

Acknowledgments

I would like to express my deepest gratitude to my advisor Dr. Oded Regev for his devoted guidance on this research. This work could not have been accomplished without him. I would also like to deeply thank Prof. Muli Safra for always being supportive and optimistic about research and life.

I would also like to thank Michal and Avi, who are always there when I need them, my family and friends for putting up with me all these years.

Abstract

We show 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 Pre-processing. The reductions are based on norm embeddings from ℓ_2 into ℓ_p . Some of the new results include hardness within any constant for the Shortest Vector Problem in the ℓ_1 norm, hardness of approximation within any constant for the Closest Vector Problem with Pre-processing in the ℓ_∞ norm, this is the first known hardness of approximation result for this problem in the ℓ_∞ norm.

Contents

1	Introduction	7
1.1	Lattices	7
1.2	Embedding	10
1.3	Our techniques	12
2	Preliminaries	13
3	Embedding	17
3.1	Embedding in the ℓ_p norm for $p < \infty$	17
3.2	Embedding in the ℓ_∞ norm	18
3.3	Combining Several Embeddings in ℓ_∞	20
4	Applications - reductions among lattice problems	23
4.1	SVP reductions	23
4.2	CVP reductions	25
4.3	CVPP reductions	27
4.4	SIVP and SBP reductions	30
5	Extensions	33
5.1	Deterministic Reductions	33
5.2	Reduction from ℓ_p to ℓ_q for $q \in [1, p]$, $p \leq 2$	33
5.3	Full rank	33
5.4	Norm Reduction via Dimension Preserving Random Rotation	34
6	Open Problems	39
	Bibliography	40
7	Appendix	45
7.1	Proof of Theorem 3.2.3	45

Chapter 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 say that n is the *rank* of the lattice and that m is 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 closest lattice vector 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 $\|x\|_\infty = \max_{i=1}^n |x_i|$.

One may also consider the approximation version of these problems, in which one would like to find not necessarily the shortest vector, but one that is known to be longer than the shortest vector by at most some given *approximation factor*. Similarly, for CVP, one would like to find a lattice vector whose distance to the target point is at most some given approximation factor. The complexity of those problems was investigated extensively, yielding ingenious algorithms on the one hand [LLL82], and computational hardness results on the other hand [DKRS03], [Kho04]. The complexity of these problems was also investigated for the average case [Ajt96], yielding cryptosystems [AD97], [Reg03b]. For a more thorough introduction to lattices see [MG02].

The earliest work on SVP is attributed to Gauss [Gau66] who gave in 1801 an algorithm that solves 2-dimensional SVP. The best known algorithm for computing the exact SVP was obtained by Kannan [Kan87], whose running time is $2^{O(n \log n)}$ where n is the dimension. A randomized algorithm whose running time is $2^{O(n)}$ that computes exactly SVP in the ℓ_2 norm was presented by Ajtai et al. [AKS01].

Back to the n -dimensional case and polynomial time limits, a major algorithmic step was attained by Lenstra, Lenstra and Lovász (LLL) [LLL82] in 1982 who showed a polynomial algorithm

for approximating n -dimensional SVP to within a factor of $2^{n/2}$. LLL's celebrated algorithm had a significant impact on many important problems such as integer programming [LLL82], solving low-density subset-sum problems, breaking knapsack based cryptosystems [LO85], approximation algorithms for CVP [Bab86], simultaneous diophantine approximation, and factoring polynomials over the rationals [LLL82]. Schnorr [Sch85] presented an improved algorithm with an approximation factor of $2^{n(\log \log n)^2 / \log n}$. Using Schnorr's techniques and [AKS01], this approximation factor was later improved by [AKS01] to an approximation factor of $2^{n \log \log n / \log n}$.

As for hardness of approximation, van Emde Boas [vEB81] showed in 1981 that SVP in the ℓ_∞ norm is NP-hard. For many years SVP was not known to be hard in any other norm. In 1998, Ajtai [Ajt98] showed hardness of SVP in the ℓ_p norm for any $p \geq 1$ assuming $\text{NP} \not\subseteq \text{RP}$ and also hardness of approximation to within $(1 + \frac{1}{2^{O(n)}})$. This was later improved, under the same complexity assumption, by Cai and Nerurkar [CN99] who showed hardness of approximation to within $(1 + \frac{1}{n^\epsilon})$. The first constant factor hardness result for SVP in the ℓ_p norm for any $1 \leq p < \infty$ was obtained by Micciancio [Mic01] who proved for every $\epsilon > 0$ a hardness of approximation factor of $\sqrt[p]{2} - \epsilon$ under the assumption that $\text{NP} \not\subseteq \text{RP}$. Dinur [Din02] boosted the hardness of SVP in the ℓ_∞ norm to $n^{\frac{O(1)}{\log \log n}}$, might be an indication that ℓ_p might be harder as p grows to infinity. A breakthrough result by Khot [Kho04] shows that SVP in the ℓ_p norm for any $1 < p \leq \infty$ is hard to approximate to within any constant unless $\text{NP} \subseteq \text{RP}$. Assuming a stronger assumption, that $\text{NP} \not\subseteq \text{RTIME}(2^{\text{poly}(\log n)})$, SVP in the ℓ_p norm is hard to approximate to within $2^{(\log n)^{1/2-\epsilon}}$ for any $\epsilon > 0$. Khot's result does not apply for SVP in the ℓ_1 norm where the best hardness known is Micciancio's $2 - \epsilon$. The techniques developed herein allow us to stretch Khot's hardness result to the ℓ_1 norm.

From the above results, the behavior of the hardness of SVP as a function of p seems unclear. On one hand, Micciancio's hardness factor decreases with p and becomes trivial for $p = \infty$. This might suggest that SVP is "easier" as p increases. On the other hand, Khot's recent result does not improve the hardness factor for $p = 1$, suggesting that SVP in the ℓ_1 norm might be easier than in other norms. Moreover, an earlier result of Khot [Kho03] shows that the hardness of these problems increases as p grows that might indicate that SVP gets harder as p increases. Finally, the NP-hardness of SVP in the ℓ_2 norm was known long before its hardness in any other ℓ_p norm [vEB81]. Our result clarifies the picture by showing that Khot's recent hardness result for the ℓ_2 norm implies all other results of other norms.

Much of the recent interest in SVP is a result of the work of Ajtai [Ajt96]. Ajtai proved that a certain average-case problem (namely, a variant of subset sum) is as hard on the *average-case* as approximating SVP in the *worst-case*. In other words, if there exists an algorithm that solves the average-case problem with some non-negligible probability then there exists an algorithm that approximates *any instance* of SVP (the approximation ratios here are of the form n^c for some constant c). This result has an important cryptographic application: it immediately implies the existence of a certain one-way function (based on subset-sum) whose hardness is based on the *worst-case* hardness of a lattice problem. Essentially all other known one-way function are based on an *average-case* assumption, such as the assumption that it is hard to factor a number taken from some distribution.

Later, Ajtai and Dwork [AD97] constructed a public key cryptosystem based on hardness of approximating unique-SVP. There are other cryptosystems relying on worst-case complexity of lattice problems such as [Kan87], [CN97], [Cai99] and [Reg03b] who constructed a cryptosystem based on $n^{1.5}$ -hardness of unique-SVP. Hopefully such hardness result should not be obtained

as indicated in [Hås88], [LLS90], [Ban93], [AR04] results showing that if such hardness will be obtained, then the polynomial time hierarchy collapses to a certain level. The latter showed that approximating SVP within a factor of \sqrt{n} is in $\text{NP} \cap \text{coNP}$.

Other Lattice Problems More is known about the hardness of the CVP problem. Already in 1981, Van Emde Boas [vEB81] showed that CVP is NP-hard under every 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 hard to approximate to within $n^{(c/\log \log n)}$ for some constant $c > 0$. For $p < \infty$, this result is due to Dinur, Kindler, Raz, and Safra [DKRS03]; Dinur [Din02] used those techniques to obtain the result for $p = \infty$.

The hardness of CVP is the core of many lattice based cryptographic systems. Breaking the cryptosystem can be done by solving CVP where the lattice is a public key. It is then natural to raise the question whether there is a representation of the lattice, a new basis for the lattice, in which CVP is efficiently solvable (like the Korkin-Zolotareff basis to SVP). Answering this question positively implies that if one invested time and computation resources she might break the cryptosystem. Alternatively we can ask if there are lattices such that given unlimited computation time for pre-processing, one can answer efficiently queries, where a query is a target vector t . This problem is called CVPP. Micciancio and Feige showed hardness of $\sqrt[5]{3} - \epsilon$. Regev [Reg03a] proved a hardness factor of $\sqrt[3]{3} - \epsilon$ for CVPP in the l_p norm. Alekhnovich, Khot, Kindler and Vishnoi [AKKV05] recently improved this result showing that CVPP is NP-hard to approximate by any constant. They also showed hardness of $(\log n)^{1/p-\epsilon}$ for any $p \geq 1$ under the assumption that $\text{NP} \not\subseteq \text{DTIME}(2^{\text{poly} \log(n)})$. Our method improves the two results and also shows a hardness of approximation factor for CVPP in the l_∞ norm (in both papers [Reg03a] and [AKKV05] there is no gap for l_∞).

Two more lattice problems were investigated due to Ajtai's result [Ajt96]. The *Shortest Independent Vector Problem* (SIVP) is the problem of finding n linearly independent vectors while minimizing the longest vector in the set. The approximation version of this problem is to find such vectors for which the length of the longest vector is at most γ times the best possible length.

The second problem is the *Shortest Basis Problem* (SBP), this is the problem of finding a basis of \mathcal{L} while minimizing the longest vector in the basis. The approximation version of this problem is to find such basis for which the length of the longest basis vector is at most γ times the best possible length.

Ajtai [Ajt96] showed a reduction from the worst case of SVP or SIVP or SBP to average case of SVP. Ajtai's result raised question, how hard these problems are. Blömer and Seifert [BS99] showed that SIVP and SBP under the l_2 norm is NP-hard to approximate within any constant and under the complexity assumption $\text{NP} \not\subseteq \text{DTIME}(2^{\text{poly} \log(n)})$ they showed hardness of approximation within $2^{\log^{1-\epsilon} n}$ for any $\epsilon > 0$. Our method shows that such hardness of approximation factor can be achieved for SIVP and SBP in the l_∞ norm (this is the first hardness of approximation result known for SIVP and SBP in the l_∞ norm).

