

A Note on the Distribution of the Distance from a Lattice

Ishay Haviv*

Vadim Lyubashevsky[†]

Oded Regev[‡]

June 7, 2008

Abstract

Let \mathcal{L} be an n -dimensional lattice and let x be a point chosen uniformly from a large ball in \mathbb{R}^n . In this note we consider the distribution of the distance from x to \mathcal{L} , normalized by the largest possible such distance (i.e., the covering radius of \mathcal{L}). By definition, the support of this distribution is $[0, 1]$. We show that there exists a universal constant α_2 that provides a natural ‘threshold’ for this distribution in the following sense. For any $\varepsilon > 0$ there exists a $\delta > 0$ such that for any lattice this distribution has mass at least δ on $[\alpha_2 - \varepsilon, 1]$; moreover, there exist lattices for which the distribution is tightly concentrated around α_2 (and so the mass on $[\alpha_2 + \varepsilon, 1]$ can be arbitrarily small). We also provide several bounds on α_2 and its extension to other ℓ_p norms. We end with an application from the area of computational complexity. Namely, we show that α_2 is exactly the approximation factor of a certain natural AM protocol for the Covering Radius Problem.

1 Introduction

Preliminaries. A (full-rank) *lattice* is defined as the set of all integer combinations

$$\mathcal{L}(b_1, \dots, b_n) = \left\{ \sum_{i=1}^n x_i b_i : x_i \in \mathbb{Z} \text{ for all } 1 \leq i \leq n \right\}$$

of n linearly independent vectors b_1, \dots, b_n in \mathbb{R}^n . This vector set is called a *basis* of the lattice and we often represent it by a matrix B having the basis vectors as columns. We say that a bounded region $P \subseteq \mathbb{R}^n$ is a *fundamental region* of a lattice \mathcal{L} if every element of \mathbb{R}^n can be written uniquely as the sum of an element from P and a lattice vector from \mathcal{L} . In other words, the translates of a fundamental region by lattice vectors tile \mathbb{R}^n . The volume of any fundamental region of a lattice \mathcal{L} is uniquely determined by \mathcal{L} . We define the *determinant* of \mathcal{L} , denoted by $\det(\mathcal{L})$, to be this volume.

One fundamental region is given by the ℓ_p *Voronoi cell*, defined as

$$\text{Vor}_p(\mathcal{L}) = \{x \in \mathbb{R}^n : \text{for all } v \in \mathcal{L}, \text{dist}_p(x, 0) \leq \text{dist}_p(x, v)\},$$

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

[†]University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0404, USA. Supported by NSF ITR 0313241.

[‡]Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by an Alon Fellowship, by the Binational Science Foundation, by the Israel Science Foundation, and by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848.

where $\text{dist}_p(\cdot, \cdot)$ denotes the distance with respect to the ℓ_p norm. In the ℓ_2 case, the Voronoi cells are polytopes and tile \mathbb{R}^n in a face-to-face manner. Another fundamental region of a lattice is given by its *basic parallelepiped*, defined as

$$\mathcal{P}(B) = \left\{ \sum_{i=1}^n x_i b_i : 0 \leq x_i < 1 \text{ for all } 1 \leq i \leq n \right\},$$

where $B = (b_1, \dots, b_n)$ is a basis of the lattice. Notice that the volume of this region equals $|\det(B)|$ and hence $\det(\mathcal{L}) = |\det(B)|$.

The *covering radius* of a lattice is defined as follows.

Definition 1.1 *The covering radius of a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ with respect to the ℓ_p norm is defined as*

$$\rho_p(\mathcal{L}) = \max_{x \in \mathbb{R}^n} \text{dist}_p(x, \mathcal{L}).$$

We define the *normalized distance* of a point x from a lattice \mathcal{L} as the distance of x from \mathcal{L} divided by the covering radius of \mathcal{L} . Finally, for an n -dimensional lattice \mathcal{L}_1 and an n' -dimensional lattice \mathcal{L}_2 we define their direct sum $\mathcal{L}_1 \oplus \mathcal{L}_2$ as the lattice

$$\mathcal{L}_1 \oplus \mathcal{L}_2 = \{(y_1, \dots, y_n, y'_1, \dots, y'_{n'}) : (y_1, \dots, y_n) \in \mathcal{L}_1 \text{ and } (y'_1, \dots, y'_{n'}) \in \mathcal{L}_2\}.$$

We let $\mathcal{L}^{\oplus k}$ denote the direct sum of k copies of \mathcal{L} .

Distance from a lattice. In this note we analyze the distribution of the normalized distance from a lattice of a point x chosen uniformly from a large ball in \mathbb{R}^n . For a lattice \mathcal{L} , a fundamental region P of \mathcal{L} and some $1 \leq p \leq \infty$ let the random variable $Z_{\mathcal{L},p} \in [0, 1]$ be the normalized ℓ_p distance from \mathcal{L} of a point distributed uniformly in P . More formally,

$$Z_{\mathcal{L},p} = \frac{\text{dist}_p(x, \mathcal{L})}{\rho_p(\mathcal{L})},$$

where x is chosen uniformly from P (see Figure 1). It can be shown that the distribution of $Z_{\mathcal{L},p}$ is independent of the choice of P . In particular, by choosing P to be the ℓ_p Voronoi cell $\text{Vor}_p(\mathcal{L})$, the definition of $Z_{\mathcal{L},p}$ simplifies to $Z_{\mathcal{L},p} = \frac{\|x\|_p}{\rho_p(\mathcal{L})}$. Moreover, one can equivalently define $Z_{\mathcal{L},p}$ as the limit as $R \rightarrow \infty$ of the distribution of the normalized distance from \mathcal{L} of a point chosen from a ball of radius R . Finally, note that the distribution of $Z_{\mathcal{L},p}$ is invariant under scaling and orthogonal transformations of \mathcal{L} .

