$$

where we use that $f^{-1}(v) \in \mathcal{L}$. We have obtained a NO instance of $\text{GapSVP}_{\gamma'}^p$. \blacksquare

Combining Khot's result [Kho04] and Theorem 4.2 we obtain the following corollaries:

Corollary 4.3 *SVP in the ℓ_1 norm is hard to approximate to within any constant unless $\text{NP} \subseteq \text{BPP}$ and hard to approximate to within $2^{(\log n)^{1/2-\epsilon}}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{BPTIME}(2^{\text{poly}(\log n)})$.*

In Section 6.1 we review a result showing a deterministic embedding of ℓ_2^n into ℓ_1^m where $m = O(n^2)$ achieving constant distortion of $\sqrt{3}$. When using this result in the proof we obtain:

Corollary 4.4 *SVP in the ℓ_1 norm is hard to approximate to within any constant unless $\text{NP} \subseteq \text{RP}$ and hard to approximate to within $2^{(\log n)^{1/2-\epsilon}}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{RTIME}(2^{\text{poly}(\log n)})$.*

4.2 CVP reductions

Theorem 4.5 *For any $\epsilon > 0$, $\gamma > 1$ and any $1 \leq p \leq \infty$, there is a randomized Cook reduction from CVP_γ^2 to $\text{CVP}_{\gamma'}^p$ where $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

Proof: The reduction is given an n -dimensional lattice \mathcal{L} , and a target point t . The reduction randomly chooses $f \sim \mathcal{F}(p, n, m)$ that is used to embed ℓ_2^n in ℓ_p^m where m is as in Equation (5). The reduction then invokes the oracle of $\text{CVP}_{\gamma'}^p$ with $f(\mathcal{L})$ and $f(t)$ and obtains some $v' \in f(\mathcal{L})$. The output of the reduction is the lattice vector $f^{-1}(v') \in \mathcal{L}$.

We now prove the correctness of the reduction. Let $w \in \mathcal{L}$ be such that $\|w - t\|_2 = \text{dist}_2(t, \mathcal{L})$. By Theorem 3.2 (for $p < \infty$) or Theorem 3.6 (for $p = \infty$) with probability at least $1 - 2^{-\Omega(n)}$ over $f \sim \mathcal{F}(p, n, m)$, for any $x \in \mathbb{R}^n$,

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p. \tag{8}$$

In addition, for the vector $w - t$, with probability at least $1 - 1/n$ (and in fact $1 - 2^{-\Omega(n)}$ for $p < \infty$),

$$\|f(w - t)\|_p \leq (1 + \epsilon) \|w - t\|_2. \quad (9)$$

So with probability at least, say, $1 - \frac{2}{n}$, both (8) and (9) hold. We complete the proof by showing that in such case $\|f^{-1}(v') - t\|_2 \leq \gamma \cdot \text{dist}_2(t, \mathcal{L})$, where v' is the vector returned by the oracle. By definition, $\|v' - f(t)\|_p \leq \gamma' \cdot \text{dist}_p(f(t), f(\mathcal{L}))$. Indeed,

$$\begin{aligned} \|f^{-1}(v') - t\|_2 &\stackrel{(8)}{\leq} \frac{1}{1 - \epsilon} \|v' - f(t)\|_p \leq \frac{1}{1 - \epsilon} \gamma' \cdot \text{dist}_p(f(t), f(\mathcal{L})) \\ &\leq \frac{1}{1 - \epsilon} \gamma' \|f(w) - f(t)\|_p \stackrel{(9)}{\leq} \frac{1 + \epsilon}{1 - \epsilon} \gamma' \|w - t\|_2 = \gamma \cdot \text{dist}_2(t, \mathcal{L}) \end{aligned}$$

where the second inequality follows from the choice of v' , and the third inequality holds since $f(w) \in f(\mathcal{L})$. \blacksquare

Theorem 4.6 *For any $\epsilon > 0$, $\gamma > 1$ and any $1 \leq p \leq \infty$ there is a randomized Karp reduction from GapCVP_γ^2 to $\text{GapCVP}_{\gamma'}^p$, where $\gamma' = \frac{1 - \epsilon}{1 + \epsilon} \gamma$.*

Proof: The input of the reduction is a GapCVP_γ^2 instance (\mathcal{L}, t, d) , i.e., a lattice \mathcal{L} , a target point t and a parameter d . The reduction randomly chooses $f \sim \mathcal{F}(p, n, m)$ where m is as in Equation (5). The output of the reduction is the $\text{GapCVP}_{\gamma'}^p$ instance $(f(\mathcal{L}), f(t), d')$ where $f(\mathcal{L})$ is the embedded lattice, $f(t)$ is the embedded target vector, and d' is a parameter $d' = (1 + \epsilon)d$.

We now prove the correctness of the reduction. First we consider a YES instance of GapCVP_γ^2 where there exists $u \in \mathcal{L}$ such that $\|u - t\|_2 \leq d$. With probability at least $(1 - 1/n)$ (and even $1 - 2^{-\Omega(n)}$ for $p < \infty$), it follows from Theorem 3.2 (for $p < \infty$) or Theorem 3.6 (for $p = \infty$) that $\|f(u - t)\|_p \leq (1 + \epsilon) \|u - t\|_2 \leq (1 + \epsilon)d = d'$. In particular this implies that $(f(\mathcal{L}), f(t), d')$ is a YES instance for $\text{GapCVP}_{\gamma'}^p$.