Our Results

Our main theorem is the following.

Theorem 1.1.1 *For any $1 \leq p \leq \infty$ there exists a randomized polytime reduction from SVP in the l_2 norm to SVP in the l_p norm. A similar result holds for CVP, SIVP, SBP. For any $1 \leq p \leq \infty$*

there exists a deterministic reduction from CVPP in the ℓ_2 norm to CVPP in the ℓ_p norm.

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 CVPP in the ℓ_p norm for $2 \leq p \leq \infty$ is hard to approximate
 - to within any constant unless $\text{NP} \subseteq \text{P}$, and
 - to within $(\log n)^{1/2-\epsilon}$, for any $\epsilon > 0$, unless $\text{NP} \subseteq \text{DTIME}(2^{\text{poly} \log(n)})$.

Previously, no hardness of approximation was known for this problem in the ℓ_∞ norm. For $1 \leq \ell < \infty$, this problem was known to be hard to approximate

- to within any constant unless $\text{NP} \subseteq \text{P}$ [AKKV05], and
- to within $(\log n)^{1/p-\epsilon}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{DTIME}(2^{\text{poly} \log(n)})$ [AKKV05].
- 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 $\text{NP} \subseteq \text{P}$, and
 - to within $2^{(\log n)^{1-\epsilon}}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{DTIME}(2^{\text{poly} \log(n)})$.

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

It also follows from the theorem that it is enough to find approximation algorithms for SVP, SIVP, or CVP in, say, the ℓ_1 norm or the ℓ_∞ norm. Any such algorithm yields an algorithm in the ℓ_2 norm. Obtaining algorithms in the ℓ_1 and ℓ_∞ 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 .

1.2 Embedding

In this section we review some of the embedding techniques that are used in our reductions and other embedding results that are the base of our embedding result. The trivial bound states that $\|x\|_p \leq \|x\|_2 \leq n^{1/2-1/p} \|x\|_p$ for $p > 2$ and that $n^{1/p-1/2} \|x\|_p \leq \|x\|_2 \leq \|x\|_p$ for $1 \leq p < 2$. This bound does not preserve hardness factor results that are sensitive to polynomial factors, thus useless for achieving hardness factors for lattice problems such as SVP, SIVP, CVP, CVPP. To do better we use embedding techniques.

For an n -dimensional normed space $X = (\mathbb{R}^n, \|\cdot\|_p)$ equipped with the ℓ_p norm and an m -dimensional normed space $Y = (\mathbb{R}^m, \|\cdot\|_q)$ equipped with the ℓ_q norm, we call an embedding function, $f: X \rightarrow Y$ an *embedding* of X in Y . If an embedding function, $f: X \rightarrow Y$ satisfies that for any two points $x, x' \in X$,

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

then we say that the *distortion* of $f(\cdot)$ is $\frac{1+\epsilon}{1-\epsilon}$.

Low distortion embeddings between metric spaces has fascinated mathematicians for many decades. One very important embedding is embedding Banach space (normed complete space) of finite dimension into other Banach space of finite dimension. The important result of Figiel, Lindenstrauss and Milman [FLM77] showed, based on the celebrated theorem of Dvoretzky [Dvo61], that for a given m -dimensional Banach space X and for some suitable $n < m$ (often not much smaller than m), most n -dimensional subspaces of X are $(1 + \epsilon)$ -close to ℓ_2^n . In particular, ℓ_2^n can be randomly embedded into ℓ_p^m where the dimension m is not much larger than n (for $p < \infty$). An example of such a random embedding is shown later. Figiel et al. [FLM77] used the important phenomenon of concentration of measure in high dimensional spaces. In high dimensions most of the measure is concentrated around the median. Let $(X, \|\cdot\|)$ be an n -dimensional Banach space with some norm and let $\|\cdot\|_2$ be the inner product norm on X where $a \|x\|_2 \leq \|x\| \leq b \|x\|_2$, for all $x \in X$ and for suitable $0 < a \leq b < \infty$. Let M_r be the median of $r(x) = \|x\|$ on the n -dimensional sphere, S^{n-1} , with respect to the rotation invariant measure on S^{n-1} , formally $\mu\{x, \|x\| \geq M_r\} \geq 1/2$ and $\mu\{x, \|x\| \leq M_r\} \geq 1/2$. Let $A = \{x, \|x\| = M_r\}$ and let $A_\epsilon = \{x, \text{dist}_2(x, A) \leq \epsilon\}$ where $\text{dist}_2(\cdot)$ is the standard Euclidean norm ℓ_2 and $\text{dist}_2(x, A) \leq \epsilon$ means that there is a point in A whose distance from x is at most ϵ . Surprisingly by taking the set A_ϵ we almost take all $x \in X$ up to a set of exponentially small measure. Next they show that the norm of every $x \in S^{n-1}$ is bounded by $|M_r - b\epsilon|$. They also show that for suitable n and m , most isometries U from ℓ_2^n onto $(X, M_r^{-1} \|\cdot\|_2)$ satisfies $M_r - b\epsilon \leq \|Ux\| \leq M_r + b\epsilon$. It then follows that most n -dimensional slices of X are close up to $b\epsilon$ to ℓ_2^n after normalizing in M_r . For a friendly survey on this topic see [Bal97] Section 8 and Section 9.

Johnson and Schechtman [JS82] complemented in some sense the result of [FLM77]. They showed that for every $1 \leq q < p \leq 2$, ℓ_p^m can be embedded into ℓ_q^m with distortion $1 + \epsilon$ for $m \leq \beta n$ where $\beta(\epsilon, p, q)$ is a constant independent of n .

Recently, embedding of norms was found extremely useful in many applications in computer science. It was then used in various fields, such as approximation algorithms for geometric problems [Ind01], [Ind00a], learning [BDES02]. More applications can be found in the survey of Indyk [Ind01].

In some problems such as finding the diameter of a set of points and finding the maximum spanning tree, the problem in the ℓ_∞ norm seems simpler than in the ℓ_2 norm. However, [FLM77] showed that for any constant $c > 1$, embedding ℓ_2^n into ℓ_∞^m , m is of size exponentially in n to achieve distortion c . They also showed that this is tight, i.e., that any embedding of ℓ_2^n into ℓ_∞^m with distortion c must satisfy $m = \exp(n)$. A way to bypass this problem was suggested by Indyk in [Ind03] via *asymmetric* random embeddings. The embedding is asymmetric in terms of the probability of contraction and expansion. This embedding takes multiple random mappings $z_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $z_i(x) = \langle x, z_i \rangle$. Each random mapping is a vector whose coordinates are chosen according to the standard normal distribution and uses the asymmetric property of the tail of $\max_i \langle z_i, x \rangle$. The left tail is sharp whereas the right tail is not very sharp. This embedding was used in algorithms of the Nearest Neighbor Problem, Farthest Neighbor Problem in [Ind03].

The embedding results mentioned so far are probabilistic and therefore do not imply an explicit embedding. There are some known explicit embeddings of ℓ_2 into ℓ_1 , but they either have high distortion such as $\sqrt{3}$ in [Ber97] or forced to embed into a space with superpolynomial dimension, e.g., [Ind00b],[LLR95].

1.3 Our techniques

For the proof of our main theorem we use two embedding techniques. The first is the one by Figiel et al. [FLM77], and is used for all reductions to the ℓ_p norm for $p < \infty$. For ℓ_∞ , we show an embedding of ℓ_2 into ℓ_∞ that does not contract *any* point in \mathbb{R}^n and that for any point in \mathbb{R}^n there is a polynomially small probability of expansion. This embedding is based on the work of Indyk [Ind03] and could be useful in other applications. Almost all our reductions to the ℓ_∞ norm use this embedding as is. The only exception is our reduction in the case of CVPP, where we need to find some sort of embedding that is good for all $x \in \mathbb{R}^n$ *simultaneously*. For this purpose, we show how to choose a set of embeddings into ℓ_∞ with the following property. For any point $x \in \mathbb{R}^n$, there is an embedding function in the set that does not expand it by “too much” and no embedding function in the set contracts it by “too much”.

Chapter 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 2.0.1 (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.

Definition 2.0.2 (θ -net) We define a θ -net, K_0 on the $(n - 1)$ -dimensional sphere, $S^{n-1} = \{x, \|x\|_2 = 1\}$, to be a set of points $\{y_i\}_{i=1}^N \in S^{n-1}$ such that for every $x \in \mathbb{R}^n$ there is a point $y_0 \in K_0$ for which $\text{dist}_2(x, y_0) \leq \theta$.

Definition 2.0.3 (c -Dimension Blowup) An embedding from $(X, \|\cdot\|_p)$ where X is n -dimensional Banach space into $(Y, \|\cdot\|_q)$ where Y is m -dimensional Banach space and $m(n)$, we say the embedding has m -dimension blowup. For example, if m is linear/exponential in n , we say it has a linear/exponential - dimension blowup.

We show hardness for following problems:

Definition 2.0.4 (SVP $_\gamma$ Approximation Problem) For any $\gamma \geq 1$ the SVP $_\gamma$ problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is to approximate shortest nonzero vector in the lattice up to a factor of γ , namely find $0 \neq v \in \mathcal{L}$ such that $\|v\|_p \leq \gamma \cdot \lambda_1^p(\mathcal{L})$.

Definition 2.0.5 (GapSVP Problem) For $\gamma \geq 1$ the GapSVP problem is a decision problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and $d > 0$ is to determine between two instances. It is a Yes Instance if $\lambda_1^p(\mathcal{L}) \leq d$, it is a No Instance $\lambda_1^p(\mathcal{L}) \geq \gamma \cdot d$

Definition 2.0.6 (CVP $_\gamma$ Approximation version) For any $\gamma \geq 1$ the CVP $_\gamma$ problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and a target vector $t \in \mathbb{R}^n$, is to approximate the closest vector up to γ , namely find $v \in \mathcal{L}$ such that $\|v - t\|_p \leq \gamma \cdot \text{dist}_p(t, \mathcal{L})$.

Definition 2.0.7 (GapCVP) For $\gamma \geq 1$ the GapCVP is a decision problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$, a vector target $t \in \mathbb{R}^n$ and a parameter $d > 0$, is to determine between two instances. It is a Yes Instance if $\exists v \in \mathcal{L}, \|v - t\|_p \leq d$, it is a No Instance $\forall v \in \mathcal{L}, \|v - t\|_p \geq \gamma d$.

