Hardness of the Binary Covering Radius Problem in Large ℓ_p Norms

Huck Bennett
University of Colorado Boulder, CO, USA

Peter Ly
University of Colorado Boulder, CO, USA

Abstract

We study the hardness of the γ-approximate decisional Covering Radius Problem on lattices in the ℓ_p norm (γ-GapCRP$_p$). Specifically, we prove that there is an explicit function $\gamma(p)$, with $\gamma(p) > 1$ for $p > p_0 \approx 35.31$ and $\lim_{p\to\infty} \gamma(p) = 9/8$, such that for any constant $\varepsilon > 0$, $(\gamma(p) - \varepsilon)$-GapCRP$_p$ is NP-hard. This shows the first hardness of GapCRP$_p$ for explicit $p < \infty$. Work of Haviv and Regev (CCC, 2006 and CJTCS, 2012) previously showed Π_2-hardness of approximation for GapCRP$_p$ for all sufficiently large (but non-explicit) finite p and for $p = \infty$.

In fact, our hardness results hold for a variant of GapCRP called the *Binary Covering Radius Problem* (BinGapCRP). The Binary Covering Radius Problem trivially reduces to both GapCRP and the decisional Linear Discrepancy Problem (LinDisc) in any norm in an approximation-preserving way. We also show Π_2-hardness of $(9/8 - \varepsilon)$-BinGapCRP in the ℓ_∞ norm for any constant $\varepsilon > 0$.

Our work extends and heavily uses the work of Manurangsi (IPL, 2021), which showed Π_2-hardness of $(9/8 - \varepsilon)$-LinDisc in the ℓ_∞ norm.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness

Keywords and phrases Covering radius problem, linear discrepancy, hardness of approximation

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.10

Category APPROX

Related Version *Full Version*: https://arxiv.org/abs/2603.03219

Funding Supported by NSF Award No. 2432132.

Acknowledgements We thank Ishay Haviv, Pasin Manurangsi, Daniele Micciancio, and Noah Stephens-Davidowitz for helpful comments and references.

1 Introduction

The *covering radius* $\mu(\mathcal{L})$ of a lattice $\mathcal{L}$ is the maximum distance of a point in the span of $\mathcal{L}$ to a vector in $\mathcal{L}$. It is one of the most fundamental lattice quantities. Accordingly, the (decisional) Covering Radius Problem ((Gap)CRP) is among the most fundamental lattice problems. Indeed, it is connected to the main applications of lattices in computer science. Approximate GapCRP is one of the worst-case problems from which there is a reduction to the average-case problems used in lattice-based cryptography [32], and it is related to fast algorithms for integer programming (see, e.g., [8]).

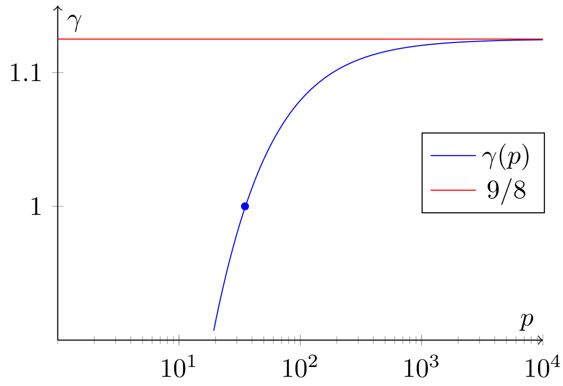

Figure 1 A linear-log plot of the approximation factor $\gamma = \gamma(p)$ for which we show NP-hardness of $(\gamma - \varepsilon)$-BinGapCRP in the ℓ_p norm (in blue; see Theorem 1). The function $\gamma(p)$, specified in Equation (1), is greater than 1 for all $p > p_0 = 35.310188...$ and satisfies $\lim_{p\to\infty} \gamma(p) = 9/8$. (The blue dot is at the point $(p, \gamma) = (p_0, 1)$, where p_0 is the unique value of p such that $\gamma(p) = 1$. The horizontal asymptote of $9/8$ is shown in red.) This factor of $9/8$ also appears in our Π_2-hardness result for $\mathrm{BinGapCRP}_\infty$ (Theorem 2) and in the work of Manurangsi [29] on LinDisc, which is the foundation of our work..