We now consider a NO instance of GapCVP_γ^2 in which $\text{dist}_2(t, \mathcal{L}) > \gamma d$. By Theorem 3.2 (for $p < \infty$) or Theorem 3.6 (for $p = \infty$), we conclude that with probability at least $1 - 2^{-\Omega(n)}$, for all $v \in f(\mathcal{L})$, $\|v - f(t)\|_p \geq (1 - \epsilon) \|f^{-1}(v) - t\|_2 > (1 - \epsilon)\gamma d = \gamma' d'$ and thus $(f(\mathcal{L}), f(t), d')$ is a NO instance of $\text{GapCVP}_{\gamma'}^p$, and the correctness of the reduction follows. \blacksquare

4.3 SIVP and SBP reductions

We next show that the Shortest Independent Vector Problem SIVP under the ℓ_2 norm is not harder than SIVP in the ℓ_p norm, for $1 \leq p \leq \infty$. An essentially identical reduction works for the SBP problem and is omitted.

Theorem 4.7 *For any $\epsilon > 0$, $\gamma > 1$ and any $1 \leq p \leq \infty$ there is a randomized Karp reduction from GapSIVP_γ^2 to $\text{GapSIVP}_{\gamma'}^p$, where $\gamma' = \frac{1 - \epsilon}{1 + \epsilon} \gamma$.*

Proof: Given a GapSIVP_γ^2 instance (\mathcal{L}, d) , the reduction randomly chooses $f \sim \mathcal{F}(p, n, m)$ where m is chosen as in Equation (5) except for the case $p = \infty$:

$$m = \begin{cases} \epsilon^{-2} n & 1 \leq p < 2 & \text{by Theorem 3.2} \\ \epsilon^{-p} p^{-p/2} n^{p/2} & 2 \leq p < \infty & \text{by Theorem 3.2} \\ (n \log n)^{O(1/\epsilon)} & p = \infty & \text{by setting } \delta = 1/n^2 \text{ in Theorem 3.6} \end{cases}.$$

The reduction outputs the $\text{GapSIVP}_{\gamma'}^p$ instance $(f(\mathcal{L}), d')$ where $d' = (1 + \epsilon)d$.

We now analyze the reduction. We first consider a YES instance, where there is a set $U = \{u_1, u_2, \dots, u_n\}$ of n linearly independent vectors in \mathcal{L} whose length is at most d and show that there is a set $V = \{v_1, v_2, \dots, v_n\}$ of n linearly independent vectors in $f(\mathcal{L})$ whose length is at most $d(1 + \epsilon)$. For all i define $v_i = f(u_i)$. By union bound over all vectors in U it follows that with probability at least $1 - 1/n$ (and in fact $1 - 2^{-\Omega(n)}$ for $p < \infty$), for all i , $\|v_i\|_p \leq (1 + \epsilon) \|u_i\|_2 \leq (1 + \epsilon)d = d'$. Moreover, by Fact 2.2 v_1, \dots, v_n are linearly independence. Hence, $(f(\mathcal{L}), d')$ a YES instance of $\text{GapSIVP}_{\gamma'}^p$.

We now consider a NO instance where any set $U = \{u_1, u_2, \dots, u_n\}$ of n linearly independent vectors in \mathcal{L} is of length greater than γd and show that any set $V = \{v_1, \dots, v_n\}$ of n linearly independent vectors in $f(\mathcal{L})$, is of length greater than $\gamma' d(1 + \epsilon) = \gamma' d'$ where $\gamma' = \frac{1 + \epsilon}{1 - \epsilon} \gamma$. By Theorem 3.6 (for $p = \infty$) or Theorem 3.2 (for $p < \infty$) we obtain that with probability at least $1 - 2^{-\Omega(n)}$, f does not contract any $x \in \mathbb{R}^n$ by more than $1 - \epsilon$. So assume f has this property. Let $V = \{v_1, \dots, v_n\}$ be a set of n linearly independent vectors in $f(\mathcal{L})$. Since $f^{-1}(V)$ is a set of n linearly independent vectors in ℓ_2^n then there is i such that $\|f^{-1}(v_i)\|_2 > \gamma d$. Therefore, $\|v_i\|_p \geq (1 - \epsilon) \|f^{-1}(v_i)\|_2 > (1 - \epsilon)\gamma d = \gamma' d'$. We obtained a NO instance of $\text{GapSIVP}_{\gamma'}^p$. ■

Theorem 4.8 *For any $\epsilon > 0$, $\gamma > 1$ and any $1 \leq p \leq \infty$ there is a randomized Cook reduction from SIVP_{γ}^2 to $\text{SIVP}_{\gamma'}^p$, where $\gamma' = \frac{1 - \epsilon}{1 + \epsilon} \gamma$.*

The proof is very similar to that of Theorem 4.7 and is omitted.

4.4 CVPP reductions

We now present our reductions for the case of CVPP. As mentioned before, we focus on the non-uniform variant of the problem. It is possible to apply similar ideas to the uniform case. However, because of their probabilistic nature, our reductions do not quite fit the standard uniform formulation of CVPP. While this can be mended by an appropriate modification of the model, the resulting statement is not so elegant and we therefore choose to ignore this issue.

Before we go on, we should explain how to define a reduction between preprocessing problems. For example, a Cook reduction between $\text{GapCVPP}_{\gamma}^2$ to $\text{GapCVPP}_{\gamma'}^p$ in the non-uniform model is simply a solution to $\text{GapCVPP}_{\gamma}^2$ whose proof of existence is based on the assumption that a solution to $\text{GapCVPP}_{\gamma'}^p$ exists. A Karp reduction between $\text{GapCVPP}_{\gamma}^2$ and $\text{GapCVPP}_{\gamma'}^p$ is a function R that maps any lattice \mathcal{L} to a circuit of size polynomial in the size of \mathcal{L} . The circuit maps any instance (\mathcal{L}, t) of $\text{GapCVPP}_{\gamma}^2$ to an instance (\mathcal{L}', t') of $\text{GapCVPP}_{\gamma'}^p$, such that YES instances are mapped to YES instances and NO instances are mapped to NO instances. Crucially, the lattice \mathcal{L}' is fixed and is independent of t .