Definition 2.0.8 (CVPP $_\gamma$ Approximation version) For $\gamma \geq 1$, we say that (\mathbf{P}, \mathbf{D}) solves CVPP $_\gamma$ if the following is satisfied. The input to the function \mathbf{P} (pre-processing) is the lattice \mathcal{L} (given as some basis \mathbf{B} of \mathcal{L}). The function \mathbf{P} generates a representation of the lattice such that $|\mathbf{P}(\mathcal{L})|$ is polynomial in the size of the dimension of \mathcal{L} . There are no computational requirements from \mathbf{P} . The function \mathbf{D} (a decoding function) given the new description $\mathbf{P}(\mathcal{L})$ and a target vector t , outputs a vector $v \in \mathcal{L}$ such that $\|v - t\|_p \leq \gamma \cdot \text{dist}_p(t, \mathcal{L})$.

Definition 2.0.9 (GapCVPP) For $\gamma \geq 1$, we say that (\mathbf{P}, \mathbf{D}) as above solves GapCVPP $_\gamma$ if given a target vector $t \in \mathbb{R}^n$ and $d > 0$ it can distinguish between the two instances: Yes Instance: where $\text{dist}_p(t, \mathcal{L}) \leq d$ and No Instance where $\text{dist}_p(t, \mathcal{L}) \geq \gamma \cdot d$

We now view some versions of the Shortest Independent Vectors Problem (SIVP). The SIVP problem on $\mathcal{L} \subseteq \mathbb{R}^n$ is to find a set of linearly independent vectors $\{v_i\}_1^n, v_i \in \mathcal{L}$ such that $\max_i \|v_i\| = \lambda_n(\mathcal{L})$.

Definition 2.0.10 (SIVP Search Problem) The SIVP problem on $\mathcal{L} \subseteq \mathbb{R}^n$ is to find a set of linearly independent vectors $\{v_i\}_{i=1}^n, v_i \in \mathcal{L}$ such that $\max_i \|v_i\|$ is minimized.

Definition 2.0.11 (SIVP $_\gamma$ Approximation version) For $\gamma \geq 1$ the SIVP $_\gamma$ problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is to find a set of linearly independent vectors $\{u_i\}_{i=1}^n, v_i \in \mathcal{L}$ such that $\max_i \|u_i\|$ approximates $\lambda_n(\mathcal{L})$ up to γ i.e., $\max_i \|u_i\|_p \leq \gamma \cdot \lambda_n(\mathcal{L})$.

Definition 2.0.12 (GapSIVP) For $\gamma \geq 1$, the GapSIVP problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and a parameter $d > 0$ is to distinguish between the two instances: Yes Instance, where $\lambda_n(\mathcal{L}) \leq d$ and No Instance, where $\lambda_n(\mathcal{L}) \geq \gamma d$.

We now view some versions of the Shortest Basis Problem (SBP). The (SBP) problem on $\mathcal{L} \subseteq \mathbb{R}^n$ is to find a basis of \mathcal{L} , $\{v_i\}_1^n, v_i \in \mathcal{L}$ such that $\max_i \|v_i\|$ is minimized.

Definition 2.0.13 (SBP $_\gamma$ Approximation version) Let $\{v_i\}_1^n$ be a basis of \mathcal{L} such that $\max_i \|v_i\|$ is the smallest possible. The SBP $_\gamma$ problem on an n -rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is to find a basis of \mathcal{L} , $\{u_i\}_{i=1}^n$, such that $\max_i \|u_i\|$ approximates the maximum length of $\max_i \|v_i\|$ up to γ .

Definition 2.0.14 (GapSBP) Let $\{v_i\}_1^n$ be a basis of \mathcal{L} such that $\max_i \|v_i\|$ is the smallest possible and let r denote $\max_i \|v_i\|$. The GapSBP problem on $\mathcal{L} \subseteq \mathbb{R}^n$ and a parameter d is for any $\gamma > 1$, to distinguish between the two instances: Yes Instance, where $r \leq d$ and No Instance, where $r \geq \gamma \cdot d$.

Random Reduction

By a randomized Karp reduction r from A to B , we mean a two-sided error reduction such that

- For every $x \in A$, $\Pr[r(x) \in B] \geq 2/3$
- For every $x \notin A$, $\Pr[r(x) \in B] \leq 1/3$

By a *randomized Cook reduction* r from A to B , we mean that $r(\cdot)$ is a randomized polynomial time Turing machine that solves A by using an oracle to B . The solution of the reduction satisfies $\Pr[r(x) \text{ is the correct solution to } A] \geq 2/3$.

In both reductions the probability is over the random coins used by the reduction Turing machine.

Chapter 3

Embedding

In this chapter we introduce families of embedding functions that will be used later to prove the main theorem. Our reductions use low distortion embeddings of ℓ_2^n into ℓ_p^m for $1 \leq p \leq \infty$. For $p < \infty$ we use the results of [FLM77] which are discussed in Chapter 3.1. Then in Chapter 3.2 we discuss the result of Indyk [Ind03] for embedding into ℓ_∞^m we use and generalize it. Namely we show that with high probability there is an embedding of ℓ_2^n into ℓ_∞^m that does not contract **any** point and that for every point there is a polynomially small probability of expansion. This embedding might be useful in other applications. The above embedding results will be enough for the reductions of SVP, CVP, SIVP, SBP. For the reduction of CVPP we need to combine several embeddings and this is shown in Chapter 3.3.

3.1 Embedding in the ℓ_p norm for $p < \infty$

Definition 3.1.1 (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 random¹ 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.

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

Theorem 3.1.2 [FLM77] There exists $\nu_p(n, m)$ such that for all $p < \infty$, $\epsilon > 0$ and for all n there is m such that with probability at least $1 - 2^{-\Omega(n)}$, a randomly chosen embedding function $f \sim \mathcal{F}(p, n, m)$ satisfy that for all $x \in \mathbb{R}^n$

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p \leq (1 + \epsilon) \|x\|_2. \quad (3.1)$$

For a given dimension n and distortion parameter $\epsilon > 0$, the dimension m should be chosen to be

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

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

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 independent vectors from S^{m-1} . With probability 1, the vectors are linearly independent and we can therefore apply the Gram-Schmidt orthogonalization process to them. 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 . We will approximate ν_p up to ϵ by randomly choosing a vector in \mathbb{R}^n and calculating $Ax/\|x\|_2$. Since $f(x) = \nu_p \cdot Ax$ and since with probability at least $1 - 2^{-\Omega(n)}$ Equation (3.1) holds then by randomly choosing a vector $x \in \mathbb{R}^n$ we get that with probability at least $1 - 2^{-\Omega(n)}$, ν_p is $Ax/\|x\|_2$ up to $1 \pm \epsilon$.

3.2 Embedding in the ℓ_∞ norm

Definition 3.2.1 (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 $\nu_\infty = \frac{1}{\sqrt{2 \ln 2mn}}$ is a normalization factor.

Definition 3.2.2 (Probability of expansion and contraction) For a family \mathcal{F} of embedding functions from ℓ_p^n to ℓ_q^m and a vector $v \in \ell_p^n$, we define the probability of expansion by more than $(1 + \epsilon)$ as

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

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

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

Theorem 3.2.3 [Ind03] For any positive reals $\epsilon, \delta, \eta > 0$ and any large enough integer n , 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 7.1.1.

Notice that the embedding is asymmetric in sense 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.

Notice also that, unlike Theorem 3.1.2, 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. 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 3.2.8 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 3.2.7).

Theorem 3.2.4 For any $\epsilon, \delta > 0$, and any large enough n , 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 $1 - 2^{-\Omega(n)}$ 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: We start with two standard claims.

Claim 3.2.5 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 3.2.6 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 3.2.7 Let f be chosen randomly from $\mathcal{F}(\infty, n, m)$ for some $m \leq \text{poly}(n)$. Then, with probability $1 - 2^{-\Omega(n)}$, for any $x \in \mathbb{R}^n$,

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

Proof: We have

$$\|f(x)\|_\infty = \frac{1}{\sqrt{2 \ln 2mn}} \max_i |z_i x| \leq \max_i \|z_i\|_2 \|x\|_2.$$

By a union bound and Claim 3.2.6,

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

This probability is of the form $2^{-\Omega(n)}$ and the lemma follows. ■

Lemma 3.2.8 *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 $\theta = \frac{1}{n}$, $\eta = 2^{-10n \log n}$, $\delta = n^{-1}$, and $m = (\delta^{-1} + n \log n + 1/\epsilon)^{O(1/\epsilon)}$ as in Theorem 3.2.4. The bound on the expansion probability of $\mathcal{F}(\infty, n, m)$ follows immediately from Theorem 3.2.4. In the rest of the proof we prove the second claim.

Let K_0 be some θ -net of cardinality $|K_0| \leq (1 + 2/\theta)^n$. Such a net exists by Claim 3.2.5. By Theorem 3.2.3, 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 3.2.7, we know that with probability $1 - 2^{-\Omega(n)}$, 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 3.2.8 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 is the preprocessing step. 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 choose $k = O(n)$ 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 3.3.1 *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 $x \in \mathbb{R}^n$,*

$$\exists i, \|f_i(x)\|_\infty \leq (1 + \epsilon) \|x\|_2 \quad \text{and} \quad (3.2)$$

$$\forall j, \|f_j(x)\|_\infty \geq (1 - \epsilon) \|x\|_2. \quad (3.3)$$

Proof: Choose $m = (n \log n)^{O(1/\epsilon)}$ as given by Theorem 3.2.4 with ϵ and $c = 1$. For any fixed j , Theorem 3.2.4 says that the probability that there exists an x such that $\|f_j(x)\|_\infty < (1 - \epsilon) \|x\|_2$ is at most $2^{-\Omega(n)}$. By the union bound, we obtain that (3.3) holds with probability $1 - k2^{-\Omega(n)} = 1 - 2^{-\Omega(n)}$.

In the remainder of the proof we prove (3.2). 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 3.2.5. Fix some $y \in K_0$. By Theorem 3.2.4, 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 $1 - 2^{-\Omega(n)}$, (3.2) 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 3.2.7, we have that with probability $1 - 2^{-\Omega(n)}$ no f_i expands any point by more than $2\sqrt{n}$.