Yet, the hardness of GapCRP is perhaps the least well understood among major lattice problems. In particular, it is not even known that *exact* GapCRP in the ℓ_2 norm is NP-hard, even though it is plausibly NP-hard to approximate within any constant $\gamma \geq 1$ and Π_2-hard to approximate to within a small constant $\gamma > 1$. This stands in contrast to, e.g., the decisional Shortest Vector Problem (γ-GapSVP), which a long line of elegant work [2, 31, 25, 24, 36, 21, 33] showed to be NP-hard in any ℓ_p norm for any constant $\gamma \geq 1$.[1]

Seminal work of Haviv and Regev [20] was the first – and prior to this work, only – to study the hardness of GapCRP. (Other work has studied the complexity of GapCRP more generally, including limitations on its hardness [15, 28]. See Section 1.2.) Let γ-GapCRP_p be γ-approximate GapCRP in the ℓ_p norm. Haviv and Regev [20] showed that for any constant $\varepsilon > 0$, $(3/2 - \varepsilon)$-GapCRP_∞ is Π_2-hard, and that for all sufficiently large p there exists $\gamma = \gamma(p) > 1$ such that γ-GapCRP_p is Π_2-hard. However, the cutoff for "sufficiently large" in their work is non-explicit and presumably quite high, and so in some sense not much more is known about the hardness of GapCRP_p for any explicit, finite p than for $p = 2$.

1.1 Our Results

In this work, we address the problem above by showing NP-hardness of $(\gamma - \varepsilon)$-GapCRP_p for all p greater than an *explicit* value $p_0 \approx 35.31$ with an explicit approximation factor $\gamma = \gamma(p) > 1$. This is the first new hardness result for GapCRP since [20], which was first published 20 years ago.

In fact, we show hardness for a new promise variant of γ-GapCRP_p called γ-$\mathrm{BinGapCRP}_p$ where in the YES case every vector in the fundamental parallelepiped $\mathcal{P}(\mathbf{B}) := \mathbf{B} \cdot [0, 1]^n$ of the input basis $\mathbf{B} \in \mathbb{R}^{m\times n}$ is close in ℓ_p distance to a lattice vector $\mathbf{B}\boldsymbol{x}$ with a *binary* coefficient vector $\boldsymbol{x} \in \{0, 1\}^n$. That is, in the YES case, for every $\boldsymbol{w} \in [0, 1]^n$ there exists $\boldsymbol{x} \in \{0, 1\}^n$ such that $\|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p$ is small. This implies that the covering radius $\mu_p(\mathcal{L}(\mathbf{B}))$

[1] Formally, this NP-hardness is only known under *randomized reductions*, although several works – some of which are very recent [16, 22, 17, 38]! – have derandomized the reductions in some cases.

is small. (See Equation (4) and the surrounding text for a formal definition of the covering radius and a discussion of why this is true.) In the NO case, there exists a vector in $\mathcal{P}(\mathbf{B})$ that is far from *every* lattice vector $\mathbf{B}\boldsymbol{x}$ for $\boldsymbol{x} \in \mathbb{Z}^n$, as with standard γ-GapCRP$_p$.

BinGapCRP is reminiscent of the "binary coefficient vector" variant of the Closest Vector Problem often used in hardness reductions (see, e.g., [13, Definition 7] and [3, Definition 3.1]), and may be of independent interest. It has the YES instances of the Linear Discrepancy Problem (LinDisc) and the NO instances of GapCRP. It therefore trivially reduces to both problems in any norm in an approximation-preserving way (see Section 2.2 for definitions of these problems). In particular, our result not only implies new hardness for GapCRP$_p$ with $p < \infty$, but also for LinDisc in the ℓ_p norm for $p < \infty$ (LinDisc in finite ℓ_p norms does not seem to have been studied before).

Our main result is the following theorem. See also Figure 1.

▶ **Theorem 1** (NP-hardness of approximation for BinGapCRP$_p$). *Let*

$$\gamma(p) := \left(\frac{9^p + 159 \cdot 3^p}{8^{p+2} + 96 \cdot 6^p}\right)^{1/p}, \tag{1}$$

and let $p_0 = 35.310188\ldots$ be the unique value of $p \geq 1$ such that $\gamma(p) = 1$. Then for all $p > p_0$ and any constant $\varepsilon > 0$, $(\gamma(p) - \varepsilon)$-BinGapCRP$_p$ is NP*-hard. Therefore, $(\gamma(p) - \varepsilon)$-GapCRP$_p$ and $(\gamma(p) - \varepsilon)$-LinDisc$_p$ for such p and ε are also* NP*-hard.*

We note that the expression $\gamma(p)$ in Equation (1) satisfies $\lim_{p\to\infty} \gamma(p) = 9/8$. Our second result shows how to achieve Π_2-hardness (rather than just NP-hardness) of $(9/8 - \varepsilon)$-BinGapCRP in the ℓ_∞ norm.