Theorem 4.9 *For any $\epsilon > 0$, $\gamma > 1$ and any $1 \leq p \leq \infty$, there is a deterministic reduction from $\text{GapCVPP}_{\gamma}^2$ to $\text{GapCVPP}_{\gamma'}^p$, where $\gamma' = \frac{1 - \epsilon}{1 + \epsilon} \gamma$.*

The proof for the case $p < \infty$ is very similar to previous proofs, and we describe it here briefly.

Lemma 4.10 *For any $\epsilon > 0$, $\gamma > 1$ and any $1 \leq p < \infty$, there is a deterministic Karp reduction from $\text{GapCVPP}_{\gamma}^2$ to $\text{GapCVPP}_{\gamma'}^p$, where $\gamma' = \frac{1 - \epsilon}{1 + \epsilon} \gamma$.*

Proof: Given an n -dimensional lattice \mathcal{L} , the reduction function R maps it to a circuit defined as follows. Let m be the value given by Theorem 3.2 for n, ϵ, p . Let f be a (fixed) embedding in $\mathcal{F}(p, n, m)$ for which Theorem 3.2 holds. The output of R is the circuit that maps an instance (\mathcal{L}, t) to the instance $(f(\mathcal{L}), f(t))$. The proof of the correctness is very similar to the proof of Theorem 4.5 and is therefore omitted. \blacksquare

Lemma 4.11 *For any $\epsilon > 0$, $\gamma > 1$, there is a deterministic Cook reduction from GapCVPP_γ^2 to $\text{GapCVPP}_{\gamma'}^\infty$ where $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

Proof: Assume there exists a solution P_∞ to $\text{GapCVPP}_\gamma^\infty$. We now describe a solution P_2 to GapCVPP_γ^2 . Assume we are given an n -dimensional lattice \mathcal{L} . Let f_1, f_2, \dots, f_k for $k = 10n$ be a sequence of embedding functions for which Theorem 3.11 holds. Notice that for each i , $P_\infty(f_i(\mathcal{L}))$ is a circuit that solves $\text{GapCVPP}_{\gamma'}^\infty$ instances of the form $(f_i(\mathcal{L}), t)$. The preprocessing function P_2 outputs the following circuit. Notice that we hardwire f_i and $P_\infty(f_i(\mathcal{L}))$ into the circuit.

Decoding circuit for GapCVPP_γ^2 with the lattice \mathcal{L} :

1. Input: t, d
2. For every $i = 1, \dots, k$ apply the circuit $P_\infty(f_i(\mathcal{L}))$ to the instance $(f_i(t), d')$ where $d' = (1 + \epsilon)d$.
3. If at least one of the circuits returns YES then output YES, else output NO.

We now prove the correctness of the reduction. First assume we are given a YES instance, that is, there is $u \in \mathcal{L}$ such that $\|u - t\|_2 \leq d$. From (3) it follows that there exists an i for which f_i does not expand $u - t$ by more than $(1 + \epsilon)$. Hence,

$$\|f_i(u) - f_i(t)\|_\infty = \|f_i(u - t)\|_\infty \leq (1 + \epsilon) \|u - t\|_2 \leq (1 + \epsilon)d = d'.$$

Since $f_i(u) \in f_i(\mathcal{L})$ we obtain that at least one of the circuits returns YES and therefore we output YES. Now assume that we are given a NO instance, that is, for all $u \in \mathcal{L}$, $\|u - t\|_2 > \gamma \cdot d$. By (4), we have that for any $x \in \mathbb{R}^n$, no f_j contracts x by more than $(1 - \epsilon)$. It then follows that

$$\forall j, \forall v \in f_j(\mathcal{L}), \quad \|v - f_j(t)\|_\infty = \left\| f_j(f_j^{-1}(v) - t) \right\|_\infty \geq (1 - \epsilon) \left\| f_j^{-1}(v) - t \right\|_2 > (1 - \epsilon)\gamma d = \gamma' d'.$$

Therefore all circuits return NO and we output NO. \blacksquare

As always, it is possible to extend this reduction to the search variant of the problem.

Theorem 4.12 *For any $\epsilon > 0$, $\gamma > 1$ and $1 \leq p \leq \infty$, there is a deterministic Cook reduction from CVPP_γ^2 to $\text{CVPP}_{\gamma'}^p$ for $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

As before, we first handle the case $p < \infty$ and then discuss the $p = \infty$ case.

Lemma 4.13 *For any $\epsilon > 0$, $\gamma > 1$, $p < \infty$ there is a Cook reduction from CVPP_γ^2 to $\text{CVPP}_{\gamma'}^p$, where $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

Proof: Assume there exists a solution P_p to $\text{CVPP}_{\gamma'}^p$. We now describe a solution P_2 to CVPP_γ^2 . Given an n -dimensional lattice \mathcal{L} , P_2 maps it to a circuit defined as follows. Let m be the value given by Theorem 3.2 for n, ϵ, p . Let f be a (fixed) embedding in $\mathcal{F}(p, n, m)$ for which Theorem 3.2 holds.

Decoding circuit for CVPP $_{\gamma}^2$ with the lattice \mathcal{L} :

1. Input: t, d
2. Apply the circuit $P_p(f(\mathcal{L}))$ to $f(t)$ and let $v \in f(\mathcal{L})$ be the result.
3. Output $f^{-1}(v)$.

We omit the proof of the correctness, as it is very similar to previous proofs. ■

Lemma 4.14 *For any $\epsilon > 0$, $\gamma > 1$ there is a Cook reduction from CVPP $_{\gamma}^2$ to CVPP $_{\gamma'}^{\infty}$ where $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$.*

Proof: Assume there exists a solution P_{∞} to CVPP $_{\gamma'}^{\infty}$. We now describe a solution P_2 to CVPP $_{\gamma}^2$. Assume we are given an n -dimensional lattice \mathcal{L} . Let f_1, f_2, \dots, f_k for $k = 10n$ be a sequence of embedding functions for which Theorem 3.11 holds. The preprocessing function P_2 outputs the following circuit.

