

Hardness of the Covering Radius Problem on Lattices*

Ishay Haviv[†]

Oded Regev[‡]

August 7, 2012

Abstract

We provide the first hardness result for the Covering Radius Problem on lattices (CRP). Namely, we show that for any large enough $p \leq \infty$ there exists a constant $c_p > 1$ such that CRP in the ℓ_p norm is Π_2 -hard to approximate to within any constant factor less than c_p . In particular, for the case $p = \infty$, we obtain the constant $c_\infty = 3/2$. This gets close to the factor 2 beyond which the problem is not believed to be Π_2 -hard (Guruswami et al., Computational Complexity, 2005).

1 Introduction

A *lattice* is the set of all integer combinations of some linearly independent vectors in \mathbb{R}^n . Given a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and some $1 \leq p \leq \infty$, the covering radius of \mathcal{L} in the ℓ_p norm is the smallest number d , such that ℓ_p balls of radius d centered around all lattice points in \mathcal{L} cover the entire space. Equivalently, the covering radius is the smallest d such that *for any point in \mathbb{R}^n there exists a lattice point within distance at most d* . In the Covering Radius Problem in the ℓ_p norm (CRP ^{p}), given a lattice and some value d , we are supposed to decide if the covering radius in the ℓ_p norm is at most d . It follows from the definition that CRP ^{p} is in the complexity class Π_2 of the second level of the polynomial-time hierarchy. However, not much is known about its hardness.

In the last decade computational problems on lattices have been extensively studied, and there are many known hardness results in this area. Some of the main and natural lattice problems are the Shortest Vector Problem (SVP), the Closest Vector Problem (CVP) and the Shortest Independent Vectors Problem (SIVP). All these problems are known to be NP-complete for any ℓ_p norm.¹ Moreover, they are all hard to approximate to within some super-constant factors. For example, SVP and CVP are hard to approximate to within $n^{c/\log \log n}$ for some constant $c > 0$ for any ℓ_p norm [5, 12, 9]. Arguably, CRP is the only natural lattice problem that has no known hardness result. One indication that the problem is hard is given by the fact that the analogous problem for *linear codes* is Π_2 -hard in its exact and approximation variants [15, 6]. The problem on lattices is not even known to be NP-hard.

*A preliminary version of the results presented here appeared as part of a paper in the proceedings of the 21st IEEE Conference on Computational Complexity, 2006.

[†]School of Computer Science, The Academic College of Tel Aviv-Yaffo, Israel. Work done while at Tel Aviv University.

[‡]Blavatnik School of Computer Science, Tel Aviv University, and CNRS, Département d'Informatique, École normale supérieure, Paris. Supported by an Alon Fellowship, by the Binational Science Foundation, by the Israel Science Foundation, by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848, and by the European Research Council (ERC) Starting Grant.

¹To be precise, for SVP with $p < \infty$ this is only known under randomized reductions.

The study of the covering radius on lattices from a computational point of view was initiated by Guruswami et al. in [6]. Among other things, it was shown there that approximating CRP^2 to within $\gamma(n)$ can be done in exponential time $2^{O(n)}$ for any constant $\gamma(n) > 1$ and in polynomial-time for some $\gamma(n) = 2^{O(n \log \log n / \log n)}$.² In addition, it was shown that CRP^2 is in AM for $\gamma(n) = 2$, in coAM for $\gamma(n) = \sqrt{n/\log n}$, and in $\text{NP} \cap \text{coNP}$ for $\gamma(n) = \sqrt{n}$. As indicated in [6], the AM protocol for $\gamma(n) = 2$ can be extended to CRP^p for any $1 \leq p \leq \infty$. An extension to arbitrary $p \geq 2$ of the containment in coNP was shown in [20] and that of the containment in NP was shown in [7]. In another somewhat related result, Kannan [11] showed that for any fixed dimension, the problem of computing the covering radius of a lattice with respect to a given input norm, defined by a convex polytope specified as a system of linear inequalities, can be solved in polynomial time.

In this paper we provide the first hardness proof of CRP on lattices, solving an open question of [6]. Namely, we show that CRP^p is Π_2 -hard to approximate to within some constant factor for any large enough norm $p \leq \infty$. For small p , such as the interesting case of ℓ_2 , the problem remains open. We remark that this is not the first time that a lattice problem is shown hard only for large norms. For example, SVP was shown to be NP-hard in the ℓ_∞ norm already in 1981 by van Emde Boas [23], while the hardness question in any other norm remained open till the work of Ajtai in 1998 [1]. As another example, the hardness of approximating SVP to within arbitrarily large constants was first established by Dinur for the ℓ_∞ norm [4], then by Khot [13] for large enough norms, and only then was extended to all norms [12].