Lemma 3.3.2 *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 ,

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

■

■

Chapter 4

Applications - reductions among lattice problems

4.1 SVP reductions

Theorem 4.1.1 *For all $\epsilon > 0$ and for all $1 \leq p \leq \infty$, there is a randomized Cook reduction from SVP_γ in the ℓ_2 norm to $\text{SVP}_{\gamma'}$ in the ℓ_p norm where $\gamma' = (1 - \epsilon)\gamma$.*

We divide the reduction into two cases, the first case is from SVP in the ℓ_2 norm to SVP_γ in the ℓ_p norm for $1 \leq p < \infty$ and the second to $\text{SVP}_{\gamma'}$ in the ℓ_∞ norm.

Lemma 4.1.2 (*Approximation SVP_γ*) *For any $\epsilon > 0$, $p < \infty$, and $\gamma' \geq 1$, there is a randomized Cook reduction from SVP_γ in the ℓ_2 norm to $\text{SVP}_{\gamma'}$ in the ℓ_p norm where $\gamma = (1 + \epsilon)\gamma'$.*

Proof: The input of the reduction is a lattice \mathcal{L} . It first chooses $f \sim \mathcal{F}(p, n, m)$ where m is as in Theorem 3.1.2 and then invokes the oracle of $\text{SVP}_{\gamma'}$ in the ℓ_p norm 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. From Theorem 3.1.2, it follows that with high probability over the choice of f , for every $x \in \mathbb{R}^n$,

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p \leq (1 + \epsilon) \|x\|_2. \quad (4.1)$$

So in the rest of the proof, assume that f satisfies (4.1). Let $w \in \mathcal{L}$ be a vector that achieves $\|w\|_2 = \lambda_1^2(\mathcal{L})$, and let $s' \in f(\mathcal{L})$ be the output of the oracle. By assumption, s' satisfies $\|s'\|_p \leq \gamma' \cdot \lambda_1^p(f(\mathcal{L}))$. Using (4.1) twice, we obtain

$$\|f^{-1}(s')\|_2 \leq \frac{1}{1 - \epsilon} \|s'\|_p \leq \frac{1}{1 - \epsilon} \cdot \gamma' \cdot \lambda_1^p(f(\mathcal{L})) \leq \frac{1}{1 - \epsilon} \cdot \gamma' \cdot \|f(w)\|_p \leq \frac{1 + \epsilon}{1 - \epsilon} \cdot \gamma' \|w\|_2 = \frac{1 + \epsilon}{1 - \epsilon} \cdot \gamma' \cdot \lambda_1^2(\mathcal{L}).$$

Therefore, the reduction outputs a $\frac{1+\epsilon}{1-\epsilon} \cdot \gamma'$ approximation to the shortest vector in \mathcal{L} , and since ϵ is arbitrary, this completes the proof. \blacksquare

Lemma 4.1.3 *For any $\epsilon > 0$ there is a randomized Cook reduction from SVP_γ in the ℓ_2 norm to $\text{SVP}_{\gamma'}$ in the ℓ_∞ norm, for $\gamma' = 1 - \epsilon$.*

Proof: The input to the reduction is a lattice, \mathcal{L} . The reduction randomly chooses $f \sim \mathcal{F}(\infty, n, m)$ and embeds ℓ_2^n into ℓ_∞^m , where $m = (n \log n)^{O(1/\epsilon)}$ by setting $\delta = 1/n$ in Theorem 3.2.4. The reduction then invokes the oracle of SVP in the ℓ_∞ norm with $f(\mathcal{L})$ to obtain some $s' \in f(\mathcal{L})$. The output of the reduction is the lattice vector $f^{-1}(s') \in \mathcal{L}$.

We now prove the correctness of the reduction. By Theorem 3.2.4, with probability $1 - 2^{-\Omega(n)}$, f does not contract any $x \in \mathbb{R}^n$. In the rest of the proof we assume that this is the case. We now show that $\|f^{-1}(s')\|_2 \leq \gamma \cdot \lambda_1^2(\mathcal{L})$.

Let s' be the vector returned by the oracle of $\text{SVP}_{\gamma'}$ in the ℓ_∞ norm. s' at most γ' times the shortest vector in $f(\mathcal{L})$ under the ℓ_∞ norm, i.e., $\|s'\|_\infty \leq \gamma' \cdot \lambda_1^\infty(f(\mathcal{L}))$. Then since $f^{-1}(s') \in \mathbb{R}^n$ then in particular f does not contract $f^{-1}(s')$. We obtain that:

$$\|f^{-1}(s')\|_2 \leq \frac{1}{1-\epsilon} \|s'\|_\infty \leq \frac{1}{1-\epsilon} \gamma' \cdot \lambda_1^\infty(f(\mathcal{L}))$$

In particular for some $w \in \mathcal{L}$ such that $\|w\|_2 = \lambda_1^2(\mathcal{L})$

$$\frac{1}{1-\epsilon} \gamma' \cdot \lambda_1^\infty(f(\mathcal{L})) \leq \frac{1}{1-\epsilon} \gamma' \|f(w)\|_\infty \leq \frac{1+\epsilon}{1-\epsilon} \gamma' \|w\|_2 \stackrel{1}{=} (1+\epsilon) \gamma \cdot \lambda_1^2(\mathcal{L})$$

$$(1) \gamma = \frac{1+\epsilon}{1-\epsilon} \gamma' \quad \blacksquare$$

Theorem 4.1.4 *For any ϵ there is a randomized Karp reduction from GapSVP_γ in the ℓ_2 norm to $\text{GapSVP}_{\gamma'}$ in the ℓ_p norm for $1 \leq p \leq \infty$, for $\gamma' = (1-\epsilon)\gamma$.*

We omit the reduction from GapSVP_γ in the ℓ_2 norm to $\text{GapSVP}_{\gamma'}$ for finite p , which is very similar to the reduction of SVP_γ in the approximation version and show the reduction from GapCVP_γ in the ℓ_2 norm to $\text{GapCVP}_{\gamma'}$ in the ℓ_∞ norm.

Lemma 4.1.5 *For any $\epsilon > 0$, $\gamma \geq 1$, there is a Karp randomized reduction from GapSVP_γ in the ℓ_2 norm to $\text{GapSVP}_{\gamma'}$ in the ℓ_∞ norm for $\gamma' = (1-\epsilon)\gamma$.*

Proof: The input of the reduction is a lattice \mathcal{L} and a parameter d . The reduction randomly chooses $f \sim \mathcal{F}(\infty, n, m)$ where $m = (n \log n)^{O(1/\epsilon)}$ by setting $\delta = 1/n$ in Theorem 3.2.4. The output of the reduction is the embedded lattice, $f(\mathcal{L})$ and a parameter $d' = (1+\epsilon)d$.

We now prove the correctness of the reduction. First we consider a YES instance where there exists $u \in \mathcal{L}$ such that $\|u\|_2 \leq d$. With high probability ($\geq 1 - 1/\text{poly}(n)$), it follows from Theorem 3.2.4 that $\|f(u)\|_\infty \leq (1+\epsilon)\|u\|_2 \leq (1+\epsilon)d = d'$. Since $f(u) \in f(\mathcal{L})$, $\lambda_1^\infty(f(\mathcal{L})) \leq (1+\epsilon)d = d'$ and hence the output is a YES instance. Namely, with probability $1 - n^{-O(1)}$ over a choice of f , there is $v \in f(\mathcal{L})$, such that $\|v\|_\infty \leq (1+\epsilon)d$. This is a YES instance for SVP in the ℓ_∞ norm.

We now consider a NO instance where $\lambda_1^2(\mathcal{L}) \geq \gamma d$. By Theorem 3.2.4, with probability $1 - 2^{-\Omega(n)}$ over the choice of the $f \sim \mathcal{F}(\infty, n, m)$, 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 $\forall u \in \mathcal{L}$, $\|u\|_2 \geq \gamma d$ then with probability $\geq 1 - 2^{-\Omega(n)}$,

$$\forall v \in f(\mathcal{L}), v \neq 0 \quad \|v\|_\infty \geq (1-\epsilon) \|f^{-1}(v)\|_2 \geq (1-\epsilon)\gamma d$$

where we use that $f^{-1} \in \mathcal{L}$. Since all $v \in f(\mathcal{L})$ have a source $u = f^{-1}(v)$ and by our assumption that any $v \in f(\mathcal{L})$, $(1-\epsilon)\gamma d \leq \|v\|_\infty$, we obtain that:

$$\|v\|_\infty \leq \gamma' \cdot (1+\epsilon)d$$

We have obtained a NO instance for SVP in the ℓ_∞ norm. \blacksquare

Combining Khot's [Kho04] result and Theorem 4.1.4 we obtain the following corollaries:

Corollary 4.1.6 *SVP in the ℓ_1 norm is NP-hard to approximate within a constant for every constant unless $\text{NP} \subseteq \text{BPP}$.*

Corollary 4.1.7 *SVP in the ℓ_1 norm is NP-hard to approximate within $2^{(\log n)^{1/2-\epsilon}}$ for every ϵ unless $\text{NP} \subseteq \text{BPPTIME}(2^{\text{poly}(\log n)})$.*

In Chapter 5.1 we review a result showing a deterministic embedding of ℓ_2^m 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 4.1.8 *SVP in the ℓ_1 norm is NP-hard to approximate within any constant unless $\text{NP} \subseteq \text{RP}$.*

Corollary 4.1.9 *SVP in the ℓ_1 norm is NP-hard to approximate 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 4.2.1 *For any $\epsilon > 0$, $\gamma > 1$ and for $1 \leq p \leq \infty$, there is a randomized Cook reduction from CVP_γ in the ℓ_2 norm to $\text{CVP}_{\gamma'}$ in the ℓ_p norm where $\gamma' = (1 - \epsilon)\gamma$.*

We divide the reduction into two cases, the first case is from CVP in the ℓ_2 norm to CVP_γ in the ℓ_p norm for $1 \leq p < \infty$ and the second is from CVP_γ in the ℓ_2 norm to CVP'_γ in the ℓ_∞ norm.