Decoding circuit for CVPP $_{\gamma}^2$ with the lattice \mathcal{L} :

1. Input: t
2. For every $i = 1, \dots, k$ apply the circuit $P_{\infty}(f_i(\mathcal{L}))$ to $f_i(t)$ where $d' = (1 + \epsilon)d$ and let $v_i \in f_i(\mathcal{L})$ be the result.
3. Let $u_i = f_i^{-1}(v_i)$. Output the u_i that minimizes $\|t - u_i\|_2$.

We now prove correctness. Let $w \in \mathcal{L}$ be a vector in \mathbb{R}^n for which $\|w - t\|_2 = \text{dist}_2(t, \mathcal{L})$. By (3), there exists an i such that f_i does not expand $w - t$ by more than $1 + \epsilon$. For this i ,

$$\|f_i(t) - f_i(w)\|_{\infty} = \|f_i(t - w)\|_{\infty} \leq (1 + \epsilon) \|t - w\|_2 = (1 + \epsilon) \text{dist}_2(t, \mathcal{L}).$$

In particular,

$$\text{dist}_{\infty}(f_i(t), f_i(\mathcal{L})) \leq (1 + \epsilon) \text{dist}_2(t, \mathcal{L}) \tag{10}$$

By our assumption, v_i is such that

$$\|v_i - f_i(t)\|_{\infty} \leq \gamma' \cdot \text{dist}_{\infty}(f_i(t), f_i(\mathcal{L})) \tag{11}$$

Applying (4),

$$\|f_i^{-1}(v_i) - t\|_2 \leq \frac{1}{1 - \epsilon} \|v_i - f_i(t)\|_{\infty} \tag{12}$$

Combining (10), (11) and (12) we obtain

$$\|f_i^{-1}(v_i) - t\|_2 \leq \frac{1 + \epsilon}{1 - \epsilon} \gamma' \cdot \text{dist}_2(t, \mathcal{L}) = \gamma \cdot \text{dist}_2(t, \mathcal{L})$$

and the correctness of the reduction follows. ■

By combining Theorem 4.9 and [AKKV05] we conclude the following result:

Corollary 4.15 *The CVPP problem in the ℓ_p norm for $2 \leq p \leq \infty$ is hard to approximate to within any constant unless NP has polynomial size circuits (i.e., unless $\text{NP} \subseteq \text{P/poly}$), and to within $(\log n)^{1/2-\epsilon}$, for any $\epsilon > 0$, unless NP has quasi-polynomial size circuits.*

5 Norm Reduction via Dimension Preserving Random Rotation

In this section we present another embedding result with the property that it preserves the dimension. By plugging this embedding into the reductions shown in the previous section, one can obtain the dimension-preserving reductions of Theorem 1.2.

The idea in our dimension-preserving embedding is quite simple: the embedding simply consists of a random rotation. At first, this seems strange, since a random rotation is clearly not a low-distortion embedding: the ‘bad’ directions (in which ℓ_p differs considerably from ℓ_2) are obviously still there, they are just shifted around. However, it turns out that a random rotation is enough for lattice problems such as SVP where all that we care about is the length of the shortest vector as compared with the lengths of other vectors. Indeed, as we shall see below, with high probability, a random rotation brings the shortest vector (or any other fixed vector) into an area where the ℓ_p norm is almost as short as possible. Other lattice vectors can become longer, but this does not matter for us.

The following definition is essentially the special case of Definition 3.1 with $n = m$. For convenience, we do not include the normalization factor.

Definition 5.1 (Dimension Preserving Embedding family) *For any n, p , we define a distribution $\mathcal{F}(p, n)$ over embedding functions $f : \ell_2^n \rightarrow \ell_p^n$. Choosing a function f from $\mathcal{F}(p, n)$ is done by choosing n orthonormal vectors in S^{n-1} uniformly at random. We define $f(x) = Ax$ where A is the $n \times n$ matrix whose columns are the chosen vectors.*

We now give two embedding results, the first into ℓ_∞ and the second into ℓ_p for finite p . We define two probability measures. The first is μ , the standard Gaussian measure on \mathbb{R}^n , with density

$$(2\pi)^{-n/2} e^{-\|x\|_2^2/2}.$$

The second is σ which is the unique rotationally invariant measure on S^{n-1} .

Embedding ℓ_2 into ℓ_∞

Theorem 5.2 *Let f be chosen from $\mathcal{F}(\infty, n)$. Then for any $x \in \mathbb{R}^n$,*

$$\sqrt{\frac{1}{n}} \|x\|_2 \leq \|f(x)\|_\infty.$$

Moreover, for any $x \in \mathbb{R}^n$, with probability at least $1 - n^{-\Omega(1)}$

$$\|f(x)\|_\infty \leq \sqrt{\frac{32 \ln n}{n}} \|x\|_2.$$

Proof: The first inequality

$$\sqrt{\frac{1}{n}} \|x\|_2 \leq \|f(x)\|_\infty$$

follows from $\|f(x)\|_2 = \|x\|_2$. For the second inequality, notice that it is enough to prove it for x such that $\|x\|_2 = 1$. Moreover, for any such fixed x , $f(x)$ is distributed according to σ . It is therefore enough to prove that for x chosen according to σ , $\|x\|_\infty \leq \sqrt{(32 \ln n)/n}$ with probability $1 - n^{-\Omega(1)}$.

It turns out to be more convenient to work with the Gaussian measure μ instead of σ . To this end, notice that if $x \sim \mu$ then $x/\|x\|_2$ is distributed according to σ . This follows from the rotational invariance of μ and the uniqueness of σ . It is therefore enough to prove that an x chosen according to μ satisfies with probability at least $1 - n^{-\Omega(1)}$ that

$$\left\| \frac{x}{\|x\|_2} \right\|_\infty \leq \sqrt{32 \frac{\ln n}{n}}.$$

We prove this in two steps. We first establish that $\|x\|_2 \geq \sqrt{n}/2$ with high probability. We then complete the proof by showing that $\|x\|_\infty \leq \sqrt{8 \ln n}$ also holds with high probability.