For any $1 \leq p < \infty$, we define $\alpha_p(\mathcal{L})$ as the p th norm of the random variable $Z_{\mathcal{L},p}$, i.e.,

$$\alpha_p(\mathcal{L}) = \sqrt[p]{\mathbf{E} \left[Z_{\mathcal{L},p}^p \right]}.$$

It is easy to see that in addition to being invariant under scaling and orthogonal transformations, $\alpha_p(\mathcal{L})$ is also invariant under the direct sum operation, i.e.,

$$\alpha_p(\mathcal{L}^{\oplus k}) = \alpha_p(\mathcal{L})$$

for any $k \geq 1$ (a lattice parameter satisfying these properties is called a *geometrical invariant* in [7]). Finally, we denote by α_p the infimum of the $\alpha_p(\mathcal{L})$ over all lattices \mathcal{L} . Obviously, it is enough to take the infimum over all lattices \mathcal{L} satisfying $\rho_p(\mathcal{L}) = 1$. We extend the definition to the ℓ_∞ norm by

$$\alpha_\infty = \frac{1}{2}.$$




Figure 1: The density function of the distribution of $Z_{\mathbb{Z}^k, 2}$ for $k = 1, 2, 3, 4$ (from top left to bottom right). For large k the distribution is concentrated around $\frac{1}{\sqrt{3}} \approx 0.577$.

The following theorem proves two fundamental properties of $Z_{\mathcal{L}, p}$ in terms of α_p . First, it shows that for arbitrarily small $\varepsilon > 0$ and any lattice \mathcal{L} , $Z_{\mathcal{L}, p}$ must have some positive mass in the segment $[\alpha_p - \varepsilon, 1]$. Second, it shows that this cannot be improved as there are lattices for which $Z_{\mathcal{L}, p}$ is tightly concentrated around α_p . The proof appears in Section 2.

Theorem 1.2 *For any $1 \leq p \leq \infty$ and any $\varepsilon > 0$ the following holds.*

1. *There exists a $\delta > 0$ such that for any lattice \mathcal{L} ,*

$$\Pr [Z_{\mathcal{L}, p} \geq \alpha_p - \varepsilon] \geq \delta.$$

2. *For arbitrarily large n there exists an n -dimensional lattice \mathcal{L} such that*

$$\Pr [|Z_{\mathcal{L}, p} - \alpha_p| \geq \varepsilon]$$

is exponentially small in n .

Related work. A somewhat related parameter was studied by Conway and Sloane [5, Chapter 21] and by Barnes and Sloane [3] (among others) in the context of quantizing data that is uniformly distributed over a large region of \mathbb{R}^n . This parameter, known as the “normalized second moment”, is defined for an n -dimensional lattice \mathcal{L} as

$$G(\mathcal{L}) = \frac{\mathbf{E}_{x \in \text{Vor}_2(\mathcal{L})} [\|x\|_2^2]}{n \cdot \det(\mathcal{L})^{2/n}}.$$

In other words, $G(\mathcal{L})$ is the expected squared ℓ_2 distance from the lattice divided by the dimension, after the lattice is normalized to have determinant 1. In [5, Chapter 21] formulas for $G(\mathcal{L})$ were provided for some well-known lattices such as root lattices and their duals. In [3] a formula for $G(\mathcal{L})$ for any 3-dimensional lattice \mathcal{L} was given, a result that we will use later in this note (see Theorem 3.6).

The parameters $G(\mathcal{L})$ and $\alpha_2^2(\mathcal{L})$ might seem very similar as they both depend on the second moment of the Voronoï cell, $\mathbf{E}_{x \in \text{Vor}_2(\mathcal{L})} [\|x\|_2^2]$. However, there is one main difference between them: whereas in $G(\mathcal{L})$ we normalize the Voronoï cell to have volume 1, in $\alpha_2^2(\mathcal{L})$ we normalize it to have circumradius 1. This difference is crucial, and is best demonstrated by considering the task of minimizing each of these parameters. Since the ball has the smallest second moment among all bodies of volume 1, we expect lattices whose Voronoï cell is as ‘spherical’ as possible to be minimizers of G . On the other hand, the ball is definitely *not* the minimizer of the second moment among all bodies of circumradius 1: in the limit of large dimension, the second moment of a ball of radius 1 approaches 1, whereas that of a cube of circumradius 1 is $\frac{1}{3}$. This indicates that lattices whose Voronoï cell is ‘cubical’ (such as \mathbb{Z}^n) achieve a smaller α_2 than those whose Voronoï cell is ‘spherical’.

Bounds on α_p . In Section 3 we provide several upper and lower bounds on α_p . For any real $1 \leq p < \infty$ we show that

$$\frac{1}{2} \leq \alpha_p \leq \frac{1}{\sqrt[p]{p+1}}, \quad (1)$$

where the right inequality is achieved by the lattice \mathbb{Z}^n for any $n \geq 1$. Clearly, (1) implies that $\alpha_1 = \frac{1}{2}$. For any $1 \leq p < \infty$ we show the following upper bound, which is better than (1) for any $p > 2$ (see Figure 2).

$$\alpha_p \leq \min_{0 < \beta \leq 1} \frac{1}{2} \cdot \sqrt[p]{\frac{(1+\beta)^{p+2} - (1-\beta)^{p+2}}{\beta^2(p+1)(p+2)}}. \quad (2)$$

Combining the lower bound from (1) and the upper bound from (2) for $\beta = \frac{1}{p}$ we get that

$$\lim_{p \rightarrow \infty} \alpha_p = \frac{1}{2}.$$

Moreover, by choosing $\beta = 1$ one gets

$$\alpha_p \leq \sqrt[p]{\frac{4}{(p+1)(p+2)}}. \quad (3)$$

For the ℓ_2 norm, the three bounds (1), (2) and (3) provide an upper bound of $\frac{1}{\sqrt{3}}$ on α_2 . This gives some evidence to the following conjecture.