Lemma 4.2.2 *For any $\epsilon > 0$, $\gamma > 1$ and $1 \leq p < \infty$ there is a randomized Cook reduction from CVP_γ in the ℓ_2 norm to $\text{CVP}_{\gamma'}$ in ℓ_p where $\gamma' = (1 - \epsilon)\gamma$.*

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^m in ℓ_p^m where m is given in Theorem 3.1.2. The reduction then invokes the oracle of CVP in the ℓ_p norm 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 $\|w - t\|_2 = \text{dist}_2(t, \mathcal{L})$. First notice that the vector $v' \in f(\mathcal{L})$ is such that,

$$\|v' - f(t)\|_p \leq \gamma' \cdot \text{dist}_p(f(t), f(\mathcal{L})). \quad (4.2)$$

More over by Theorem 3.1.2, it follows that with high probability over the choice of f , for any $x \in \mathbb{R}^n$,

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p \leq (1 + \epsilon) \|x\|_2. \quad (4.3)$$

By (4.3) and (4.2) we obtain that:

$$\begin{aligned} \|f^{-1}(v') - t\|_2 &\leq \frac{1}{1 - \epsilon} \cdot \|v' - f(t)\|_p && \text{by (4.3)} \\ &\leq \frac{1}{1 - \epsilon} \cdot \gamma' \cdot \text{dist}_p(f(t), f(\mathcal{L})) && \text{by (4.2)} \\ &\leq \frac{1}{1 - \epsilon} \cdot \gamma' \cdot \|f(w - t)\|_p && f(w) \in \mathcal{L} \\ &\leq \frac{1 + \epsilon}{1 - \epsilon} \cdot \gamma' \cdot \|w - t\|_2 = \gamma \cdot \text{dist}_2(t, \mathcal{L}) && \text{by (4.3)} \end{aligned}$$

where we define $\gamma = \frac{1-\epsilon}{1+\epsilon}\gamma'$. ■

Lemma 4.2.3 *For any $\epsilon > 0$, $\gamma > 1$ there is a randomized Cook reduction from CVP_γ in the ℓ_2 norm to $\text{CVP}_{\gamma'}$ in the ℓ_∞ norm where $\gamma = (1 - \epsilon)\gamma'$*

Proof: The input of the reduction is an n -dimensional lattice \mathcal{L} and a target point t . The reduction randomly chooses $f \sim \mathcal{F}(\infty, n, m)$ that is used to embed ℓ_2^n into ℓ_∞^m where $m = (n \log n)^{O(1/\epsilon)}$ is determined by setting $\delta = 1/n$ in Theorem 3.2.4. The reduction then invokes the oracle of $\text{CVP}_{\gamma'}$ in the ℓ_∞ norm with $f(\mathcal{L})$ and $f(t)$ to obtain $v' \in f(\mathcal{L})$. 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 a vector achieving the minimum distance $\text{dist}_2(t, \mathcal{L})$. By Theorem 3.2.4 with probability $1 - 2^{-\Omega(n)}$ over $f \sim \mathcal{F}(\infty, n, m)$, for any $x \in \mathbb{R}^n$,

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_\infty. \quad (4.4)$$

In addition, for the vector $w - t$, with probability $1 - 1/\text{poly}(n)$

$$\|f(w - t)\|_\infty \leq (1 + \epsilon) \|w - t\|_2. \quad (4.5)$$

We now show that $\|f^{-1}(v') - t\|_2 \leq \gamma \cdot \text{dist}_2(t, \mathcal{L})$. Let v' be the vector returned by the oracle, i.e., $\|v' - f(t)\|_\infty \leq \gamma' \cdot \text{dist}_\infty(f(t), f(\mathcal{L}))$. Then for the vector $(f^{-1}(v') - t) \in \mathbb{R}^n$ with exponentially high probability, f does not contract $(f^{-1}(v') - t)$. Thus,

$$\begin{aligned} \|f^{-1}(v') - t\|_2 &\stackrel{(4.4)}{\leq} \frac{1}{1 - \epsilon} \|v' - f(t)\|_\infty \stackrel{\text{assumption on } v'}{\leq} \frac{1}{1 - \epsilon} \gamma' \cdot \text{dist}_\infty(f(t), f(\mathcal{L})) \leq \\ &\leq \frac{1}{1 - \epsilon} \gamma' \|f(w) - f(t)\|_\infty \stackrel{(4.5)}{\leq} \frac{1 + \epsilon}{1 - \epsilon} \gamma' \|w - t\|_2 \stackrel{\gamma = \frac{1+\epsilon}{1-\epsilon}\gamma'}{=} \gamma \cdot \text{dist}_2(t, \mathcal{L}) \end{aligned}$$
■

Theorem 4.2.4 *For any $\epsilon > 0$, $\gamma > 1$ there is a randomized Karp reduction from GapCVP_γ in the ℓ_2 norm to $\text{GapCVP}_{\gamma'}$ in ℓ_p for $1 \leq p \leq \infty$, for $\gamma' = (1 + \epsilon)\gamma$.*

We omit the reduction from GapCVP_γ in the ℓ_2 norm to $\text{GapCVP}_{\gamma'}$ for $1 \leq p < \infty$ which is very similar to that of Lemma 4.2.2 and we show the reduction from GapCVP_γ in the ℓ_2 norm to $\text{GapCVP}_{\gamma'}$ in the ℓ_∞ norm.

Lemma 4.2.5 *For any $\epsilon > 0$, $\gamma \geq 1$, there is a randomized Karp reduction from GapCVP_γ in the ℓ_2 norm to $\text{GapCVP}_{\gamma'}$ in the ℓ_∞ norm.*

Proof: The input to the reduction is a lattice \mathcal{L} , a target point t and a parameter d . The reduction randomly chooses $f \sim \mathcal{F}(\infty, n, m)$ where $m = (n \log n)^{O(1/\epsilon)}$ by setting $\delta = 1/n$ in Theorem 3.2.4. The output of the reduction is the embedded lattice, $f(\mathcal{L})$, the embedded target vector, $f(t)$ and a parameter $d' = (1 + \epsilon)d$.

We now prove correctness of the reduction. First we consider a YES instance where there exists $u \in \mathcal{L}$ such that $\|u - t\|_2 \leq d$. With probability at least $(1 - 1/\text{poly}(n))$, it follows from Theorem

3.2.4 that $\|f(u-t)\|_\infty \leq (1+\epsilon)\|u-t\|_2 \leq (1+\epsilon)d = d'$. In particular this implies that this is a YES instance for CVP in the ℓ_∞ norm.

We now consider a NO instance in which $\text{dist}(t, \mathcal{L}) \geq \gamma d$. By Theorem 3.2.4, we conclude that with probability $\geq 1 - 2^{-\Omega(n)}$, for all $v \in f(\mathcal{L})$, $\|v - f(t)\|_\infty \geq (1-\epsilon)\|f^{-1}(v) - t\|_2 \geq (1-\epsilon)\gamma d = \gamma' d'$ and since all $v \in f(\mathcal{L})$ have a source $u = f^{-1}(v)$ then we obtained a NO instance for CVP in the ℓ_∞ norm and the correctness of the reduction follows. ■

4.3 CVPP reductions

Theorem 4.3.1 *For any $\epsilon > 0, \gamma > 1$ and $1 \leq p \leq \infty$, there is a deterministic Karp reduction from GapCVPP_γ in the ℓ_2 norm with parameter γ to $\text{GapCVPP}_{\gamma'}$ in the ℓ_p norm with parameter $\gamma' = (1-\epsilon)\gamma$.*

We first prove the reduction to $\text{GapCVPP}_{\gamma'}$ in ℓ_p norm for $p < \infty$ and then for ℓ_∞ norm.

Lemma 4.3.2 *For any $\epsilon > 0, \gamma > 1$ and $1 \leq p < \infty$, there is a gap preserving reduction from GapCVPP_γ in the ℓ_2 to $\text{GapCVPP}_{\gamma'}$ in the ℓ_p where $\gamma' = (1-\epsilon)\gamma$.*

Proof: Denote by D_p, P_p the decoding function and the pre-processing function of CVPP in the ℓ_p norm. Let m be the value given by Theorem 3.1.2 for n, ϵ . The reduction construct (D_2, P_2) using (D_p, P_p) we define P_2, D_2 that solve GapCVPP_γ in the ℓ_2 as follows:

P_2 - Pre-processing

1. Input: A lattice $\mathcal{L}(B)$ in the form of a basis B .
2. Find an embedding f in $\mathcal{F}(p, n, m)$ for which Theorems 3.1.2 holds. (recall that there are no time limitations on the pre-processing function).
3. Invoke $P_p(f(\mathcal{L}))$ to obtain a new description of $f(\mathcal{L})$ - the embedded lattice in ℓ_p^m .
4. Output $P_p(f(\mathcal{L}))$ and $f(\cdot)$.

D_2 - Decoding

1. Input $t, d, P_p(f(\mathcal{L})), f(\cdot)$
2. Compute $f(t) \in \mathbb{R}^m$ where t is the target vector and $f(\cdot)$ is the embedding function used in the pre-processing step.
3. Return $D_p(f(t), d', P_p(f(\mathcal{L})), f(\cdot))$. Where $d' = (1+\epsilon)d$ and $D_p(\cdot, \cdot, \cdot, \cdot)$ returns TRUE or FALSE. If the distance of $f(t)$ from the lattice $f(\mathcal{L})$ is less than d' it returns TRUE and if it is more than $\gamma' d'$ it returns FALSE, where $\gamma' = \frac{1+\epsilon}{1-\epsilon}$.

The proof of the correctness is very similar to the proof of Lemma 4.2.2 and therefore is omitted. ■

Lemma 4.3.3 *For any $\epsilon > 0, \gamma > 1$, there is a deterministic Cook reduction from GapCVPP_γ in the ℓ_2 norm to $\text{GapCVPP}_{\gamma'}$ in the ℓ_∞ norm where $\gamma' = (1-\epsilon)\gamma$.*

Proof: The construction of (D_2, P_2) using (D_∞, P_∞) is done in the following way:

P_2 - Pre-processing

1. Let f_1, f_2, \dots, f_k (k will be set later) be a sequence of embedding functions for which Theorem 3.3.1 holds. Since there are no time limitations in the pre-processing stage, such a set can be found.
2. Invoke $P_\infty(\cdot)$ on $f_i(\mathcal{L})$ to obtain k new descriptions of the embedded lattices in ℓ_∞^m .
3. Output: $\{P_\infty(f_i(\mathcal{L}))\}_{i=1}^k, \{f_i(\cdot)\}_{i=1}^k$

 D_2 - Decoding

1. Input $t, d, \{P_\infty(f_i(\mathcal{L}))\}_{i=1}^k, \{f_i(\cdot)\}_{i=1}^k$
2. For every $i = 1, \dots, k$ invoke $D_\infty(f_i(t), d', P_\infty(f_i(\mathcal{L})), f_i(\cdot))$ where $d' = (1+\epsilon)d$ for solving $\text{GapCVPP}_{\gamma'}$ in the ℓ_∞ norm.
3. If one of the calls $D_\infty(f_i(t), d', P_\infty(f_i(\mathcal{L})), f_i(\cdot))$ returns TRUE then return TRUE, else return FALSE.