Claim 5.3 *For a vector x chosen from the standard n -dimensional normal distribution,*

$$\Pr[\|x\|_2 \leq \sqrt{n}/2] \leq (1.37)^{-n}.$$

Proof: By a change of variable, we see that

$$\int_{\mathbb{R}^n} e^{-\|x\|_2^2/2} dx = 2^{-n} \int_{\mathbb{R}^n} e^{-\|x\|_2^2/2} dx.$$

On the other hand,

$$\int_{\mathbb{R}^n} e^{-\|x\|_2^2/2} dx \geq \int_{\|x\|_2 \leq \frac{1}{2}\sqrt{n}} e^{-\|x\|_2^2/2} dx \geq e^{-3/8n} \int_{\|x\|_2 \leq \frac{1}{2}\sqrt{n}} e^{-\|x\|_2^2/2} dx$$

where we use that $e^{-\|x\|_2^2/2}/e^{-\|x\|_2^2/2} = e^{-3/2\|x\|_2^2} \geq e^{-3/8n}$ whenever $\|x\|_2 \leq \frac{1}{2}\sqrt{n}$. The claim follows by combining the two equations and using $(2/e^{3/8})^n \geq (1.37)^n$. ■

Hence, to complete the proof it is enough to show that $\|x\|_\infty \leq \sqrt{8 \ln n}$ with high probability. Let $\Phi(t)$ denote $\frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-x^2/2} dx$, that is, the distribution function of the standard normal distribution $N(0, 1)$. We need the following standard estimate on the tail of the normal distribution (see, e.g., [Fel68] Chapter VII Lemma 2).

Fact 5.4 *For any $t > 2$,*

$$e^{-t^2/2}/2t \leq 1 - \Phi(t) \leq e^{-t^2/2}/t.$$

Let M_∞ be the median of $\|x\|_\infty = \max |x_i|$ with respect to μ , defined as the value of M_∞ for which

$$\mu\{x, \|x\|_\infty \leq M_\infty\} = 1/2 \quad \text{and} \quad \mu\{x, \|x\|_\infty \geq M_\infty\} = 1/2.$$

In other words, the cube $[-M_\infty, M_\infty]^n$ has Gaussian measure $\frac{1}{2}$. We start by calculating a bound on M_∞ . Since the Gaussian measure is a product measure, for any $M > 0$ we have

$$\mu([-M, M]^n) = \left(\frac{1}{\sqrt{2\pi}} \int_{-M}^M e^{-t^2/2} dt \right)^n.$$

By Fact 5.4 we obtain that

$$\frac{1}{\sqrt{2\pi}} \int_{-M}^M e^{-t^2/2} dt = 1 - 2(1 - \Phi(M)) \geq 1 - 2e^{-M^2/2}/M.$$

Thus for $M = \sqrt{2 \ln n}$,

$$\frac{1}{\sqrt{2\pi}} \int_{-M}^M e^{-t^2/2} dt \geq 1 - 2e^{-\ln n} / \sqrt{2 \ln n} = 1 - \frac{2}{n\sqrt{2 \ln n}}$$

and hence

$$\mu([-M, M]^n) \geq \left(1 - \frac{2}{n\sqrt{2 \ln n}}\right)^n > 1/2.$$

We thus obtain the bound $M_\infty \leq \sqrt{2 \ln n}$.

We show that with high probability a randomly chosen $x \sim \mu$ satisfies $\|x\|_\infty \leq M_\infty + \sqrt{2 \ln n}$. To do so we state Theorem 8.1 in [Bal97].

Lemma 5.5 *For a measurable set $A \subseteq \mathbb{R}^n$ with $\mu(A) = \frac{1}{2}$,*

$$\mu(A_\epsilon) \geq 1 - 2e^{-\epsilon^2/4} \tag{13}$$

where $A_\epsilon = \{x \mid \text{dist}_2(x, A) \leq \epsilon\}$.

By definition, $\mu(\{x \mid \|x\|_\infty \leq M_\infty\}) = 1/2$. Since $\|\cdot\|_\infty$ is 1-Lipschitz (i.e., $|\|x\|_\infty - \|y\|_\infty| \leq \|x - y\|_2$), we can apply Lemma 5.5 with $\epsilon = \sqrt{2 \ln n}$ to obtain that

$$\mu(\{x \mid \|x\|_\infty \leq M_\infty + \sqrt{2 \ln n}\}) \geq 1 - 2e^{-\epsilon^2/4} = 1 - \frac{2}{\sqrt{n}}.$$

Thus,

$$\Pr_{x \sim \mu} \left[\|x\|_\infty \leq \sqrt{8 \ln n} \right] \geq 1 - \frac{2}{\sqrt{n}}$$

and the theorem follows. ■

Embedding ℓ_2 into ℓ_p

Theorem 5.6 *Fix some $2 \leq p < \infty$ and let f be chosen from $\mathcal{F}(p, n)$. Then for all $x \in \mathbb{R}^n$,*

$$\|f(x)\|_p \geq n^{1/p-1/2} \|x\|_2.$$

Moreover, for any $x \in \mathbb{R}^n$, with probability at least $1 - 2^{-n^{\Omega(1)}}$ over the choice of f ,

$$\|f(x)\|_p \leq (1 + o(1)) \cdot c(p) \cdot n^{1/p-1/2} \|x\|_2$$

where $c(p) = \left(\frac{2^{\frac{1+p}{2}} \Gamma(\frac{1+p}{2})}{\sqrt{2\pi}} \right)^{1/p}$.