Conjecture 1.3 $\alpha_2 = \frac{1}{\sqrt{3}}$.

As an additional evidence to this conjecture, we show in Subsection 3.2 that $\alpha_2(\mathcal{L}) \geq \frac{1}{\sqrt{3}}$ for all lattices of dimension at most three. So any counterexample to the conjecture must be of dimension at least 4. We also show that any lattice of dimension at most three that achieves $\alpha_2(\mathcal{L}) = \frac{1}{\sqrt{3}}$ is generated by a basis of orthogonal vectors. In addition, we verified that some well-studied lattices (such as the lattices from [5, Page 61, Table 2.3]) satisfy $\alpha_2(\mathcal{L}) \geq \frac{1}{\sqrt{3}}$. We also note that α_2 for the well-known lattices A_n and D_n (see [5, Pages 108,117] for the definition) tends to $\frac{1}{\sqrt{3}}$ from above when n tends to infinity, as follows easily from the parameters calculated in [5, Chapter 21].




Figure 2: Upper bounds on α_p as a function of p including the bound from (1) and bounds from (2) for various values of $\beta \in (0, 1]$.

The covering radius problem. Our original motivation for studying α_p comes from the area of computational complexity, as explained in Section 4. The Covering Radius Problem (CRP) in the ℓ_p norm with approximation factor $\gamma \geq 1$ is that of distinguishing between YES instances, which are lattices with ℓ_p covering radius at most d , and NO instances, which are lattices with ℓ_p covering radius bigger than $\gamma \cdot d$. The complexity of CRP was first studied by Guruswami et al. [8]. They presented a simple and natural protocol by which a prover can convince a (randomized) verifier that an instance of CRP is a YES instance for $\gamma = 2$. This protocol implies that CRP with factor $\gamma = 2$ is in the complexity class AM, which very roughly speaking is not much wider than NP. Their analysis is not necessarily tight. Based on Theorem 1.2, we observe in Section 4 that the approximation factor achieved by the protocol is essentially $\frac{1}{\alpha_p}$. This implies the following theorem.

Theorem 1.4 *For any $1 \leq p \leq \infty$ and $\gamma > \frac{1}{\alpha_p}$, CRP in the ℓ_p norm with factor γ lies in AM.*

Notice that assuming Conjecture 1.3, Theorem 1.4 yields that CRP in the ℓ_2 norm with factor γ is in AM for any $\gamma > \sqrt{3}$.

Open questions. The main open question raised by this note is that of determining α_p for any $1 < p < \infty$. Of special interest is the Euclidean norm case — is $\alpha_2 = \frac{1}{\sqrt{3}}$? One possible approach is to show that the second moment of any polytope P with circumradius 1 that has some of the properties satisfied by Voronoi cells is at least $\frac{1}{3}$. For instance, one can prove this for centrally symmetric polytopes whose facets are centrally symmetric, and in which each vector pointing from the center of the cell to the center of a facet is perpendicular to the facet.

A possibly easier task is showing a lower bound of $\frac{1}{\sqrt{3}}$ on $\alpha_2(\mathcal{L})$ for certain families of lattices \mathcal{L} . For example, it will be interesting to show it for all lattices that are associated with positive definite quadratic forms lying in the closure of what is known as Voronoi's principal domain of the first type. Such a result was obtained for the lattice covering problem using techniques of convex optimization and semidefinite programming [11, Chapter 7].

Another interesting open question is to find the smallest $\gamma \geq 1$ for which CRP in the ℓ_p norm with factor γ is in AM. As mentioned before, the AM protocol of [8] achieves a factor of $\frac{1}{\alpha_p}$, but the existence of better protocols is indeed possible.

2 Proof of Theorem 1.2

Lemmas 2.1 and 2.2 imply Theorem 1.2 for any *finite* $p \geq 1$.

Lemma 2.1 *For any $1 \leq p < \infty$ and $\varepsilon > 0$ there exists a constant $\delta > 0$ such that for any lattice \mathcal{L} ,*

$$\Pr \left[Z_{\mathcal{L},p}^p \geq \alpha_p^p - \varepsilon \right] \geq \delta.$$

Proof: The random variable $Z_{\mathcal{L},p}$ satisfies $0 \leq Z_{\mathcal{L},p}^p \leq 1$ and by definition $\mathbf{E} \left[Z_{\mathcal{L},p}^p \right] \geq \alpha_p^p$. By Markov's inequality it follows that $Z_{\mathcal{L},p}^p \geq \alpha_p^p - \varepsilon$ with probability at least $\delta = \frac{\varepsilon}{1 - \alpha_p^p + \varepsilon}$. ■

Lemma 2.2 *For any $1 \leq p < \infty$ and $\varepsilon > 0$ there exists an integer n_0 and a lattice sequence $\{\mathcal{L}_k\}_{k=1}^\infty$, such that the dimension of \mathcal{L}_k is $n_0 k$ and $\Pr \left[\left| Z_{\mathcal{L}_k,p}^p - \alpha_p^p \right| \geq \varepsilon \right]$ is exponentially small in k .*