Set $\gamma' = \frac{1-\epsilon}{1+\epsilon}\gamma$. According to Theorem 3.3.1, f_1, f_2, \dots, f_k are such that

$$\forall x \in \mathbb{R}^n \left[\exists i, \|f_i(x)\|_\infty \leq (1+\epsilon)\|x\|_2 \quad \bigwedge \quad \forall j, \|f_j(x)\|_\infty \geq (1-\epsilon)\|x\|_2 \right]. \quad (4.6)$$

We now prove correctness. First assume we are given a YES instance, that is, there is $u \in \mathcal{L}$ such that $\|u - t\|_2 \leq d$, then from (4.6) it follows that there is i such that 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 calls to $D_\infty(f_i(t), d', P_\infty(f_i(\mathcal{L})), f_i(\cdot))$ will return True and therefore this is a Yes instance for $f_i(\mathcal{L})$. We now assume a NO instance, that is, for all $u \in \mathcal{L}, \|u - t\|_2 \geq \gamma \cdot d$. By the choice of $\{f_j\}_{j=1}^k$ we have that for any $x \in \mathbb{R}^n$, no $f_j(\cdot)$ contracts x in more than $(1-\epsilon)$. It then follows that

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

For the choice of $\gamma' = \frac{1-\epsilon}{1+\epsilon}$ and $d' = (1+\epsilon)d$, we obtain a NO instance for every $f_j(\mathcal{L})$. ■

Theorem 4.3.4 *For any $\epsilon > 0$, $\gamma > 1$ and $1 \leq p \leq \infty$, there is a Cook reduction from CVPP_γ in the ℓ_2 norm with parameter γ to $\text{CVPP}_{\gamma'}$ in the ℓ_p norm with parameter $\gamma' > 1$ for $\gamma' = (1-\epsilon)\gamma$.*

The proof is very similar to the proof of Lemma 4.3.2 and Lemma 4.3.3 and so we only state the differences between the two. We divide the proof to two lemmas one for the reduction into ℓ_p^m for finite p and the other for $p = \infty$:

Lemma 4.3.5 *For any $\epsilon > 0$, $\gamma > 1$ there is a Cook reduction from CVPP_γ in the ℓ_2 norm to $\text{CVPP}_{\gamma'}$ in the ℓ_p norm for $1 \leq p < \infty$ (finite p) where $\gamma' = (1-\epsilon)\gamma$.*

Proof: The reduction uses the same pre-processing as in Lemma 4.3.2 and only differs on the output of $D_p(\mathcal{L}, t)$. Recall that in the approximation version, $D_p(\mathcal{L}, t)$ outputs a vector which is guaranteed to be an approximation of the shortest vector up to γ . Thus D_2 returns $f^{-1}(D_p(P_p(f(\mathcal{L})), f(t)))$. The correctness proof is similar to that in Lemma 4.2.2 and is omitted. ■

Lemma 4.3.6 *For any $\epsilon > 0$, $\gamma > 1$ there is a Cook reduction from CVPP_γ in the ℓ_2 norm to $\text{CVPP}_{\gamma'}$ in the ℓ_∞ norm where $\gamma' = (1 - \epsilon)\gamma$.*

Proof: The reduction uses the same pre-processing step as in Lemma 4.3.3 and only differs on the output of $D_p(\cdot)$. Let $u_i = f_i^{-1}(D_\infty(P_\infty(f_i(\mathcal{L})), f_i(t)))$. The reduction invokes all $(D_\infty(P_\infty(f_i(\mathcal{L})), f_i(t)))$ to compute all $\{u_i\}_{i=1}^k$. The output of the reduction is u_i that minimizes $\|t - u_i\|_2$.

We now prove correctness. We need to show that there is an embedding function f_i such that the vector v_i returned by $D_\infty(P_\infty(f_i(\mathcal{L})), f_i(t))$ satisfies the following

$$\|f_i^{-1}(v_i) - t\|_2 \leq \gamma \cdot \text{dist}_2(t, \mathcal{L}). \quad (4.7)$$

Let $w \in \mathcal{L}$ be a vector in \mathbb{R}^n achieving $\text{dist}_2(t, \mathcal{L})$. By Theorem 3.3.1, there is an embedding function f_i that does not expand $(w - t) \in \mathbb{R}^n$ by more than $1 + \epsilon$. Let f_i be such a function. We show that $D_\infty(P_\infty(f_i(\mathcal{L})), f_i(t))$ returns a vector that satisfies (4.7).

Since f_i does not expand $w - t$ and by using the linearity of f_i we obtain

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

In particular,

$$\text{dist}_\infty(f_i(t), f_i(\mathcal{L})) \leq (1 + \epsilon) \text{dist}_2(t, \mathcal{L}) \quad (4.8)$$

By our assumption v_i is such that

$$\|v_i - f_i(t)\|_\infty \leq \gamma' \cdot \text{dist}_\infty(f_i(t), f_i(\mathcal{L})) \quad (4.9)$$

By the choice of $\{f_j(\cdot)\}_{j=1}^k$ (that satisfies Theorem 3.3.1) it is guaranteed that for every $f_j(\cdot)$ and for any $v \in f_j(\mathcal{L})$

$$\|f_j^{-1}(v) - t\|_2 \leq \frac{1}{(1 - \epsilon)} \|v - f_j(t)\|_\infty$$

Therefore in particular for f_i and v_i we obtain

$$\|f_i^{-1}(v_i) - t\|_2 \leq \frac{1}{(1 - \epsilon)} \|v_i - f_i(t)\|_\infty \quad (4.10)$$

Combining (4.8), (4.9) and (4.10) we obtain

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

for the choice of $\gamma' = \frac{1 - \epsilon}{1 + \epsilon} \gamma$, and the correctness of the reduction follows. ■

By combining Theorem 4.3.1 and [AKKV05] we conclude the following result:

Corollary 4.3.7 *The CVPP problem in the ℓ_p norm for $2 \leq p \leq \infty$ is hard to approximate*

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

4.4 SIVP and SBP reductions

We next show that the Shortest Independent Vector Problem SIVP (respectively SBP) under the ℓ_2 norm is not harder than SIVP (respectively SBP) in the ℓ_p norm, for $1 \leq p \leq \infty$. The reduction is shown for SIVP but the same reduction applies for SBP.

Theorem 4.4.1 *For any $\epsilon > 0$, $\gamma > 1$ and $1 \leq p \leq \infty$ there is a randomized Karp reduction from GapSIVP_γ in the ℓ_2 norm to $\text{GapSIVP}_{\gamma'}$ in the ℓ_p norm, where $\gamma' = (1 + \epsilon)\gamma$.*

We divide the proof into two part, the first is a reduction from SIVP in the ℓ_2 norm to SIVP in the ℓ_p , for finite p and the second to SIVP in the ℓ_∞ norm. Given a lattice \mathcal{L} , and a parameter d , the reduction randomly chooses $f \sim \mathcal{F}(p, n, m)$ satisfying Theorem 3.1.2 (for $p < \infty$) or Theorem 3.2.4 (for $p = \infty$ where m is obtained by setting $\delta = 1/n^2$) and then outputs $f(\mathcal{L})$ and a parameter $d' = (1 + \epsilon)d$.

Definition 4.4.2 (*length of a set of vectors - $\text{le}_p(\cdot)$*) *Let the length of a set of vectors $V = \{v_1, v_2, \dots, v_n\}$ be the length of the longest vector in the set measured in some given norm. More formally, $\text{le}_p(V) = \max_i \|v_i\|_p$.*

Claim 4.4.3 *Let $\{v_i\}_{i=1}^n$ be a set of linearly independent vectors in \mathbb{R}^n and let f be an embedding function $f \sim \mathcal{F}(p, n, m)$, then $\{f(v_i)\}_{i=1}^n$ is a set of linearly independent vectors in \mathbb{R}^m .*

Proof: In the case $p < \infty$, the embedding function is a matrix of n orthonormal column vectors from \mathbb{R}^m that are therefore linearly independent. In the $p = \infty$ case, the embedding function is a matrix of n randomly chosen column vectors from \mathbb{R}^m . The probability that they are not in general position is zero and therefore we can assume that the matrix columns are linearly independent too. A set of linearly independent vectors remains linearly independent after multiplication by a non-singular matrix. ■

We now analyze the reduction to GapSIVP in the ℓ_p norm for $1 \leq p < \infty$ (finite p). The reduction randomly chooses $f \sim \mathcal{F}(n, m, p)$ to embed ℓ_2^n into ℓ_p^m where m is give in Theorem 3.1.2. From Theorem 3.1.2, it follows that with probability at least $1 - 2^{-\Omega(n)}$ over the choice of f , for every $x \in \mathbb{R}^n$,

$$(1 - \epsilon) \|x\|_2 \leq \|f(x)\|_p \leq (1 + \epsilon) \|x\|_2. \quad (4.11)$$

We now consider a YES instance where there is a set of linearly independent lattice vectors $U = \{u_i\}_{i=1}^n$, $u_i \in \mathcal{L}$ such that, $\text{le}_p(U) \leq d$. It then follows by from (4.11) that with probability at least $1 - 2^{-\Omega(n)}$ over the choice of the f , all vectors in U satisfy

$$(1 - \epsilon) \|u_i\|_2 \leq \|f(u_i)\|_p \leq (1 + \epsilon) \|u_i\|_2 \leq (1 + \epsilon)d = d'. \quad (4.12)$$

and we obtain a YES instance in $f(\mathcal{L})$. Now consider a NO instance in which every set of n -linearly independent vectors $U = \{u_i\}_{i=1}^n$, $u_i \in \mathcal{L}$ is such that $\text{le}_2(U) \geq \gamma d$. Using (4.12) again we obtain that for the vector u_i achieving $\text{le}_2(U)$ the following holds

$$\|f(u_i)\|_p \geq (1 - \epsilon) \|u_i\|_2 \geq (1 - \epsilon)\gamma d = \gamma' d'$$

therefore correctness of the reduction for $p < \infty$ follows.