Proof: The first inequality

$$\|f(x)\|_p \geq n^{1/p-1/2} \|x\|_2$$

follows from $\|f(x)\|_2 = \|x\|_2$. As in Theorem 5.2, to prove the second inequality it is enough to show that for $x \sim \sigma$, $\|x\|_p \leq (1 + o(1))c(p)n^{1/p-1/2}$ with probability $1 - 2^{-n^{\Omega(1)}}$. We start by bounding the expectation of $\|x\|_p$ for $x \sim \sigma$. We will later see that the median of $\|x\|_p$ is close to this expectation and that with high probability, for a point $x \sim \sigma$, $\|x\|_p$ is close to that median.

The above expectation is given by $\int_{S^{n-1}} \|x\|_p d\sigma$. By using polar coordinates we can write

$$\int_{S^{n-1}} \|x\|_p d\sigma = \frac{\Gamma(n/2)}{\sqrt{2}\Gamma((n+1)/2)} \int_{\mathbb{R}^n} \|x\|_p d\mu(x) = (1 + o(1)) \frac{1}{\sqrt{n}} \int_{\mathbb{R}^n} \|x\|_p d\mu(x).$$

In addition,

$$\begin{aligned} \int_{\mathbb{R}^n} \|x\|_p d\mu(x) &= \int_{\mathbb{R}^n} \left(\sum_{i=1}^n |x_i|^p \right)^{1/p} d\mu(x) \\ &\leq \left(\int_{\mathbb{R}^n} \sum_{i=1}^n |x_i|^p d\mu(x) \right)^{1/p} && \text{(by Jensen inequality)} \\ &= \left(n \int_{\mathbb{R}} |x|^p d\mu(x) \right)^{1/p} && \text{(by symmetry)} \\ &= n^{1/p} \cdot \left(\frac{2^{\frac{1+p}{2}} \Gamma(\frac{1+p}{2})}{\sqrt{2\pi}} \right)^{1/p}. \end{aligned}$$

We thus obtain a bound on the expectation of $\|x\|_p$. We would like to obtain a similar bound on the median of $\|x\|_p$ defined as the value of M_p for which $\sigma(\{x \mid \|x\|_p \leq M_p\}) = 1/2$. For this, we use the following lemma which says that the median is close to the expectation.

Lemma 5.7 (Lemma 2.7 in [FLM77]) *Let M_p be the median of $\|\cdot\|_p$. Then for some absolute constant $c > 0$,*

$$\left| \int_{S^{n-1}} \|x\|_p d\sigma(x) - M_p \right| \leq \frac{c}{\sqrt{n}}.$$

Using this lemma, we obtain the bound

$$M_p \leq \int_{S^{n-1}} \|x\|_p d\sigma(x) + \frac{c}{\sqrt{n}} \leq (1 + o(1)) \cdot c(p) \cdot n^{1/p-1/2}.$$

The following concentration of measure lemma shows that for x chosen from σ , $\|x\|_p$ is very close to M_p with high probability.

Lemma 5.8 (Equation 2.6 in [FLM77]) *Let M_p be the median of $\|\cdot\|_p$. Then,*

$$\sigma \left(\left\{ x; \left| \|x\|_p - M_p \right| \leq \epsilon \right\} \right) \geq 1 - 4e^{-n\epsilon^2/2}.$$

We complete the proof by applying this lemma with $\epsilon = n^{1/p-1/2-\xi}$ for some small enough $\xi > 0$ to obtain

$$\Pr_{x \sim \sigma} \left[\|x\|_p \leq (1 + o(1))c(p)n^{1/p-1/2} \right] \geq 1 - 2^{-n^{\Omega(1)}}.$$

■

6 Other Results

6.1 Deterministic Reductions

All the embedding results we have used so far were probabilistic. There are a few known deterministic (explicit) embeddings. One such embedding is given in [Ber97]. It is an explicit embedding of ℓ_2^n into ℓ_1^m where $m = O(n^2)$ whose distortion is $\sqrt{3}$. Another explicit embedding is given by Indyk [Ind00] and is an embedding of ℓ_2^n into ℓ_1^m for $m = n^{O(\log n)}$ whose distortion is $1 + \epsilon$ for an arbitrarily small constant ϵ .

By using these embeddings in the reductions of Section 4, we can obtain *deterministic* reductions among lattice problems. More specifically, using the result of [Ber97], we obtain a deterministic polynomial reduction from the ℓ_2 norm to the ℓ_1 norm with a loss of $\sqrt{3}$ in the approximation factor. Using the result of [Ind00], we obtain a deterministic quasi-polynomial reduction with a loss of only $1 + \epsilon$ in the approximation factor.

6.2 Reduction from ℓ_p to ℓ_q for $q \in [1, p]$, $p \leq 2$

All the embeddings we have used so far were from the ℓ_2 norm. We now present an embedding, due to Johnson and Schechtman [FLM77], from ℓ_p into ℓ_q for $1 \leq q < p \leq 2$:

Theorem 6.1 ([JS82]) *For all $\epsilon > 0$ and $1 \leq q < p \leq 2$, ℓ_p^n can be embedded into ℓ_q^m with distortion $(1 + \epsilon)$ whenever $n \leq \beta m$ for some constant $\beta = \beta(p, q, \epsilon)$.*

By using this embedding in the reductions of Section 4, we obtain

Theorem 6.2 *For any $\epsilon > 0, \gamma \geq 1$ and $1 \leq q < p \leq 2$, there is a randomized reduction from SVP_{γ}^p to $\text{SVP}_{\gamma'}^q$, where $\gamma' = (1 - \epsilon)\gamma$. A similar result holds for other lattice problems such as CVP, CVPP, SIVP, SBP and for the decision variants.*

7 Discussion

We presented reductions between lattice problems in different norms. Our reductions have proven useful in deriving new inapproximability results. It would be interesting to find algorithmic applications of these reductions.