Theorem 1.1 *For any large enough $p \leq \infty$, there exists a constant $c_p > 1$ such that CRP^p is Π_2 -hard to approximate within any factor less than c_p .*

The proof of Theorem 1.1 is based on a reduction from a problem known as GroupColoring, which is a Π_2 -variant of the usual graph coloring problem. The standard decision version of this problem was shown to be Π_2 -hard by Král and Nejedlý [14], and this fact suffices to prove the case $p = \infty$ of our theorem. For the general case of sufficiently large p , we require the Π_2 -hardness of an approximation version of GroupColoring. We prove this Π_2 -hardness result by modifying the reduction in [14] to use a certain bounded occurrence version of $\forall\exists$ -3-SAT which was shown Π_2 -hard in [10].

Open Questions

Our work raises some interesting open questions.

- The main open question is whether CRP^p is Π_2 -hard for small values of $p \geq 1$. Of special interest is CRP^2 in the Euclidean ℓ_2 norm, which is conjectured in [16] to be Π_2 -hard.
- It was shown in [6] that CRP on linear codes is NP-hard to approximate to within any constant factor. It would be very interesting to show a similar result for CRP on lattices.
- It is interesting to find the largest value of γ for which approximating CRP^∞ to within γ is Π_2 -hard. In the current paper it is shown that it is at least $3/2$. On the other hand, the AM protocol of [6], whose exact performance was studied in [8], yields that approximating CRP^∞ to within $\gamma \geq 2$ is not Π_2 -hard unless the polynomial-time hierarchy collapses [3].

²To be precise, [6] only describe *randomized* approximation algorithms for these two tasks. The deterministic algorithms follow by replacing the use of the randomized algorithm of [2] with the deterministic one in [19].

Outline

The rest of the paper is organized as follows. In Section 2, we introduce some basic definitions and notions and give background about lattices and group colorings. In Section 3, we prove Theorem 1.1. Finally, in Section 4 we prove the Π_2 -hardness of the approximation variant to the GroupColoring problem.

2 Preliminaries

2.1 General

Let \mathbb{R} , \mathbb{Q} and \mathbb{Z} be the sets of the reals, the rationals and the integers, respectively. If $S \subseteq \mathbb{R}^n$ is an arbitrary region of space, and $x \in \mathbb{R}^n$ is a vector, $S + x = \{y + x : y \in S\}$ denotes a copy of S shifted by x . By \mathbb{Z}_q we denote the ring of integers modulo q , which is the cyclic Abelian group of order q . The ℓ_p norm of a vector $x \in \mathbb{R}^n$ is $\|x\|_p = \sqrt[p]{\sum |x_i|^p}$, and its ℓ_∞ norm is $\|x\|_\infty = \max_i |x_i|$. The associated distance between two vectors $x, y \in \mathbb{R}^n$ is $\text{dist}_p(x, y) = \|x - y\|_p$. The ℓ_p distance of a point x from a set of points S is denoted by $\text{dist}_p(x, S) = \inf_{y \in S} \text{dist}_p(x, y)$.

2.2 Lattices

A *lattice* is a discrete additive subgroup of \mathbb{R}^n . Equivalently, it is the set of all integer combinations

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

of m linearly independent vectors b_1, \dots, b_m in \mathbb{R}^n ($n \geq m$). This set of vectors is called a *basis* of the lattice. A basis can be represented by a matrix having the basis vectors as columns. If the rank m equals the dimension n , then the lattice is called *full rank*. All lattices in this paper are full rank.

Definition 2.1 *The covering radius in the ℓ_p norm of a full-rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is defined as*

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

Hence, $\rho_p(\mathcal{L}) \leq d$ means that for any $x \in \mathbb{R}^n$ there exists a lattice point $y \in \mathcal{L}$ such that $\text{dist}_p(x, y) \leq d$. Conversely, $\rho_p(\mathcal{L}) > d$ means that there exists some $x \in \mathbb{R}^n$ such that any lattice point $y \in \mathcal{L}$ satisfies $\text{dist}_p(x, y) > d$. For any real $1 \leq p \leq \infty$ and any approximation factor $\gamma \geq 1$ we define the following computational problem.