Proof: Fix $\varepsilon > 0$ and let \mathcal{L}_0 be an n_0 -dimensional lattice satisfying $\alpha_p^p(\mathcal{L}_0) \leq \alpha_p^p + \frac{\varepsilon}{2}$. For any $k \in \mathbb{N}$ we define the $n_0 k$ -dimensional lattice $\mathcal{L}_k = \mathcal{L}_0^{\oplus k}$. Observe that the random variable $Z_{\mathcal{L}_k,p}^p$ is an average of k independent and identically distributed random variables $Z_{\mathcal{L}_0,p}^p$ and that its expectation equals to $\alpha_p^p(\mathcal{L}_0)$. Hence, by the standard estimate of Chernoff (see for example, [1] Appendix A) we have

$$\Pr \left[\left| Z_{\mathcal{L}_k,p}^p - \alpha_p^p \right| \geq \varepsilon \right] \leq \Pr \left[\left| Z_{\mathcal{L}_k,p}^p - \alpha_p^p(\mathcal{L}_0) \right| \geq \frac{\varepsilon}{2} \right] \leq 2 \cdot e^{-\frac{\varepsilon^2 k}{8}},$$

as required. ■

Lemmas 2.3 and 2.4 imply Theorem 1.2 for the ℓ_∞ norm. The former is due to [8] and holds for any $1 \leq p \leq \infty$.

Lemma 2.3 ([8]) *For any $1 \leq p \leq \infty$ and lattice \mathcal{L} ,*

$$\Pr \left[Z_{\mathcal{L},p} \geq \frac{1}{2} \right] \geq \frac{1}{2}.$$

Proof: Let h be a point of ℓ_p distance $\rho_p(\mathcal{L})$ from the lattice (such a point is known as a *deep-hole*) and let P be some fundamental region for \mathcal{L} . By the triangle inequality, for any point x we have

$$\rho_p(\mathcal{L}) = \text{dist}_p(\mathcal{L}, h + \mathcal{L}) \leq \text{dist}_p(x, \mathcal{L}) + \text{dist}_p(x, h + \mathcal{L}). \quad (4)$$

Therefore, we have

$$\begin{aligned} 1 &= \Pr_{x \in P} \left[\frac{\text{dist}_p(x, \mathcal{L}) + \text{dist}_p(x, h + \mathcal{L})}{\rho_p(\mathcal{L})} \geq 1 \right] \\ &\leq \Pr_{x \in P} \left[\frac{\text{dist}_p(x, \mathcal{L})}{\rho_p(\mathcal{L})} \geq \frac{1}{2} \right] + \Pr_{x \in P} \left[\frac{\text{dist}_p(x, h + \mathcal{L})}{\rho_p(\mathcal{L})} \geq \frac{1}{2} \right] \\ &= 2 \cdot \Pr \left[Z_{\mathcal{L},p} \geq \frac{1}{2} \right], \end{aligned}$$

and we are done. ■

Lemma 2.4 For any $0 < \varepsilon < \frac{1}{2}$ and $n \geq 1$ there exists an n -dimensional lattice \mathcal{L} for which $\Pr \left[\left| Z_{\mathcal{L}, \infty} - \frac{1}{2} \right| \geq \varepsilon \right]$ is exponentially small in n .

Proof: Fix $0 < \varepsilon < \frac{1}{2}$ and let D_n be the “checkerboard lattice” defined as

$$D_n = \left\{ x = (x_1, \dots, x_n) \in \mathbb{Z}^n : \sum_{i=1}^n x_i = 0 \pmod{2} \right\}.$$

Clearly, $\rho_\infty(D_n) = 1$ and this is achieved by $(1, 0, \dots, 0)$. To calculate the probability from the lemma we can consider a point chosen uniformly from $P = [0, 2) \times [0, 1)^{n-1}$ since P is a fundamental region of D_n . A necessary condition for the ℓ_∞ distance from D_n to be at least ε -far from $\frac{1}{2}$ is that for each $2 \leq i \leq n$, x_i is outside the range $(\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon)$. Therefore,

$$\Pr \left[\left| Z_{\mathcal{L}, \infty} - \frac{1}{2} \right| \geq \varepsilon \right] \leq \Pr_{x \in P} \left[x_i \notin \left(\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon \right) \text{ for each } 2 \leq i \leq n \right] = (1 - 2\varepsilon)^{n-1}.$$

■

3 Bounds on α_p

The next claim gives some simple bounds on α_p based on ideas from [8]. (The lower bound in the claim can also be seen as an immediate consequence of Lemmas 2.2 and 2.3.)

Claim 3.1 For any real $1 \leq p < \infty$,

$$\frac{1}{2} \leq \alpha_p \leq \frac{1}{\sqrt[p]{p+1}}.$$

In particular,

$$\alpha_1 = \frac{1}{2}.$$

Proof: The right inequality follows from calculating $\alpha_p(\mathbb{Z})$.

$$\alpha_p(\mathbb{Z}) = \sqrt[p]{\mathbf{E} \left[Z_{\mathbb{Z}, p}^p \right]} = \sqrt[p]{\frac{\int_0^{1/2} x^p dx}{\frac{1}{2} \cdot \frac{1}{2^p}}} = \frac{1}{\sqrt[p]{p+1}}.$$

The proof of the left inequality resembles that of Lemma 2.3. Let \mathcal{L} be an arbitrary lattice with some fundamental region P , and let h be a deep-hole. By (4) we have

$$1 \leq \mathbf{E}_{x \in P} \left[\frac{\text{dist}_p(x, \mathcal{L})}{\rho_p(\mathcal{L})} \right] + \mathbf{E}_{x \in P} \left[\frac{\text{dist}_p(x, h + \mathcal{L})}{\rho_p(\mathcal{L})} \right] = 2 \cdot \mathbf{E} [Z_{\mathcal{L}, p}].$$