Another interesting open question is related to the covering radius problem. For a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ we define the covering radius as the smallest R such that any $x \in \mathbb{R}^n$ is within distance R of \mathcal{L} . More generally,

Definition 7.1 *The covering radius in the ℓ_p norm of a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is defined as*

$$\rho^p(\mathcal{L}) = \max_{x \in \text{span}(\mathcal{L})} \text{dist}_p(x, \mathcal{L}).$$

In the covering radius problem CRP_{γ}^p , our goal is to approximate $\rho^p(\mathcal{L})$ to within γ . Using the techniques of Section 4, one can obtain for any $p < \infty$ a reduction from CRP_{γ}^2 to $\text{CRP}_{\gamma'}^p$, where $\gamma' = (1 - \epsilon)\gamma$ for arbitrary small $\epsilon > 0$. However, it is not obvious how to extend this to the ℓ_{∞} norm. We leave this as an open question.

Acknowledgement

We thank Daniele Micciancio, Vitali Milman, and Muli Safra for their helpful comments.

References

- [AD97] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. 29th Annual ACM Symp. on Theory of Computing (STOC)*, pages 284–293. 1997.
- [Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, pages 99–108. ACM, New York, 1996.
- [Ajt98] M. Ajtai. Worst-case complexity, average-case complexity and lattice problems. In *Proceedings of the International Congress of Mathematicians, Vol. III (Berlin, 1998)*, pages 421–428 (electronic). 1998. ISSN 1431-0643.
- [AKKV05] M. Alekhnovich, S. Khot, G. Kindler, and N. K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing, 2005. Submitted.
- [AKS01] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, pages 601–610 (electronic). ACM, New York, 2001.
- [ALN05] S. Arora, J. R. Lee, and A. Naor. Euclidean distortion and the sparsest cut. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, pages 553–562. ACM, New York, 2005.
- [AMS99] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *J. Comput. System Sci.*, 58(1, part 2):137–147, 1999.
- [AR04] D. Aharonov and O. Regev. Lattice problems in NP intersect coNP. In *Proc. 45th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 362–371. 2004.
- [Bab86] L. Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. ISSN 0209-9683.
- [Bal97] K. Ball. An elementary introduction to modern convex geometry. In *Flavors of geometry*, volume 31 of *Math. Sci. Res. Inst. Publ.*, pages 1–58. Cambridge Univ. Press, Cambridge, 1997.
- [Ber97] B. Berger. The fourth moment method. *SIAM J. Comput.*, 26(4):1188–1207, 1997. ISSN 0097-5397.
- [BS99] J. Blömer and J.-P. Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999)*, pages 711–720 (electronic). ACM, New York, 1999.

- [CN99] J.-Y. Cai and A. Nerurkar. Approximating the SVP to within a factor $(1 + 1/\dim^\epsilon)$ is NP-hard under randomized reductions. *J. Comput. System Sci.*, 59(2):221–239, 1999. ISSN 0022-0000. 13th Annual IEEE Conference on Computation Complexity (Buffalo, NY, 1998).
- [Das99] S. Dasgupta. Learning mixtures of Gaussians. In *40th Annual Symposium on Foundations of Computer Science (New York, 1999)*, pages 634–644. IEEE Computer Soc., Los Alamitos, CA, 1999.
- [Din02] I. Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. *Theoretical Computer Science*, 285(1):55–71, 2002.
- [DKRS03] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003.
- [Dvo61] A. Dvoretzky. Some results on convex bodies and Banach spaces. In *Proc. Internat. Sympos. Linear Spaces (Jerusalem, 1960)*, pages 123–160. Jerusalem Academic Press, Jerusalem, 1961.
- [Fel68] W. Feller. *An introduction to probability theory and its applications. Vol. I.* Third edition. John Wiley & Sons Inc., New York, 1968.
- [FLM77] T. Figiel, J. Lindenstrauss, and V. D. Milman. The dimension of almost spherical sections of convex bodies. *Acta Math.*, 139(1-2):53–94, 1977. ISSN 0001-5962.
- [FM04] U. Feige and D. Micciancio. The inapproximability of lattice and coding problems with preprocessing. *J. Comput. System Sci.*, 69(1):45–67, 2004. ISSN 0022-0000.
- [Gau66] C. F. Gauss. *Disquisitiones arithmeticae*. Translated into English by Arthur A. Clarke, S. J. Yale University Press, New Haven, Conn., 1966.
- [GG00] O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. System Sci.*, 60(3):540–563, 2000. ISSN 0022-0000. 30th Annual ACM Symposium on Theory of Computing (Dallas, TX, 1998).
- [Ind00] P. Indyk. Stable distributions, pseudorandom generators, embeddings and data stream computation. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pages 189–197. IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.
- [Ind01] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)*, pages 10–33. IEEE Computer Soc., Los Alamitos, CA, 2001.
- [Ind03] P. Indyk. Better algorithms for high-dimensional proximity problems via asymmetric embeddings. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003)*, pages 539–545. ACM, New York, 2003.
- [JS82] W. B. Johnson and G. Schechtman. Embedding l_p^m into l_1^n . *Acta Math.*, 149(1-2):71–85, 1982. ISSN 0001-5962.

- [Kan87] R. Kannan. Minkowski’s convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, 1987. ISSN 0364-765X.
- [Kho03] S. Khot. Hardness of approximating the shortest vector problem in high L_p norms. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*. 2003.
- [Kho04] S. Khot. Hardness of approximating the shortest vector problem in lattices. In *Proc. 45rd Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 126–135. 2004.
- [Len83] H. W. Lenstra, Jr. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983. ISSN 0364-765X.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982. ISSN 0025-5831.
- [LLR95] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. ISSN 0209-9683.
- [LO85] J. C. Lagarias and A. M. Odlyzko. Solving low-density subset sum problems. *J. Assoc. Comput. Mach.*, 32(1):229–246, 1985. ISSN 0004-5411.
- [MG02a] D. Micciancio and S. Goldwasser. *Complexity of lattice problems*. The Kluwer International Series in Engineering and Computer Science, 671. Kluwer Academic Publishers, Boston, MA, 2002. ISBN 0-7923-7688-9. A cryptographic perspective.
- [MG02b] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, March 2002.
- [Mic01] D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035 (electronic), 2001. ISSN 1095-7111.
- [Reg03a] O. Regev. Improved inapproximability of lattice and coding problems with preprocessing. In *Proc. of 18th IEEE Annual Conference on Computational Complexity (CCC)*, pages 363–370. 2003.
- [Reg03b] O. Regev. New lattice based cryptographic constructions. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 407–416 (electronic). ACM, New York, 2003.
- [Sch85] C.-P. Schnorr. A hierarchy of polynomial time basis reduction algorithms. In *Theory of algorithms (Pécs, 1984)*, volume 44 of *Colloq. Math. Soc. János Bolyai*, pages 375–386. North-Holland, Amsterdam, 1985.
- [vEB81] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, Math Inst., University Of Amsterdam, Amsterdam, 1981.