We now analyze the reduction in the ℓ_∞ norm. We first show that if there is a set $U = \{u_1, u_2, \dots, u_n\}$ of n linearly independent vectors in \mathcal{L} whose length is $\leq d$ then there is a set of $V = \{v_1, v_2, \dots, v_n\}$ in $f(\mathcal{L})$ of n linearly independent vectors in $f(\mathcal{L})$ whose length is $\leq d(1 + \epsilon)$. By union bound over all points in U it follows that with probability at least $1 - 1/n$, f does not expand any u_i by more than $(1 + \epsilon)$. Thus $\text{le}_\infty(f(U)) \leq (1 + \epsilon)\text{le}_2(U) \leq (1 + \epsilon)d = d'$.

We now show that if every set $U = \{u_1, u_2, \dots, u_n\}$ of n linearly independent vectors in \mathcal{L} has length $\geq \gamma d$ then every set of $V = \{v_1, v_2, \dots, v_n\}$ in $f(\mathcal{L})$ of n linearly independent vectors in $f(\mathcal{L})$ is of length $\geq \gamma' d(1 + \epsilon) = \gamma' d'$ where $\gamma' = \frac{1+\epsilon}{1-\epsilon}\gamma$. Assume for all $U \subseteq \mathcal{L}$, $\text{le}_2(U) \geq \gamma d$. Choosing $f \sim \mathcal{F}(\infty, n, m)$ according to Theorem 3.2.4 we obtain that with probability $1 - 2^{-\Omega(n)}$, f does not contract any $x \in \mathbb{R}^n$ by more than $1 - \epsilon$. It then follows that with high probability

$$\forall U \subseteq \mathcal{L}, \text{le}_\infty(f(U)) \geq (1 - \epsilon)\text{le}_2(U) \geq \gamma \cdot d(1 - \epsilon).$$

Since every $v \in f(\mathcal{L})$ has a source $u = f^{-1}(v)$ we obtain that for all sets of n linearly independent vectors $V = \{v_1, v_2, \dots, v_n\}$ in $f(\mathcal{L})$

$$\text{le}_\infty(V) \geq \gamma' d'$$

and the theorem follows.

Lemma 4.4.4 *For any $\epsilon > 0$, $\gamma > 1$ there is a random Cook reduction from SIVP_γ in the ℓ_2 norm to $\text{GapSIVP}_{\gamma'}$ in the ℓ_p norm, where $\gamma' = (1 + \epsilon)\gamma$.*

The proof is very similar to that of Theorem 4.4.1 and is therefore omitted.

Chapter 5

Extensions

5.1 Deterministic Reductions

All the embedding results we have used such as Theorem 3.1.2 or Theorem 3.2.4 were probabilistic. There are some deterministic embeddings of ℓ_2^n into ℓ_1^m that can be used. A useful deterministic embedding constructed by [Ber97] and [Mul96] gives a constructive embedding where $m = O(n^2)$ and has distortion of $\sqrt{3}$. Indyk [Ind00a] showed a deterministic embedding with parameters $m = n^{O(\log n)}$ and $(1 + \epsilon)$ distortion.

The deterministic embeddings can be used in all the reductions in Chapter 4 to obtain a deterministic reduction with lost of a constant factor using [Ber97]. One can also choose instead of $f \sim (p, n, m)$, the embedding function given in [Ind00a] to obtain a super polynomial deterministic reduction without losing any factor, i.e., for any $\epsilon > 0$.

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

Theorem 5.2.1 *For any $1 \leq q < p \leq 2$, there is a randomized Karp reduction from lattice problem such as GapSVP_γ , GapCVP_γ , GapSIVP_γ in the ℓ_p norm to the corresponding lattice problem in the ℓ_q norm.*

The proof follows by replacing the embedding function in the proof of SVP in the ℓ_p norm for $p < \infty$ by the embedding function given by Johnson and Schechtman [JS82]. We omit the proof of the theorem and state the result of [JS82].

Theorem 5.2.2 *For all $\epsilon > 0$ and $1 \leq q < p \leq 2$ there is a constant $\beta = \beta(p, q, \epsilon) \geq 1$ ℓ_p^n embeds into ℓ_q^m with distortion $(1 + \epsilon)$ whenever $n \leq \beta^{-1}m$.*

5.3 Full rank

The reductions shown in Chapter 4 do not preserve a full rank lattice. Given a full rank lattice, the reduction embeds a lattice of dimension n and rank n into \mathbb{R}^m obtaining a lattice in dimension m and rank $n < m$. To maintain a full-rank lattice after the reduction we can choose the rest of the $m - n$ vectors from the orthogonal space to the n -dimensional space that contains the lattice. For problems such as SVP, CVP, CVPP we need to make sure that the vectors added are long enough (longer than some basis vector of the lattice). For problems such as SIVP, SBP we need to make

sure that the vectors added are short enough so that the longest vector is taken from the original lattice (this can be done via Cook reduction).

5.4 Norm Reduction via Dimension Preserving Random Rotation

In this chapter we present another embedding result with the property that it preserves the dimension. This is a desirable property, especially if one want to use norm reductions in the construction of efficient lattice algorithms. Most known algorithms for lattice problems such as SVP, CVP exponentially depend on the dimension of the lattice (not the rank) therefore it is crucial in those cases that the reduction embeds while maintaining the dimension of the lattice. Clearly, a polynomial blowup in the dimension often makes things impractical. As another example, assume one tries to transform an ℓ_∞ algorithm that runs in time $2^{O(n)}$ into an ℓ_2 algorithm. A square blowup (say) in the dimension implies an algorithm that runs in time $2^{O(n^2)}$.

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 the ℓ_p differs considerably from the 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 very similar to that of Definition 3.1.1 without the normalization factor and with $m = n$.

Definition 5.4.1 (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 vectors in S^{n-1} uniformly at random and then apply the Gram-Schmidt orthogonalization process to them. Let A be the $n \times n$ matrix whose columns are the chosen orthogonal vectors. Then define $f(x) = Ax$.*

We now give two embedding results, the first into ℓ_∞^n and the second into ℓ_p^n 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 5.4.2 *For $f \sim \mathcal{F}(\infty, n)$ it holds that 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 $1 - n^{-\Omega(1)}$

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

where $c(\infty)$ is a global constant.

Proof: The first inequality

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

is easy to verify. For the second we will need a claim and a lemma. Consider a point chosen from μ . Because μ is rotationally invariant and since the rotationally invariant measure on S^{n-1} is unique then $x/\|x\|_2$ is distributed according to σ . Therefore it is enough to show that for x chosen according to μ satisfies with probability at least $1 - n^{-\Omega(1)}$ that

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

We start by calculating the median M_∞ of $\|x\|_\infty = \max |x_i|$ with respect to the standard Gaussian measure on \mathbb{R}^n which is defined as the value of M_∞ for which

$$\mu\{x, \|x\|_\infty \leq M_\infty\} = 1/2 \quad \text{and} \quad \mu\{x, \|x\|_\infty \geq M_\infty\} = 1/2. \quad (5.1)$$

In other words the cube $[-M_\infty, M_\infty]^n$ has a Gaussian measure $\frac{1}{2}$. Since the cube is a “product”,

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

We now calculate

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

Using Fact 7.1.3 we obtain that

$$1 - 2(1 - \Phi(M)) \geq 1 - 2e^{-M^2/2}/M.$$

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

$$1 - 2(1 - \Phi(\sqrt{2 \log n})) \geq 1 - 2e^{-\log n}/\sqrt{2 \log n} = 1 - \frac{2}{n\sqrt{2 \log n}} = 1 - n^{-\Omega(1)}.$$

Moreover,

$$\mu([-M_\infty, M_\infty]^n) = \left(1 - n^{-\Omega(1)}\right)^n \geq 1/2.$$

Therefore $M_\infty \leq \sqrt{2 \log n}$.

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

Lemma 5.4.3 *For a measurable set $A \subset \mathbb{R}^n$ with $\mu(A) = \frac{1}{2}$,*

$$\mu(A_\epsilon) \geq 1 - 2e^{-\epsilon^2/4} \quad (5.2)$$

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

By definition $\mu(\{x; \|x\|_\infty \leq M_\infty\}) = 1/2$ and since $\|\cdot\|_\infty$ is 1-Lipschitz (i.e., $|\|x\|_\infty - \|y\|_\infty| \leq \|x - y\|_2$) then by setting $\epsilon = M_\infty$ in Lemma 5.4.3 we obtain that

$$\mu(\{x; \|x\|_\infty \leq 2M_\infty\}) \geq 1 - 2e^{-M_\infty^2/4}.$$

Thus,

$$\Pr_{x \sim \mu} \left[\|x\|_\infty \leq \sqrt{8 \log n} \right] \geq 1 - n^{-\Omega(1)} \quad (5.3)$$

We now give a lower bound on $\|x\|_2$.

Claim 5.4.4 *For a vector z chosen from the standard n -dimensional normal distribution,*

$$\Pr[\|z\|_2 \leq \frac{1}{2}\sqrt{n}] \leq (1.37)^{-n}.$$

Proof: By a change of variable, we see that

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

On the other hand,

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

where we use that $e^{-\|z\|_2^2/2}/e^{-\|z\|_2^2/2} = e^{-3/2\|z\|_2^2} \geq e^{-3/8n} \geq (0.687)^{-n}$ whenever $\|z\|_2 \leq \frac{1}{2}\sqrt{n}$. The claim follows by combining the two equations $\left(\frac{2}{e^{3/8}}\right)^n \geq (1.37)^n$. ■

From the claim we obtain that with probability at least $1 - 2^{-\Omega(n)}$, $\|x\|_2 \geq \frac{1}{2}\sqrt{n}$. Combining this with (5.3) we obtain

$$\Pr_{x \sim \mu} \left[\left\| \frac{x}{\|x\|_2} \right\|_\infty \leq \sqrt{32 \frac{\log n}{n}} \right] \geq 1 - n^{-\Omega(1)}.$$

and the theorem follows. ■

Embedding ℓ_2 into ℓ_p

Theorem 5.4.5 *For $2 \leq p < \infty$, any $f \sim \mathcal{F}(p, n)$ satisfies that 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 $f \sim \mathcal{F}(p, n)$,

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

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

Before we prove the theorem we will state two lemmas that will be used later. We need to analyze the expectation of $\|x\|_p$ over all $x \in S^{n-1}$. This expectation is given by

$$\int_{S^{n-1}} \|x\|_p d\sigma$$

where σ is the rotationally invariant probability on S^{n-1} . By using polar coordinates we can write

$$\int_{S^{n-1}} \|x\|_p d\sigma = \frac{\Gamma(n/2)}{\sqrt{2}\Gamma((n+1)/2)} \int_{\mathbb{R}^n} \|x\|_p d\mu(x) = (1 + o(1)) \frac{1}{\sqrt{n}} \int_{\mathbb{R}^n} \|x\|_p d\mu(x) \quad (5.4)$$

We now state Equation 2.6 and Lemma 2.7 in [FLM77]. We need to keep those in mind for later.