▶ **Theorem 2** (Π_2-hardness of approximation for BinGapCRP$_\infty$). *For any constant $\varepsilon > 0$, $(9/8 - \varepsilon)$-BinGapCRP$_\infty$ is Π_2-hard.*

Theorem 2 is a strengthening of the main result in [29] – it shows hardness not just of $(9/8-\varepsilon)$-LinDisc$_\infty$ but of $(9/8-\varepsilon)$-BinGapCRP$_\infty$, where LinDisc$_\infty$ is the Linear Discrepancy Problem in the ℓ_∞ norm. Theorem 2 also gives an alternative proof of Π_2-hardness of approximation for γ-GapCRP$_\infty$, albeit with a smaller approximation factor than the one achieved by [20].

1.2 Related Work

Early work of McLoughlin [30] showed that the Covering Radius Problem on error-correcting codes is Π_2-hard by a reduction from a Π_2-hard variant of the 3-Dimensional Matching Problem, which is in turn shown to be Π_2-hard by a reduction from Π_2-3-SAT.[2] Twenty years later, work of Guruswami, Micciancio, and Regev [15] studied the complexity of the Covering Radius Problem both on codes and lattices. It showed a number of results, including that the Covering Radius Problem on codes is NP-hard to approximate to within any constant and Π_2-hard to approximate to within some constant. (It did not show hardness results for GapCRP on lattices.)

The work [15] also showed that 2-GapCRP$_p$ for any p is in AM via an elegant protocol. The fact that 2-GapCRP$_p$ is in AM provides evidence that γ-GapCRP$_p$ for $\gamma \geq 2$ is unlikely to be Π_2-hard, since if it were the polynomial hierarchy would collapse. However, being in AM says nothing about the potential NP-hardness of these problems. In particular it could very well be the case that γ-GapCRP is NP-hard for every constant $\gamma \geq 1$ as is the case for the Covering Radius Problem on codes.

[2] Here and throughout this section we refer to linear error-correcting codes simply as "codes."

The factor of $\gamma = 2$ in the [15] AM protocol follows generically for GapCRP in any norm from the triangle inequality, but it is not always tight. Follow-up work by Haviv, Lyubashevsky, and Regev [19] studied the distribution of the (normalized) ℓ_p distance of a random point in a fundamental domain of a lattice to the lattice. This distribution is closely related to the [15] AM protocol, and [19] showed that γ-GapCRP$_2$ for $\gamma > \sqrt{3} \approx 1.732$ is in AM under a certain conjecture about this distribution in the ℓ_2 norm. Their conjecture was subsequently proved by Magazinov [28].

Understanding the worst-case hardness of γ-GapCRP is also motivated by its relevance to cryptography. Indeed, [32] showed that γ-GapCRP with sufficiently large polynomial $\gamma = \gamma(n)$ can be taken as the starting worst-case problem for a worst-case to average-case reduction to the problems used as the foundation of lattice-based cryptography.

The Haviv-Regev hardness reduction

The work of Haviv and Regev [20] is the only prior work that studied the worst-case hardness of GapCRP on lattices. As noted above, they showed that for every constant $\varepsilon > 0$, $(3/2-\varepsilon)$-GapCRP$_\infty$ is Π_2-hard and that for all sufficiently large p there exists $\gamma = \gamma(p)$ such that γ-GapCRP$_p$ is Π_2-hard to approximate. However, the cutoff for "sufficiently large" and the approximation factors $\gamma(p)$ for $p < \infty$ are non-explicit.

The reduction in [20] is from the Group Coloring Problem on graphs with respect to the group $\mathbb{Z}_3$, a problem that is Π_2-hard.[3] The problem is defined as follows. The input is a directed graph $G = (V, E)$. The goal is to decide whether for every assignment $\phi : E \to \mathbb{Z}_3$ of labels to edges, there exists a coloring $c : V \to \mathbb{Z}_3$ of vertices such that for every edge $e = (u, v) \in E$, $c(u) - c(v) \neq \phi(e)$. The reduction to GapCRP is as follows. Let $A = A(G) \in \mathbb{Z}^{m \times n}$ for $m := |E|$, $n := |V|$ be the incidence matrix of G. That is, A is the matrix with two non-zero entries per row defined as

$$A_{i,j} := \begin{cases} -1 & \text{if edge } e_i = (v_j, u) \text{ for some } u \in V, \\ 1 & \text{if edge } e_i = (u, v_j) \text{ for some } u \in V, \\ 0 & \text{otherwise.} \end{cases}$$

The reduction then outputs (a basis of) the lattice $\mathcal{L} := A \cdot \mathbb{Z}^n + 3\mathbb{Z}^m$. One may also view A as the generator matrix of a ternary code $\mathcal{C} \subseteq \mathbb{F}_3^m$, and $\mathcal{L}$ as the Construction-A lattice $\mathcal{C} + 3\mathbb{Z}^m$.