A Proof of Theorem 3.5

Let z_1, z_2, \dots, z_m be m random vectors in \mathbb{R}^n . Each coordinate of z_i is chosen i.i.d. according to the standard normal distribution $N(0, 1)$. Let $g(x) = (\langle z_1, x \rangle, \langle z_2, x \rangle, \dots, \langle z_m, x \rangle) \in \mathbb{R}^m$ for $m = (\delta^{-1} + \log \eta^{-1} + 1/\epsilon)^{O(1/\epsilon)}$ where δ and η are given in the following theorem.

Theorem A.1 For $T = (2 \ln(m/\sqrt{\ln m}))^{1/2}$ we have that for any $x \in \mathbb{R}^n$,

$$\Pr[\|g(x)\|_\infty \leq (1 - \epsilon)T \|x\|_2] \leq \eta$$

$$\Pr[\|g(x)\|_\infty > (1 + \epsilon)T \|x\|_2] \leq \delta$$

Theorem 3.5 follows by taking $f(\cdot) = \frac{1}{T}g(\cdot)$. Our proof closely follows that of Indyk. We first need to prove an auxiliary lemma that shows that the maximum of standard normally distributed independent variables is concentrated in some range.

Lemma A.2 Let Y_1, Y_2, \dots, Y_m be i.i.d. normal variables and let $Z = \max(|Y_1|, |Y_2|, \dots, |Y_m|)$. Then for any $\delta, \eta \in (0, 1)$

- $\Pr[Z < T_1] \leq \eta$ where $T_1 = \sqrt{2 \ln(\frac{4m}{\ln \eta^{-1} \sqrt{2 \ln m}})}$.

- $\Pr[Z > T_2] \leq \delta$ where $T_2 = \sqrt{2 \ln(2m\delta^{-1})}$.

Proof: Let $\Phi(t)$ denote $\frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-x^2/2} dx$, that is, the distribution function of the standard normal distribution $N(0, 1)$.

Applying union bound on Fact 5.4 we obtain

$$\Pr[Z > T_2] \leq m \cdot 2(1 - \Phi(t)) \leq 2me^{-T_2^2/2} \leq \delta.$$

For the other claim, by using again Fact 5.4,

$$\Pr[Z < T_1] = (1 - 2(1 - \Phi(T_1)))^m \leq (1 - e^{-T_1^2/2}/T_1)^m \leq \exp(-me^{-T_1^2/2}/T_1) \leq \eta$$

where the last inequality follows from our choice of T_1 . ■

Proof of the theorem: Since g is linear, we can assume without loss of generality that $\|x\|_2 = 1$. By the rotational symmetry of the normal distribution, each coordinate of $g(x)$ is distributed as $N(0, 1)$. We prove the theorem by showing that for large enough $m = (\delta^{-1} + \log \eta^{-1} + 1/\epsilon)^{O(1/\epsilon)}$, $T_1 \geq (1 - \epsilon)T$ and $T_2 \leq (1 + \epsilon)T$. For the first bound,

$$\begin{aligned} \frac{T_1}{T} &= \sqrt{\frac{2 \ln(\frac{2m}{\ln \eta^{-1} \sqrt{2 \ln m} \cdot 1/2})}{2 \ln(m/\sqrt{\ln m})}} = \sqrt{\frac{2 \ln m + 2 \ln 2 - 2 \ln \ln \eta^{-1} - \ln 2 - \ln \ln m + 2 \ln 2}{2 \ln m - \ln \ln m}} \\ &\geq \sqrt{\frac{2 \ln m + 2 \ln 2 - 2 \ln \ln \eta^{-1} - \ln 2 - \ln \ln m + 2 \ln 2}{2 \ln m}} = \sqrt{1 - \frac{2 \ln \ln \eta^{-1} + \ln \ln m - 2 \ln 2 - \ln 2}{2 \ln m}} \end{aligned}$$

We only need to show that

$$2 \ln \ln \eta^{-1} + \ln \ln m - 2 \ln 2 - \ln 2 \leq \epsilon(2 \ln m).$$

Since $m > (\ln \eta^{-1})^{O(1/\epsilon)}$, the latter holds and we obtain that $\frac{T_1}{T} \geq (1 - \epsilon)$. For the second bound it is enough to show that $\frac{T_2}{T} \leq 1 + \epsilon$

$$\frac{T_2}{T} = \sqrt{\frac{2 \ln(2m\delta^{-1})}{2 \ln(m/\sqrt{\ln m})}} = \sqrt{\frac{2 \ln m + 2 \ln 2 + 2 \ln \delta^{-1}}{2 \ln m - \ln \ln m}} = \sqrt{1 + \frac{\ln \ln m + 2 \ln 2 + 2 \ln \delta^{-1}}{2 \ln m - \ln \ln m}}$$

Since $m > (\delta^{-1})^{O(1/\epsilon)}$ it holds that

$$\ln \ln m + 2 \ln 2 + 2 \ln \delta^{-1} \leq \epsilon(2 \ln m - \ln \ln m),$$

and we obtained $\frac{T_2}{T} \leq (1 + \epsilon)$. ■