Definition 2.2 (Covering Radius Problem) *An instance of GapCRP_γ^p is a pair (B, d) where B is a full-rank 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 $\rho_p(\mathcal{L}(B)) > \gamma \cdot d$.*

2.3 Group Coloring

Colorings of graph vertices is one of the most popular areas in Graph Theory. One classical problem is the 3-coloring problem, where given a graph, we are asked to color its vertices by $\{0, 1, 2\}$ (equivalently \mathbb{Z}_3) in such a way that no two adjacent vertices have the same color. As is well known, 3-coloring is an

NP-complete problem. In this paper we consider a variant of this problem known as GroupColoring. Let $G = (V, E)$ be a directed graph and let A be some Abelian group. For an edge-labelling $\varphi : E \rightarrow A$ and a vertex coloring $c : V \rightarrow A$, we say that an edge $(u, v) \in E$ is *satisfied* if $c(u) - c(v) \neq \varphi(u, v)$. For a fixed edge-labelling $\varphi : E \rightarrow \mathbb{Z}_3$, we can ask whether there exists a coloring $c : V \rightarrow \mathbb{Z}_3$ that satisfies each oriented edge $(u, v) \in E$. Such a coloring is called a *legal coloring*. If for any φ there exists a legal coloring, we say that G is \mathbb{Z}_3 -colorable. Similarly, for an Abelian group A , a directed graph $G = (V, E)$ is said to be *A-colorable* if for every edge-labelling $\varphi : E \rightarrow A$ there is a vertex coloring $c : V \rightarrow A$ such that $c(u) - c(v) \neq \varphi(u, v)$ for each oriented edge $(u, v) \in E$. It is easy to see that A-colorability is in fact a property of the underlying undirected graph, and does not depend on the specific orientation of the edges.

Definition 2.3 (Group Coloring) *For an Abelian group A , the GroupColoring $_A$ problem is that of deciding whether a given (directed) graph $G = (V, E)$ is A-colorable.*

The GroupColoring problem was shown to be Π_2 -complete in [14] for any fixed Abelian group A of order at least 3. It can easily be seen that a graph G is \mathbb{Z}_2 -colorable if and only if it is a forest. Therefore, for $A = \mathbb{Z}_2$ the problem lies in P. The one-sided error approximation version of GroupColoring is defined as follows.

Definition 2.4 (GroupColoring $_A[\alpha, 1]$) *Given a graph $G = (V, E)$, define*

$$\eta_A(G) = \min_{\varphi: E \rightarrow A} \max_{c: V \rightarrow A} |\{(u, v) \in E : c(u) - c(v) \neq \varphi(u, v)\}|.$$

In words, $\eta_A(G)$ is the maximal r such that for any $\varphi : E \rightarrow A$ there exists a coloring $c : V \rightarrow A$ such that at least r of the edges are satisfied. In YES instances, G is A-colorable, and in NO instances $\eta_A(G) \leq \alpha|E|$.

In Section 4 we show that GroupColoring $_{\mathbb{Z}_3}[\alpha, 1]$ is Π_2 -hard for some $0 < \alpha < 1$.

3 Hardness of GapCRP p

In this section we prove Theorem 1.1 by a reduction from GroupColoring $_A[\alpha, 1]$. In our reduction, we only consider $A = \mathbb{Z}_q$, the cyclic group of order $q \geq 3$. In fact, our best hardness result is obtained for $q = 3$, so the reader can think of q as being 3. Let $G = (V, E)$ be an instance of GroupColoring $_{\mathbb{Z}_q}$ problem with n edges and k vertices. Our goal is to construct a lattice \mathcal{L}_G such that if G is \mathbb{Z}_q -colorable then the covering radius of \mathcal{L}_G is small and otherwise it is large. Fix some orientation of G . Every vertex coloring $c : V \rightarrow \mathbb{Z}_q$ induces an edge-labelling $\varphi : E \rightarrow \mathbb{Z}_q$ defined by $\varphi(u, v) = c(u) - c(v)$ for each oriented edge $(u, v) \in E$. The output of this reduction is (a basis of) the lattice $\mathcal{L}_G \subseteq \mathbb{Z}^n$ defined as the set of all integer vectors that, when reduced modulo q , correspond to an edge-labelling induced by some vertex coloring of G . Notice that \mathcal{L}_G is a lattice, since it is an additive subgroup of \mathbb{Z}^n .

An equivalent definition of \mathcal{L}_G is the following: Let us define a matrix $C \in \{-1, 0, 1\}^{n \times k}$ with n rows, one for each edge in E , and k columns, one for each vertex in V . Assume that the vertex set of G is $V = \{v_1, \dots, v_k\}$ and that its edge set is $E = \{e_1, \dots, e_n\}$. The entries of the matrix are defined by

$$C_{i,j} = \begin{cases} 1, & \text{if } e_i = (v_j, w) \text{ for some vertex } w, \\ -1, & \text{if } e_i = (w, v_j) \text{ for some vertex } w, \\ 0, & \text{otherwise.} \end{cases}$$

Then \mathcal{L}_G can also be defined as

$$\mathcal{L}_G = \{x \in \mathbb{Z}^n : \text{there exists } y \in \mathbb{Z}^k \text{ such that } x = Cy \pmod{q}\}.$$

A basis of \mathcal{L}_G can be easily constructed in polynomial time by duality (see the preliminaries in [18]).