Now by Jensen’s inequality we get

$$\alpha_p(\mathcal{L}) = \sqrt[p]{\mathbf{E} \left[Z_{\mathcal{L}, p}^p \right]} \geq \mathbf{E} [Z_{\mathcal{L}, p}] \geq \frac{1}{2},$$

and the claim follows. ■

Remark 3.2 We note that the upper bound in Claim 3.1 is achieved by any lattice of the form

$$c_1\mathbb{Z} \oplus c_2\mathbb{Z} \oplus \cdots \oplus c_n\mathbb{Z},$$

where n is a positive integer and $c_i > 0$ for every $1 \leq i \leq n$. In particular, for any $n \geq 1$ and any $1 \leq p < \infty$ we have

$$\alpha_p(\mathbb{Z}^n) = \frac{1}{\sqrt[p]{p+1}}.$$

3.1 The Lattices $\mathcal{L}_{n,k}$

In this subsection we consider a set of lattices that improve the upper bound on α_p from Claim 3.1 for any $p > 2$. For two positive integers $k \leq n$ the lattice $\mathcal{L}_{n,k}$ is the n -dimensional lattice defined as

$$\mathcal{L}_{n,k} = \left\{ x = (x_1, \dots, x_n) \in \mathbb{Z}^n : \sum_{i=1}^n x_i = 0 \pmod{k} \right\}.$$

Notice that $\mathcal{L}_{n,1}$ is the lattice \mathbb{Z}^n that achieves the upper bound from Claim 3.1 and the tight value of α_1 . Furthermore, $\mathcal{L}_{n,2}$ is the lattice D_n that attains the value of α_∞ . Motivated by these examples, we now give an upper bound on the value of $\alpha_p(\mathcal{L}_{n,k})$ for any real $p \geq 1$. We consider k that is linear in n , i.e., $k = \beta n$ for some $0 < \beta \leq 1$, and we think of n as tending to infinity.

Lemma 3.3 For any $1 \leq p < \infty$, $0 < \beta \leq 1$, small enough $\varepsilon > 0$ and large enough integer n for which $k = \beta n$ is an odd integer we have that

$$\alpha_p(\mathcal{L}_{n,k}) \leq \frac{1}{2} \cdot \sqrt[p]{\frac{(1+\beta)^{p+2} - (1-\beta)^{p+2}}{\beta^2(p+1)(p+2)}} + \varepsilon.$$

In order to prove Lemma 3.3 we use a theorem of Kolmogorov and Smirnov as given in [6]. It says that if we take many samples from the uniform distribution over $[0, 1]$ then asymptotically we expect them to be ‘uniformly distributed’ in $[0, 1]$. Below we use z^\uparrow to denote the vector obtained from z by sorting the coordinates in a nondecreasing order.

Theorem 3.4 ([6]) Let $z \in [0, 1]^n$ be a vector each of whose coordinates is chosen independently and uniformly from $[0, 1]$. Then for any $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \Pr \left[\exists i. \left| z_i^\uparrow - \frac{i}{n} \right| \geq \varepsilon \right] = 0.$$

We also need the following claim.

Claim 3.5 For any $1 \leq p < \infty$ and small enough $\varepsilon > 0$ the following holds. Let n be a large enough integer and $k \leq n$ be an odd number. For any $x \in \mathbb{R}^n$ let $y \in \mathbb{Z}^n$ be the integer point closest to x so that the difference $z = x - y$ is in $[-\frac{1}{2}, \frac{1}{2}]^n$. Let $\ell \in \{-\lfloor \frac{k}{2} \rfloor, \dots, \lfloor \frac{k}{2} \rfloor\}$ be such that $\ell = \sum_i y_i \pmod{k}$. If for each $1 \leq i \leq n$, $z_i^\uparrow \in (\frac{i}{n} - \frac{1}{2} - \varepsilon, \frac{i}{n} - \frac{1}{2} + \varepsilon)$ then

$$\text{dist}_p^p(x, \mathcal{L}_{n,k}) \leq \frac{n}{p+1} \left[\left(\frac{1}{2} + \frac{\ell}{n} \right)^{p+1} + \left(\frac{1}{2} - \frac{\ell}{n} \right)^{p+1} \right] + O(\varepsilon \cdot n).$$

Proof: For simplicity assume that $\ell \geq 0$. The case $\ell < 0$ is similar. Since $\mathcal{L}_{n,k}$ is invariant under permutations of the coordinates, we can assume without loss of generality that the coordinates of z appear in a nondecreasing order, i.e., $z = z^\uparrow$.

Denote by $\tilde{y} \in \mathbb{Z}^n$ the point

$$\tilde{y} = (y_1 - 1, \dots, y_\ell - 1, y_{\ell+1}, \dots, y_n).$$

Observe that \tilde{y} is a lattice point of $\mathcal{L}_{n,k}$ since

$$\sum_{i=1}^n \tilde{y}_i = \sum_{i=1}^n y_i - \ell = 0 \pmod{k}.$$

Therefore,

$$\begin{aligned} \text{dist}_p^p(x, \mathcal{L}_{n,k}) &\leq \text{dist}_p^p(x, \tilde{y}) = \sum_{i=1}^n |x_i - \tilde{y}_i|^p = \sum_{i=1}^{\ell} |z_i + 1|^p + \sum_{i=\ell+1}^n |z_i|^p \\ &\leq \sum_{i=1}^{\ell} \left(\frac{i}{n} + \frac{1}{2} + \varepsilon \right)^p + \sum_{i=\ell+1}^n \left(\left| \frac{i}{n} - \frac{1}{2} \right| + \varepsilon \right)^p \\ &\leq \sum_{i=1}^{\ell} \left(\frac{i}{n} + \frac{1}{2} \right)^p + \sum_{i=\ell+1}^n \left| \frac{i}{n} - \frac{1}{2} \right|^p + O(\varepsilon \cdot n). \end{aligned}$$