Lemma 5.4.6 (Equation 2.6 in [FLM77]) *Let M_r be the median of $\|\cdot\|_p$ on S^{n-1} , namely*

$$\sigma\{x, \|x\|_p \leq M_r\} = 1/2 \quad \text{and} \quad \sigma\{x, \|x\|_p \geq M_r\} = 1/2.$$

Then,

$$\sigma\left(\{x; \left|\|x\|_p - M_r\right| \leq \epsilon\}\right) \geq 1 - 4e^{-n\epsilon^2/2} \quad (5.5)$$

Lemma 5.4.7 (Lemma 2.7 in [FLM77]) *There is an absolute constant c such that for M_r , the median of $\|\cdot\|_p$,*

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

We need to separately calculate (5.4) and then use Lemma 5.4.7 and (5.5).

Proof of the theorem: The first inequality

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

is easy to verify. For the second we need to calculate (5.4). So,

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

By using Lemma 5.4.7 we can bound the median by,

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

By setting $\epsilon = n^{1/p-1/2-\xi}$ for small enough $\xi > 0$ in (5.5) it follows that

$$\Pr_f \left[\|f(x)\|_p \leq c(p)n^{1/p-1/2} \cdot \|x\|_2 \right] \geq 1 - 2^{-n^{\Omega(1)}}$$

■

Chapter 6

Open Problems

SVP,CVP

Approximation algorithms for SVP, CVP in the ℓ_p norm would imply approximation algorithms for SVP (CVP) in the ℓ_2 norm. There is still a big gap between the hardness of approximation of SVP (CVP) to the best known algorithm for approximating SVP (CVP). It will be interesting to find an algorithm for SVP (CVP) in the ℓ_p norm maybe ℓ_1 or ℓ_∞ .

Covering Radius Problem (CRP) hardness results

Definition 6.0.8 *The covering radius of a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is the smallest R such that any $x \in \mathbb{R}^n$ is within distance R of \mathcal{L} . More formally,*

$$\rho(\mathcal{L}) = \max_{x \in \mathbb{R}^n} d(x, \mathcal{L}) = \max_{x \in \mathbb{R}^n} \min_{u \in \mathcal{L}} d(x, u)$$

CRP can be not so naturally defined not only on full rank lattice but also for lattice of dimension $n < m$. In this definition we want the lattice points to cover each point in the subspace $\mathbb{R}^n \subseteq \mathbb{R}^m$ where the lattice is embedded into. For this definition of the problem the same reduction for SVP will yield a reduction from CRP in the ℓ_2 norm to CRP in the ℓ_p norm.

It is not clear how CRP behaves as a function of p (ℓ_p). It is a different class of lattice problems that depends on the distances in the whole space in which the lattice is embedded and not only on the distances between lattice points (such as successive minima problems $\lambda_i(\mathcal{L})$). Here the reduction needs to add more vectors so that the lattice will be a full rank lattice. It is not sure however how to add those vectors since short vectors and long vectors might change the solution of the original problem.

Bibliography

- [AD97] Miklós Ajtai and Cynthia 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] Miklós 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, New York, 1996. ACM.
- [Ajt98] Miklós 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.
- [AKKV05] Misha Alekhnovich, Subhash Khot, Guy Kindler, and Nisheeth K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing. submitted, 2005.
- [AKS01] Miklós Ajtai, Ravi 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), New York, 2001. ACM.
- [AR04] Dorit Aharonov and Oded 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.
- [Bal97] Keith 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.
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Math. Ann.*, 296(4):625–635, 1993.
- [BDES02] Shai Ben-David, Nadav Eiron, and Hans Ulrich Simon. Limitations of learning via embeddings in Euclidean half spaces. *J. Mach. Learn. Res.*, 3(Spec. Issue Comput. Learn. Theory):441–461, 2002.
- [Ber97] Bonnie Berger. The fourth moment method. *SIAM J. Comput.*, 26(4):1188–1207, 1997.

- [BS99] Johannes Blömer and Jean-Pierre 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.
- [Cai99] Jin-Yi Cai. Applications of a new transference theorem to Ajtai’s connection factor. In *Fourteenth Annual IEEE Conference on Computational Complexity (Atlanta, GA, 1999)*, pages 205–214. IEEE Computer Soc., Los Alamitos, CA, 1999.
- [CN97] J-Y. Cai and A. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *Proc. 38th Annual IEEE Symp. on Foundations of Computer Science (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. System Sci.*, 59(2):221–239, 1999. 13th Annual IEEE Conference on Computation Complexity (Buffalo, NY, 1998).
- [Din02] Irit Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. *Theoretical Computer Science*, 285(1):55–71, 2002.
- [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.
- [Dvo61] Aryeh 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.
- [Fel71] William Feller. *An introduction to probability theory and its applications. Vol. II.* Second edition. John Wiley & Sons Inc., New York, 1971.
- [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.
- [Gau66] Carl Friedrich Gauss. *Disquisitiones arithmeticae.* Translated into English by Arthur A. Clarke, S. J. Yale University Press, New Haven, Conn., 1966.
- [Hås88] J. Håstad. Dual vectors and lower bounds for the nearest lattice point problem. *Combinatorica*, 8(1):75–81, 1988.
- [Ind00a] Piotr 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.
- [Ind00b] Piotr 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] Piotr 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] Piotr 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, New York, 2003. ACM.
- [JS82] William B. Johnson and Gideon Schechtman. Embedding l_p^m into l_1^n . *Acta Math.*, 149(1-2):71–85, 1982.
- [Kan87] Ravi Kannan. Minkowski’s convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, 1987.
- [Kho03] Subhash 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] Subhash 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.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.
- [LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [LLS90] J. C. Lagarias, H. W. Lenstra, Jr., and C.-P. Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.
- [LO85] J. C. Lagarias and A. M. Odlyzko. Solving low-density subset sum problems. *J. Assoc. Comput. Mach.*, 32(1):229–246, 1985.
- [MG02] Daniele Micciancio and Shafi Goldwasser. *Complexity of lattice problems*. The Kluwer International Series in Engineering and Computer Science, 671. Kluwer Academic Publishers, Boston, MA, 2002. A cryptographic perspective.
- [Mic01] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035 (electronic), 2001.
- [Mul96] K. Mulmuley. Randomized geometric algorithms and pseudorandom generators. *Algorithmica*, 16(4-5):450–463, 1996.
- [Reg03a] Oded 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] Oded Regev. New lattice based cryptographic constructions. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 407–416 (electronic), New York, 2003. ACM.
- [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] Peter 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.

Chapter 7

Appendix

7.1 Proof of Theorem 3.2.3

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 7.1.1 *There exists $T = \sqrt{2 \ln(2m\delta^{-1})}$ such that for any $x \in \mathbb{R}^n$,*

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

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

This result imply Theorem 3.2.3 by taking $f(\cdot) = \frac{1}{T}g(\cdot)$.

We closely follow Indyk's proof. We first need to prove an auxiliary lemma about the tail of the maximum of standard normally distributed independent variables.

Lemma 7.1.2 *Let Y_1, Y_2, \dots, Y_m be i.i.d. normal variables and let $Z = \max(|Y_1|, |Y_2|, \dots, |Y_m|)$. For $\delta, \eta \in (0, 1)$*

- $\Pr[Z < T_1] \leq \eta$ for $T_1 = \sqrt{2 \ln(\frac{2m}{\ln \eta^{-1} \sqrt{2 \ln m \cdot B}})}$
where B is a positive absolute constant.
- $\Pr[Z > T_2] \leq \delta$ for $T_2 = \sqrt{2 \ln(2m\delta^{-1})}$

We use $\Phi(t)$ to 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 now need two simple facts that are given in [Fel71].

Facts 7.1.3 *The following inequalities hold:*

- For any $t > 0$ we have $1 - \Phi(t) \leq e^{-t^2/2}/t$
- There exists a constant $B > 0$ such that for any $t > 0$ we have $1 - \Phi(t) \geq Be^{-t^2/2}/t$

Proof of Lemma 7.1.2: We start with the right tail. For any $t \geq 0$

$$\Pr[Z < t] = (1 - 2(1 - \Phi(t)))^m \leq (1 - 2Be^{-t^2/2}/t)^m \leq \exp(-m2Be^{-t^2/2}/t)$$

For $t = T_1$ the latter expression is almost η . Now we proceed with T_2 . Applying union bound we obtain

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

We assume that the latter expression is not greater than δ . ■

Proof of the theorem: Since g is linear, without loss of generality we can assume that $\|x\|_2 = 1$. Set $T = T_2$ as in Lemma 7.1.2. The second bound follows from Lemma 7.1.2. We now prove the first bound. Since $\|x\|_2 = 1$, it follows by the rotational symmetry of the normal distribution that each coordinate of $g(x)$ is distributed as $N(0, 1)$. Therefore if we would have set T to be T_1 , from Lemma 7.1.2 the first bound would have followed. Therefore, it is sufficient to make sure that $(1 + \epsilon)T_1 \geq T$. Let $\xi = \theta(\epsilon)$ be such that

$$\frac{1 + \xi}{1 - 2\xi} \leq 1 + \epsilon$$

and set

$$m = \max(\delta^{-1}, \ln(\eta^{-1}), 1/B^2, 1/\xi^\alpha)^{1/\xi}$$

for a large enough constant α . Then

$$T/T_1 \leq \frac{\sqrt{2 \ln(2m\delta^{-1})}}{\sqrt{2(\ln(2m) - \ln \ln(\eta^{-1}) - 1/2 \ln(2 \ln m) - \ln(1/B))}}.$$

By dividing the numerator and the denominator by $\sqrt{2 \ln(2m)}$ and simplifying the expression, we obtain

$$T/T_1 \leq \frac{1 + \xi}{1 - 2\xi} \leq 1 + \epsilon$$

■