The main property of this reduction is the following: If G is \mathbb{Z}_q -colorable, then for any integer vector $x \in \mathbb{Z}^n$ there exists a lattice vector $y \in \mathcal{L}_G$, such that for each $1 \leq i \leq n$, $x_i \neq y_i \pmod{q}$. Moreover, if $\eta_{\mathbb{Z}_q}(G) \leq \alpha n$, there exists an integer vector $x \in \mathbb{Z}^n$ such that for any $y \in \mathcal{L}_G$, $x_i = y_i \pmod{q}$ for at least an α fraction of the coordinates $1 \leq i \leq n$.

The next lemma is the main argument in the correctness of the reduction.

Lemma 3.1 *For any graph $G = (V, E)$, integer $q \geq 3$ and $1 \leq p \leq \infty$, if G is a YES instance of $\text{GroupColoring}_{\mathbb{Z}_q}[\alpha, 1]$ then $\rho_p(\mathcal{L}_G) \leq \sqrt[p]{n} \cdot \frac{q-1}{2}$, and if G is a NO instance of $\text{GroupColoring}_{\mathbb{Z}_q}[\alpha, 1]$ then $\rho_p(\mathcal{L}_G) \geq \sqrt[p]{(1-\alpha)n} \cdot \frac{q}{2}$.*

In particular, for the case $p = \infty$ we get $\rho_\infty(\mathcal{L}_G) \leq \frac{q-1}{2}$ if G is a YES instance and $\rho_\infty(\mathcal{L}_G) \geq \frac{q}{2}$ otherwise. Note that the latter inequality is, in fact, an equality since $q \cdot \mathbb{Z}^n \subseteq \mathcal{L}_G$.

Proof: If G is a YES instance, i.e., G is \mathbb{Z}_q -colorable, then for any $x \in \mathbb{Z}^n$ there exists a lattice point $y \in \mathcal{L}_G$, such that $x_i \neq y_i \pmod{q}$ for any coordinate $1 \leq i \leq n$. For any point $x \in \mathbb{R}^n$ consider a point $z \in \mathbb{Z}^n + (\frac{q}{2}, \dots, \frac{q}{2})$ such that $|x_i - z_i| \leq \frac{1}{2}$ for each $1 \leq i \leq n$. Observe that there exists $y \in \mathcal{L}_G$ such that for each $1 \leq i \leq n$ the number $z_i - y_i$, when reduced modulo q , is not equal to $\frac{q}{2}$. By adding to y an appropriate vector from $q \cdot \mathbb{Z}^n$, it is possible to get $y' \in \mathcal{L}_G$, such that each coordinate i satisfies $|z_i - y'_i| \leq \frac{q}{2} - 1$. Thus, for this $y' \in \mathcal{L}_G$, each coordinate i satisfies

$$|x_i - y'_i| = |x_i - z_i + z_i - y'_i| \leq |x_i - z_i| + |z_i - y'_i| \leq \frac{1}{2} + \frac{q}{2} - 1 = \frac{q-1}{2}.$$

So the distance between x and the lattice \mathcal{L}_G satisfies

$$\text{dist}_p(x, \mathcal{L}_G) \leq \sqrt[p]{n} \cdot \frac{q-1}{2},$$

and this gives us the required bound for the covering radius of \mathcal{L}_G in the ℓ_p norm.

On the other hand, if G is a NO instance, then there exists an integer vector $x \in \mathbb{Z}^n$ for which any lattice vector $y \in \mathcal{L}_G$ satisfies $x_i = y_i \pmod{q}$ in at least $(1-\alpha)n$ of the coordinates $1 \leq i \leq n$. This means that the vector $z = x + (\frac{q}{2}, \dots, \frac{q}{2})$ satisfies $z_i - y_i \in \frac{q}{2} + q \cdot \mathbb{Z}$ in at least $(1-\alpha)n$ of the coordinates i for any lattice point $y \in \mathcal{L}_G$. Hence,

$$\text{dist}_p(z, \mathcal{L}_G) \geq \frac{q}{2} \cdot \sqrt[p]{(1-\alpha)n},$$

and the lemma follows. ■

We are ready to prove the main result of this section assuming the Π_2 -hardness of $\text{GroupColoring}_{\mathbb{Z}_3}[\alpha, 1]$ for some $0 < \alpha < 1$ given in the next section.

Theorem 3.2 *For any $1 \leq p < \infty$, let c_p denote $3 \cdot \sqrt[p]{c}/2$ and let $c_\infty = 3/2$ where c is a universal constant. Then, for any p for which $c_p > 1$ and any $\varepsilon > 0$, $\text{GapCRP}_{c_p-\varepsilon}^p$ is Π_2 -complete.*

Proof: For any $\gamma \geq 1$ and $1 \leq p \leq \infty$ the problem GapCRP_γ^p is in Π_2 (see [17, Page 137]).