For large enough n , we can bound the above by an integral,

$$\begin{aligned} \text{dist}_p^p(x, \mathcal{L}_{n,k}) &\leq \int_0^{\ell} \left(\frac{x}{n} + \frac{1}{2} \right)^p dx + \int_{\ell}^n \left| \frac{x}{n} - \frac{1}{2} \right|^p dx + O(\varepsilon \cdot n) \\ &= \int_0^{\ell} \left(\frac{x}{n} + \frac{1}{2} \right)^p dx + \int_{\frac{n}{2}}^{\frac{n}{2} + \ell} \left(\frac{1}{2} - \frac{x}{n} \right)^p dx + \int_{\frac{n}{2}}^n \left(\frac{x}{n} - \frac{1}{2} \right)^p dx + O(\varepsilon \cdot n) \\ &= \frac{n}{p+1} \left[\left(\frac{1}{2} + \frac{\ell}{n} \right)^{p+1} + \left(\frac{1}{2} - \frac{\ell}{n} \right)^{p+1} \right] + O(\varepsilon \cdot n), \end{aligned}$$

and the claim follows. ■

Proof of Lemma 3.3: First, consider the vector $(1, \dots, 1, 0, \dots, 0)$ with $\lfloor \frac{k}{2} \rfloor$ ones. The ℓ_p distance of this vector from $\mathcal{L}_{n,k}$ is $\sqrt[p]{\lfloor \frac{k}{2} \rfloor}$, and therefore

$$\rho_p(\mathcal{L}_{n,k}) \geq \sqrt[p]{\left\lfloor \frac{k}{2} \right\rfloor}. \quad (5)$$

Now our goal is to give an upper bound on the expectation of $\text{dist}_p^p(x, \mathcal{L}_{n,k})$ where x is chosen uniformly from some fundamental region P of $\mathcal{L}_{n,k}$. We take the fundamental region given by $P = [-\frac{k}{2}, \frac{k}{2}) \times [-\frac{1}{2}, \frac{1}{2})^{n-1}$. For x chosen uniformly from P , denote by $y \in \mathbb{Z}^n$ the integer vector closest to x so that $z = x - y$ is in $[-\frac{1}{2}, \frac{1}{2})^n$, and let $\ell \in \{-\lfloor \frac{k}{2} \rfloor, \dots, \lfloor \frac{k}{2} \rfloor\}$ be such that $\ell = \sum_i y_i \pmod{k}$. It is easy to see that ℓ and z are independent, that ℓ is uniformly distributed in $\{-\lfloor \frac{k}{2} \rfloor, \dots, \lfloor \frac{k}{2} \rfloor\}$, and that the coordinates of z are independently and uniformly distributed in $[-\frac{1}{2}, \frac{1}{2})$.

By Theorem 3.4, for large enough n with probability arbitrarily close to 1

$$\forall i. z_i^\uparrow \in \left(\frac{i}{n} - \frac{1}{2} - \varepsilon, \frac{i}{n} - \frac{1}{2} + \varepsilon \right).$$

Hence Claim 3.5 implies that with probability arbitrarily close to 1,

$$\text{dist}_p^p(x, \mathcal{L}_{n,k}) \leq \frac{n}{p+1} \left[\left(\frac{1}{2} + \frac{\ell}{n} \right)^{p+1} + \left(\frac{1}{2} - \frac{\ell}{n} \right)^{p+1} \right] + O(\varepsilon \cdot n).$$

By taking expectations, we obtain that for large enough n ,

$$\begin{aligned} \mathbf{E} [\text{dist}_p^p(x, \mathcal{L}_{n,k})] &\leq \frac{2}{k} \int_0^{\frac{k}{2}} \left(\frac{n}{p+1} \left[\left(\frac{1}{2} + \frac{\ell}{n} \right)^{p+1} + \left(\frac{1}{2} - \frac{\ell}{n} \right)^{p+1} \right] \right) d\ell + O(\varepsilon \cdot n) \\ &= \frac{(n+k)^{p+2} - (n-k)^{p+2}}{2^{p+1} k n^p (p+1)(p+2)} + O(\varepsilon \cdot n). \end{aligned}$$

Finally, using our bound on the covering radius of $\mathcal{L}_{n,k}$ from Inequality (5) and recalling that $k = \beta n$ one has

$$\alpha_p(\mathcal{L}_{n,k}) \leq \frac{1}{2} \cdot \sqrt[p]{\frac{(1+\beta)^{p+2} - (1-\beta)^{p+2}}{\beta^2(p+1)(p+2)}} + O(\varepsilon),$$

and since ε is arbitrary, we are done. \blacksquare

3.2 On α_2 and 3-dimensional Lattices

In this subsection we show that $\alpha_2(\mathcal{L}) \geq \frac{1}{\sqrt{3}}$ for any lattice \mathcal{L} of dimension at most three. We start with some background on 3-dimensional lattices. One associates with any lattice \mathcal{L} a quadratic form $f(x) = x^T(B^T B)x$, where the columns of the matrix B form a basis of \mathcal{L} . For 3-dimensional lattices there is a choice of basis for which f takes the form

$$f(x_1, x_2, x_3) = \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 \rho_{ij} (x_i - x_j)^2$$

for some nonnegative ρ_{ij} satisfying $\rho_{ii} = 0$ and $\rho_{ij} = \rho_{ji}$ where $x_0 = 0$ (see, e.g., [3]). Thus \mathcal{L} is represented by these six parameters $[\rho_{01}, \rho_{02}, \rho_{03}, \rho_{12}, \rho_{13}, \rho_{23}]$.