We now prove hardness by a reduction from the problem $\text{GroupColoring}_{\mathbb{Z}_3}[\alpha, 1]$, where α is a constant for which $\text{GroupColoring}_{\mathbb{Z}_3}[\alpha, 1]$ is Π_2 -hard. For a graph $G = (V, E)$, construct the lattice \mathcal{L}_G . Lemma 3.1 shows that if G is a YES instance of $\text{GroupColoring}_{\mathbb{Z}_3}[\alpha, 1]$ then $\rho_p(\mathcal{L}_G) \leq \sqrt[p]{n}$ and if G is a NO instance of $\text{GroupColoring}_{\mathbb{Z}_3}[\alpha, 1]$ then $\rho_p(\mathcal{L}_G) \geq \sqrt[p]{(1-\alpha)n} \cdot \frac{3}{2}$. We obtain that GapCRP^p is Π_2 -hard to approximate to within any factor less than

$$c_p = \frac{3 \cdot \sqrt[p]{(1-\alpha)}}{2},$$

which is greater than 1 for any $p > p' = -\log_{3/2}(1-\alpha)$. Notice, that for the ℓ_∞ norm we get the constant $c_\infty = 3/2$. ■

4 Hardness of Approximation of GroupColoring

The main result of this section is the following.

Theorem 4.1 *The problem $\text{GroupColoring}_{\mathbb{Z}_3}[\alpha, 1]$ is Π_2 -hard for some constant $0 < \alpha < 1$.*

Theorem 4.1 can be extended to any Abelian group of order at least 3 (as in [14]), but for simplicity, we concentrate on the case $A = \mathbb{Z}_3$, which is of greatest interest for us. The proof is essentially the same as the one in [14] except that we reduce from $\forall\exists$ -E3-SAT-B instead of $\forall\exists$ -E3-SAT as in their case. These are some of the basic approximation problems in the second polynomial-time hierarchy (see [21, 22] for a recent survey on the topic of completeness and hardness of approximation in the polynomial-time hierarchy).

Definition 4.2 ($\forall\exists$ -3-SAT[$1 - \varepsilon, 1$]) *An instance of $\forall\exists$ -3-SAT[$1 - \varepsilon, 1$] is a 3-CNF Boolean formula $\Psi(X, Y)$ over two sets of variables. We refer to variables in X as universal variables and to those in Y as existential variables. In YES instances, for every assignment to X there exists an assignment to Y such that the clauses of Ψ are all satisfied. In NO instances, there exists an assignment to X such that for every assignment to Y at most a $1 - \varepsilon$ fraction of the clauses are satisfied.*

For an integer $B > 0$ the problem $\forall\exists$ -3-SAT-B[$1 - \varepsilon, 1$] is defined similarly except that each variable occurs at most B times in Ψ . In the instances of the problem $\forall\exists$ -E3-SAT-B[$1 - \varepsilon, 1$] the number of literals in each clause is exactly 3 (as opposed to being at most 3).

Theorem 4.3 ([10]) *The problem $\forall\exists$ -E3-SAT-B[$1 - \varepsilon, 1$] is Π_2 -hard for some constants B and $\varepsilon > 0$.*

4.1 The Reduction

The construction in [14] uses some graph gadgets satisfying various properties. Figure 1 and the next two lemmas summarize slight variants of these gadgets for the special case $A = \mathbb{Z}_3$ and their properties. Each lemma contains two parts: the first one is used for the completeness proof and the second is used for the soundness proof.

The reduction outputs a graph G , that contains one universal gadget for each universal variable and one existential gadget for each existential variable. The gadgets are edge-disjoint, but they all have one special vertex w in common. Each gadget contains two disjoint sets of vertices: T for the positive literals of the variable and F for its negative literals.

For a vertex coloring $c : V \rightarrow A$ and a *partial coloring* $c' : U \rightarrow A$ for some $U \subseteq V$ in a graph $G = (V, E)$, we say that c *agrees* with c' , if for every $u \in U$, $c(u) = c'(u)$.

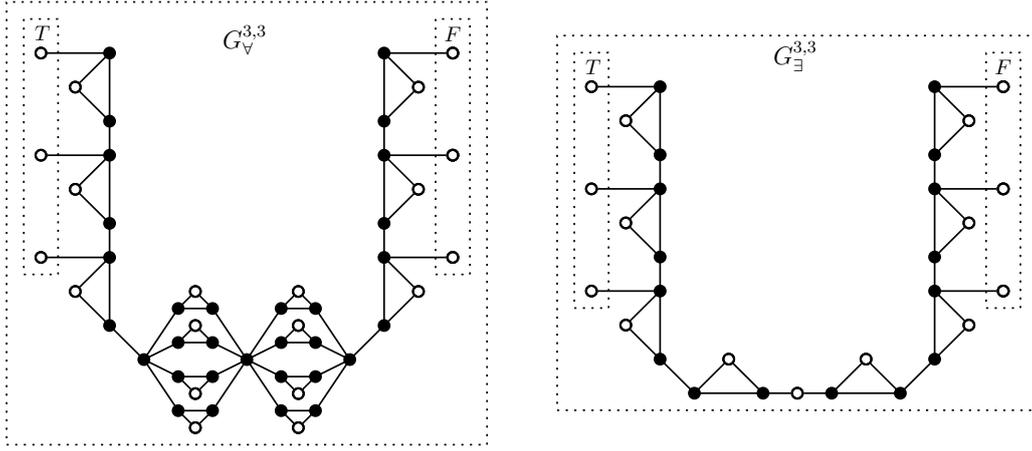


Figure 1: The universal gadget (left) and the existential gadget (right). The vertex w is not shown. All hollow vertices are connected to w .

Lemma 4.4 (Universal Gadget) For any $k, \ell \geq 0$ there exists an efficiently constructible directed graph $G_{\forall}^{k, \ell} = (V \cup \{w\}, E)$ with two disjoint subsets of vertices $T, F \subseteq V$, satisfying $|T| = k$, $|F| = \ell$, $|V| = O(k + \ell)$, $|E| = O(k + \ell)$ such that,

1. For any $\varphi : E \rightarrow \mathbb{Z}_3$ at least one of the following holds:

- (a) There exists a partial coloring $c' : T \cup \{w\} \rightarrow \mathbb{Z}_3$ with $c'(w) = 0$, such that for every partial coloring $c'' : F \rightarrow \mathbb{Z}_3$ that satisfies $c''(v) \neq \varphi(v, w)$ for all $v \in F$,³ there exists a legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ that agrees with c' and with c'' .
- (b) There exists a partial coloring $c' : F \cup \{w\} \rightarrow \mathbb{Z}_3$ with $c'(w) = 0$, such that for every partial coloring $c'' : T \rightarrow \mathbb{Z}_3$ that satisfies $c''(v) \neq \varphi(v, w)$ for all $v \in T$, there exists a legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ that agrees with c' and with c'' .

2. Both the following hold:

- (a) There exists an edge-labelling $\varphi_1 : E \rightarrow \mathbb{Z}_3$ such that every legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ colors all the vertices of F by $c(w)$.
- (b) There exists an edge-labelling $\varphi_2 : E \rightarrow \mathbb{Z}_3$ such that every legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ colors all the vertices of T by $c(w)$.

Lemma 4.5 (Existential Gadget) For any $k, \ell \geq 0$ there exists an efficiently constructible directed graph $G_{\exists}^{k, \ell} = (V \cup \{w\}, E)$ with two disjoint subsets of vertices $T, F \subseteq V$, satisfying $|T| = k$, $|F| = \ell$, $|V| = O(k + \ell)$, $|E| = O(k + \ell)$ such that,

1. For any $\varphi : E \rightarrow \mathbb{Z}_3$ both the following properties hold:

- (a) There exists a partial coloring $c' : T \cup \{w\} \rightarrow \mathbb{Z}_3$ with $c'(w) = 0$, such that for every partial coloring $c'' : F \rightarrow \mathbb{Z}_3$ that satisfies $c''(v) \neq \varphi(v, w)$ for all $v \in F$, there exists a legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ that agrees with c' and with c'' .

³This causes the edge (v, w) to be satisfied, and makes it possible to extend the coloring legally.

- (b) *There exists a partial coloring $c' : F \cup \{w\} \rightarrow \mathbb{Z}_3$ with $c'(w) = 0$, such that for every partial coloring $c'' : T \rightarrow \mathbb{Z}_3$ that satisfies $c''(v) \neq \varphi(v, w)$ for all $v \in T$, there exists a legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ that agrees with c' and with c'' .*
2. *There exists an edge-labelling $\varphi : E \rightarrow \mathbb{Z}_3$ such that every legal coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ either colors all of T by $c(w)$ or all of F by $c(w)$.*

Let $\Psi(X, Y)$ be a $\forall\exists$ -E3-SAT-B instance with m clauses. Recall that X is the set of universal variables and Y is the set of existential variables. The reduction maps it to a graph $G = (V \cup \{w\}, E)$ constructed as follows: We first take one special vertex w . Then, for each variable $x_i \in X$ contained in k positive and ℓ negative literals in Ψ , G contains a copy of $G_{\forall}^{k,\ell}$, where its w vertex is identified with the special vertex w . Similarly, for each variable $y_i \in Y$ contained in k positive and ℓ negative literals in Ψ , G contains a copy of $G_{\exists}^{k,\ell}$, again with its w vertex identified with the special vertex w . In each such gadget, the vertices in T are identified with the positive literals, and the vertices in F are identified with the negative literals. In addition, G contains m clause vertices, one for each clause in Ψ . We connect every clause vertex and the three vertices corresponding to its three literals by edges. We refer to these $3m$ edges as *clause edges*. Denote by $M = |E|$ the number of the edges in G . Notice that M is linear in the total number of variable occurrences, which is $3m$. In particular, $M \leq c \cdot m$ for some absolute constant c .