The following theorem of Barnes and Sloane [3] expresses $G(\mathcal{L})$ in terms of the ρ_{ij} . Recall that $G(\mathcal{L})$ is defined as

$$G(\mathcal{L}) = \frac{\mathbf{E}_{x \in \text{Vor}_2(\mathcal{L})} [\|x\|_2^2]}{3 \cdot \det(\mathcal{L})^{2/3}}.$$

Theorem 3.6 (Theorem 2 in [3]) Any 3-dimensional lattice \mathcal{L} satisfies

$$G(\mathcal{L}) = \frac{D \cdot S_1 + 2S_2 + K}{36D^{4/3}},$$

where¹

$$\begin{aligned} D &= \det(\mathcal{L})^2 = \sum^{(4)} \rho_{01} \rho_{02} \rho_{03} + \sum^{(3)} \rho_{01} \rho_{23} (\rho_{02} + \rho_{03} + \rho_{12} + \rho_{13}), \\ S_1 &= \rho_{01} + \rho_{02} + \rho_{03} + \rho_{12} + \rho_{13} + \rho_{23}, \\ S_2 &= \rho_{01} \rho_{02} \rho_{13} \rho_{23} + \rho_{01} \rho_{03} \rho_{12} \rho_{23} + \rho_{02} \rho_{03} \rho_{12} \rho_{13}, \end{aligned}$$

¹We use here (as in [3]) a notation for symmetric functions, so that $\sum^{(4)} \rho_{01} \rho_{02} \rho_{03}$ for example is an abbreviation for $\rho_{01} \rho_{02} \rho_{03} + \rho_{01} \rho_{12} \rho_{13} + \rho_{02} \rho_{12} \rho_{23} + \rho_{03} \rho_{13} \rho_{23}$. The number (4) indicates the number of summands.

and

$$K = \sum^{(4)} \rho_{01} \rho_{02} \rho_{03} (\rho_{12} + \rho_{13} + \rho_{23}).$$

In addition, Barnes [2] showed that the squared covering radius of a three dimensional lattice is given by

$$\rho_2^2(\mathcal{L}) = \frac{1}{4D} (D \cdot S_1 - K - 4 \min(\rho_{02} \rho_{03} \rho_{12} \rho_{13}, \rho_{01} \rho_{23} \rho_{03} \rho_{12}, \rho_{01} \rho_{23} \rho_{02} \rho_{13})).$$

In particular, we see that $\rho_2^2(\mathcal{L}) \leq \frac{1}{4D} (D \cdot S_1 - K)$. Here we are interested in bounding

$$\alpha_2(\mathcal{L}) = \sqrt{\mathbf{E} \left[Z_{\mathcal{L},2}^2 \right]} = \sqrt{\frac{\mathbf{E}_{x \in \text{Vor}_2(\mathcal{L})} [\|x\|_2^2]}{\rho_2^2(\mathcal{L})}} = \sqrt{\frac{3D^{1/3} \cdot G(\mathcal{L})}{\rho_2^2(\mathcal{L})}}. \quad (6)$$

We derive the following lemma.

Lemma 3.7 *Any lattice \mathcal{L} of dimension at most three satisfies $\alpha_2(\mathcal{L}) \geq \frac{1}{\sqrt{3}}$.*

Proof: We first note that it is enough to consider lattices of dimension exactly 3. More generally, we claim that any lower bound on $\alpha_p(\mathcal{L})$ for all $(k+1)$ -dimensional lattices holds also for all k -dimensional lattices. To show this we consider an arbitrary lattice \mathcal{L} of dimension k . Obviously, the $(k+1)$ -dimensional lattice $\mathcal{L} \oplus (\varepsilon\mathbb{Z})$ satisfies $\alpha_p(\mathcal{L} \oplus (\varepsilon\mathbb{Z})) \rightarrow \alpha_p(\mathcal{L})$ when ε tends to 0, and this completes the argument.

By Theorem 3.6 and Equation (6) we get

$$\alpha_2(\mathcal{L}) = \sqrt{\frac{D \cdot S_1 + 2S_2 + K}{12D\rho_2^2(\mathcal{L})}} \geq \frac{1}{\sqrt{3}} \cdot \sqrt{\frac{D \cdot S_1 + 2S_2 + K}{D \cdot S_1 - K}} \geq \frac{1}{\sqrt{3}}.$$

■

Remark 3.8 *The inequality in Lemma 3.7 is an equality if and only if \mathcal{L} is spanned by orthogonal vectors. Indeed, for any \mathcal{L} spanned by three orthogonal vectors, $\alpha_2(\mathcal{L}) = \frac{1}{\sqrt{3}}$. Moreover, if \mathcal{L} is a lattice of dimension 3 satisfying $\alpha_2(\mathcal{L}) = \frac{1}{\sqrt{3}}$ then $S_2 = K = 0$. This means that any multiplication of four of the ρ_{ij} equals to 0, i.e., at least three of the ρ_{ij} are zeros. It can be shown that if this is the case then \mathcal{L} is spanned by three orthogonal vectors.*

4 On the AM Protocol for CRP

In this section we relate α_p to the complexity of the Covering Radius Problem (CRP) in the ℓ_p norm. More precisely, we show a connection between α_p and the approximation factor of CRP in the ℓ_p norm for which the problem is in the complexity class AM.

Let us start with some basic definitions (see [10] for some background on computational complexity). For any $1 \leq p \leq \infty$ and any factor $\gamma \geq 1$ we define the following computational problem.