4.2 Completeness

Let $\Psi(X, Y)$, an m clause formula, be a YES instance. Hence, for any assignment to the universal variables X there exists an assignment to the existential variables Y such that the clauses of Ψ are all satisfied.

Let $\varphi : E \rightarrow \mathbb{Z}_3$ be an arbitrary edge-labelling. Consider the assignment t to X obtained from φ in the following way: According to the first part of Lemma 4.4, the corresponding gadget $G_{\forall}^{k,\ell}$ of every variable $x_i \in X$ satisfies at least one of (1a) and (1b). In the former case, set $t(x_i)$ to be False and in the latter set $t(x_i)$ to be True. For this assignment to X , by assumption, there exists an extension of t to $X \cup Y$ that satisfies Ψ . We now show the existence of a legal vertex coloring $c : V \cup \{w\} \rightarrow \mathbb{Z}_3$ in G :

- Color the special vertex w by $c(w) = 0$.
- For every variable $x_i \in X$, color the set T (in case $t(x_i) = \text{False}$) or the set F (in case $t(x_i) = \text{True}$) of its $G_{\forall}^{k,\ell}$ gadget by the partial coloring given by Lemma 4.4, part (1) (this coloring gives w the color 0).
- For every variable $y_i \in Y$, if its value is True, color the set F of its $G_{\exists}^{k,\ell}$ gadget by the partial coloring from Lemma 4.5, part (1b), and otherwise color the set T of this gadget by the partial coloring from Lemma 4.5, part (1a) (this coloring also gives w the color 0).
- Color any clause vertex v in a way that satisfies the edges adjacent to v . This is possible because for any clause the corresponding vertex is adjacent to at most two vertices that are already colored.
- For every gadget (either $G_{\forall}^{k,\ell}$ or $G_{\exists}^{k,\ell}$) do the following: First, notice that exactly one of T and F is colored. Assume T is colored. Then, color F in a way that the edges between F and w and those between F and clause vertices are satisfied. This is possible since any $v \in F$ is connected to exactly one clause vertex, so there are two constraints on v , yet it has three possible colors. Do a similar thing in the case F is colored.

- Finally, according to the first parts in Lemma 4.4 and Lemma 4.5, the coloring of the gadget vertices can be extended, in a way that satisfies all gadget edges.

To summarize, we have shown that for any $\varphi : E \rightarrow \mathbb{Z}_3$ there exists a coloring, such that the edges are all satisfied. Since this is the case for any $\varphi : E \rightarrow \mathbb{Z}_3$ we conclude that G is \mathbb{Z}_3 -colorable.

4.3 Soundness

Now assume $\Psi(X, Y)$ is a NO instance, i.e., there exists an assignment t to X such that any extension of t to $X \cup Y$ satisfies at most $(1 - \varepsilon)m$ clauses. We show that there exists an edge-labelling $\varphi : E \rightarrow \mathbb{Z}_3$ such that for any coloring $c : V \rightarrow \mathbb{Z}_3$ the fraction of satisfied edges is at most $1 - \frac{\varepsilon}{cB}$. Notice that by defining $\alpha = 1 - \frac{\varepsilon}{cB}$ the theorem will follow.

Define an edge-labelling $\varphi : E \rightarrow \mathbb{Z}_3$ as follows:

- For each clause, let φ give its three corresponding clause edges the three distinct values of \mathbb{Z}_3 .
- For every universal variable x_i , if $t(x_i)$ is True define φ on the gadget $G_{\forall}^{k,\ell}$ according to φ_1 , and otherwise according to φ_2 , where φ_1 and φ_2 are as in the second part of Lemma 4.4.
- For every existential variable y_i , define φ on the gadget $G_{\exists}^{k,\ell}$ according to φ in the second part of Lemma 4.5.

Next, for the j th clause we define C_j as the set of edges ‘related’ to this clause. Namely, C_j consists of the three clause edges corresponding to clause j , together with all edges in the gadgets corresponding to the variables in the clause j . Observe that the union $\bigcup_{j=1}^m C_j$ is the edge set of G . Moreover, every edge of G appears in at most B of these sets.