Definition 4.1 (Covering Radius Problem) *An instance of GapCRP_γ^p is a pair (B, d) where B is a lattice basis and $d \in \mathbb{Q}$ is a rational number. In YES instances $\rho_p(\mathcal{L}(B)) \leq d$ and in NO instances $\rho_p(\mathcal{L}(B)) > \gamma \cdot d$.*

The complexity class AM is that of promise problems that can be verified by a protocol, as follows. A probabilistic polynomial-time verifier generates a “challenge” based on the input, and sends it to an all-powerful prover. The prover sends back a response, and then the verifier decides whether to accept. We require that for every YES instance the prover can act in such a way that the verifier accepts with probability at least p_1 and that for any NO instance the verifier accepts with probability at most p_2 no matter what strategy is played by the prover, where $0 \leq p_2 < p_1 \leq 1$ are two constants.

Like most other lattice problems, the covering radius problem seems very hard: the best known algorithm solves GapCRP_γ^p for any $\gamma > 1$ in exponential time [8]. On the other hand, the covering radius problem exhibits some unique properties from a computational complexity point of view (see [8] for more details). One of the most interesting facts in this respect is that for any $1 \leq p \leq \infty$, GapCRP_2^p is in AM [8]. The proof of this fact is relatively simple, and follows from the following AM protocol. Given an instance (B, d) , the verifier sends to the prover a uniformly random point $x \in \mathcal{P}(B)$ and the prover has to provide a lattice point $y \in \mathcal{L}(B)$ such that $\text{dist}_p(x, y) \leq d$. Clearly, if (B, d) is a YES instance then the prover can act in a way that the verifier accepts with probability 1. On the other hand, if (B, d) is a NO instance then by Lemma 2.3, with probability at least $\frac{1}{2}$ over the choice of x there is no lattice point $y \in \mathcal{L}(B)$ such that $\text{dist}_p(x, y) \leq d$.

Moreover, it was shown in [9] that there exists a constant $0 < \delta < 1$ such that for any large enough finite p and any $\gamma < \frac{3}{2} \cdot \sqrt[p]{\delta}$, GapCRP_γ^p is hard for Π_2 , a complexity class in the second level of the polynomial-time hierarchy. For the ℓ_∞ norm, Π_2 -hardness was shown there for any γ less than $\frac{3}{2}$. It is interesting to mention in this context a result of Boppana et al. [4] that states that if a Π_2 -hard problem is in AM then the polynomial-time hierarchy collapses to the second level, an event which is considered to be highly unlikely. Thus we do not expect GapCRP_γ^p to be in AM for some γ for which it is Π_2 -hard.

Our analysis of the AM protocol above is not necessarily tight. An interesting open question raised in [8] is to determine the approximation factor achieved by this natural protocol. Here we show that this factor is essentially $\frac{1}{\alpha_p}$.

Lemma 4.2 *For any $1 \leq p \leq \infty$, the AM protocol of GapCRP_γ^p given in [8] works for any $\gamma > \frac{1}{\alpha_p}$ and fails for any $\gamma < \frac{1}{\alpha_p}$.*

Proof: Consider the AM protocol from [8] mentioned above. Clearly, if (B, d) is a YES instance then for any $x \in \mathbb{R}^n$ the prover can provide $y \in \mathcal{L}(B)$ such that $\text{dist}_p(x, y) \leq d$. By Item 1 of Theorem 1.2, for any factor γ bigger than $\frac{1}{\alpha_p}$ any NO instance is rejected by the verifier with constant probability. On the other hand, by Item 2 of Theorem 1.2, for any factor γ smaller than $\frac{1}{\alpha_p}$ there exist NO instances for which the verifier accepts with probability exponentially close to one, and the lemma follows. ■

Remark 4.3 *One may wonder whether the protocol of [8] can be improved in the following way: the verifier sends a polynomial number of challenges to the prover (instead of one) and accepts if and only if the prover is able to respond to all of them. However, by Item 2 of Theorem 1.2, there are still NO instances which the verifier accepts with probability exponentially close to one.*

Acknowledgement

We thank the anonymous referees for their useful comments.

References

- [1] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience [John Wiley & Sons], New York, second edition, 2000.
- [2] E. S. Barnes. The covering of space by spheres. *Canad. J. Math.*, 8:293–304, 1956.
- [3] E. S. Barnes and N. J. A. Sloane. The optimal lattice quantizer in three dimensions. *SIAM J. Algebraic Discrete Methods*, 4(1):30–41, 1983.
- [4] R. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25:127–132, May 1987.
- [5] J. H. Conway and N. J. Sloane. *Sphere packings, lattices and groups*. Springer Verlag, 3rd edition, 1998.
- [6] J. Durbin. *Distribution Theory for Tests Based on the Sample Distribution Function*. SIAM CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, 1973.
- [7] G. D. Forney, Jr. On the duality of coding and quantizing. In *Coding and quantization (Piscataway, NJ, 1992)*, volume 14 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 1–14. Amer. Math. Soc., Providence, RI, 1993.
- [8] V. Guruswami, D. Micciancio, and O. Regev. The complexity of the covering radius problem on lattices and codes. *Computational Complexity*, 14(2):90–121, 2005. Preliminary version in CCC’04.
- [9] I. Haviv and O. Regev. Hardness of the covering radius problem on lattices. In *Proc. of 21th IEEE Annual Conference on Computational Complexity (CCC)*, 2006.
- [10] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley Publishing Company, Reading, MA, 1994.
- [11] F. Vallentin. Sphere covering, lattices, and tilings (in low dimensions), 2003. Ph.D. thesis, Munich University of Technology.