Now, let $c : V \rightarrow \mathbb{Z}_3$ be an arbitrary vertex coloring. According to our choice of φ , for each legally colored $G_{\forall}^{k,\ell}$ gadget, if $t(x_i) = \text{True}$ then the vertex set F in the corresponding $G_{\forall}^{k,\ell}$ is colored $c(w)$, and otherwise the vertex set T is colored $c(w)$.

Extend t to $X \cup Y$ as follows: For each existential variable $y_i \in Y$, if all edges in the corresponding $G_{\exists}^{k,\ell}$ are satisfied, then by Lemma 4.5 either T or F is colored $c(w)$. In the former case set $t(y_i)$ to be False, and in the latter set it to be True. If the coloring of the gadget $G_{\exists}^{k,\ell}$ of y_i is not legal, define $t(y_i)$ arbitrarily.

Assume all the edges in C_j are satisfied for some $1 \leq j \leq m$. We claim that this implies that the j th clause is satisfied by t . Indeed, if it is not, then the clause vertex is connected to three gadget vertices that are colored with $c(w)$. Since φ assigns to three clause edges the three distinct elements of \mathbb{Z}_3 , one of the edges must be unsatisfied. Hence, we obtain that at least εm of the sets C_j contain at least one unsatisfied edge. Since an edge is contained in at most B sets C_j , we have that at most a $1 - \frac{\varepsilon}{cB}$ fraction of the edges is satisfied, as desired.

Acknowledgement

We thank Marcus Schaefer and Chris Umans for maintaining their excellent compendium of problems in the polynomial-time hierarchy [21] which taught us about group coloring problems, and Daniel Král for clarifications on [14].

References

- [1] M. Ajtai. The shortest vector problem in l_2 is NP-hard for randomized reductions. In *Proc. 30th ACM Symp. on Theory of Computing (STOC)*, pages 10–19, 1998.
- [2] M. Ajtai, R. Kumar, and D. Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *Proc. of 17th IEEE Annual Conference on Computational Complexity (CCC)*, pages 53–57, 2002.
- [3] R. B. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25:127–132, May 1987.
- [4] I. Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. *Theoretical Computer Science*, 285(1):55–71, 2002. Preliminary version in CIAC 2000.
- [5] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003. Preliminary version in FOCS 1998.
- [6] 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 2004.
- [7] I. Haviv. The remote set problem on lattices. In *APPROX-RANDOM*, volume 7408 of *Lecture Notes in Computer Science*, pages 182–193. Springer, 2012.
- [8] I. Haviv, V. Lyubashevsky, and O. Regev. A note on the distribution of the distance from a lattice. *Discrete and Computational Geometry*, 41(1):162–176, 2009.
- [9] I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 469–477, 2007.
- [10] I. Haviv, O. Regev, and A. Ta-Shma. On the hardness of satisfiability with bounded occurrences in the polynomial-time hierarchy. *Theory of Computing*, 3(3):45–60, 2007.
- [11] R. Kannan. Lattice translates of a polytope and the Frobenius problem. *Combinatorica*, 12(2):161–177, 1992.
- [12] S. Khot. Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM*, 52(5):789–808, Sept. 2005. Preliminary version in FOCS 2004.
- [13] S. Khot. Hardness of approximating the shortest vector problem in high ℓ_p norms. *Journal of Computer and System Sciences*, 72(2):206–219, Mar. 2006. Preliminary version in FOCS 2003.
- [14] D. Král and P. Nejedlý. Group coloring and list group coloring are Π_2^P -complete (extended abstract). In *Mathematical foundations of computer science 2004*, volume 3153 of *Lecture Notes in Comput. Sci.*, pages 274–286. Springer, Berlin, 2004.
- [15] A. McLoughlin. The complexity of computing the covering radius of a code. *IEEE Transactions on Information Theory*, 30:800–804, Nov. 1984.

- [16] D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM Journal on Computing*, 34(1):118–169, 2004. Preliminary version in STOC 2002.
- [17] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, MA, 2002.
- [18] D. Micciancio and O. Regev. Lattice-based cryptography. In D. J. Bernstein and J. Buchmann, editors, *Post-quantum Cryptography*. Springer, 2008.
- [19] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 351–358, 2010.
- [20] C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, 2008. Preliminary version in CCC 2007.
- [21] M. Schaefer and C. Umans. Completeness in the Polynomial-Time Hierarchy: A Compendium. *SIGACT News*, Sept. 2002.
- [22] M. Schaefer and C. Umans. Completeness in the Polynomial-Time Hierarchy: Part II. *SIGACT News*, Dec. 2002.
- [23] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, Math Inst., University Of Amsterdam, Amsterdam, 1981.