To show NP-hardness of γ-GapCRP$_p$ in some ℓ_p norm (for large γ), one might try to reduce from variants of the 3-Coloring Problem, which is the NP analog of the Π_2-complete $\mathbb{Z}_3$-Group Coloring Problem. Indeed, strong hardness of approximation results are known for Almost-3-Coloring – the optimization variant of 3-Coloring in which the goal is to minimize the number of monochromatic edges – assuming variants of the Unique Games Conjecture [9]. However, it is unclear how to get this approach to work.

Combinatorial and linear discrepancy

The (combinatorial) *discrepancy* of a set system $S_1, \ldots, S_m \subseteq \{x_1, \ldots, x_n\}$ is defined as $\min_{x_1,\ldots,x_n \in \{-1,1\}} \max_{i \in [m]} |\sum_{j \in S_i} x_j|$, and the goal of the *combinatorial discrepancy problem* is, given the set system as input, to find $x_1, \ldots, x_n \in \{-1, 1\}$ that achieve this minimum.

[3] Their reduction actually works with $\mathbb{Z}_q$ for any $q \geq 3$, but they achieve the best hardness of approximation with $q = 3$.

Equivalently, if $\mathbf{A} \in \{0,1\}^{m \times n}$ is the incidence matrix of the input set system, the goal is to find $\boldsymbol{x} \in \{-1,1\}^n$ that minimizes $\|\mathbf{A}\boldsymbol{x}\|_\infty$. Combinatorial discrepancy and other variants of discrepancy are used widely in computer science and math (see, e.g., [7]).

Linear discrepancy was originally defined and related to other notions of discrepancy in work of Lovász, Spencer, and Vestergombi [27], and it appears as a key ingredient in certain approximation algorithms (see, e.g., [11, 37, 23]). Formally, we define the linear discrepancy $\mathrm{lindisc}(\mathbf{B})$ of a matrix $\mathbf{B}$ in the ℓ_p norm as

$$\mathrm{lindisc}_p(\mathbf{B}) := \max_{\boldsymbol{w} \in [0,1]^n} \min_{\boldsymbol{x} \in \{0,1\}^n} \|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p \ . \tag{2}$$

This definition is syntactically very similar to the definition of the covering radius in the ℓ_p norm, μ_p. The difference is that in the definition of μ_p, the minimum is taken over $\boldsymbol{x} \in \mathbb{Z}^n$ rather than $\boldsymbol{x} \in \{0,1\}^n$. See Sections 2.1 and 2.2.

Combinatorial discrepancy is captured, up to translation and scaling, by fixing $\boldsymbol{w} := \frac{1}{2} \cdot \mathbf{1}$ in Equation (2) instead of taking the max over all $\boldsymbol{w} \in [0,1]^n$ there. Linear discrepancy was first investigated from a complexity perspective by Li and Nikolov [26], who provided both algorithms and hardness results. In particular, they showed that the exact Linear Discrepancy Problem in the ℓ_∞ norm, $\mathrm{LinDisc}_\infty$, is NP-hard and that it is contained in Π_2. Their hardness result was subsequently improved by Manurangsi [29], who showed that γ-$\mathrm{LinDisc}_\infty$ is Π_2-hard for $\gamma := 9/8 - \varepsilon$ for any constant $\varepsilon > 0$. Our work makes extensive use of [29].

"Standard" linear discrepancy corresponds to $\mathrm{lindisc}_\infty$, and, to the best of our knowledge, our work is the first to extend the definition to finite ℓ_p norms. This extension is useful for our work. Moreover, we note that although linear discrepancy does not seem to have been studied in finite ℓ_p norms before, both combinatorial discrepancy and hereditary discrepancy have (see, e.g., [34]).

Lattice hardness results in large ℓ_p norms

We recall that our main hardness result in Theorem 1 only holds for GapCRP_p with $p > p_0 \approx 35.31$, whereas likely the most interesting case is $p = 2$. We note that it has been the case several times that progress on showing hardness of lattice problems in large ℓ_p norms has precipitated progress for the Euclidean case of $p = 2$. In particular, Khot [25] showed randomized NP-hardness of approximation for the Shortest Vector Problem (SVP) in large ℓ_p norms, and then shortly thereafter showed such hardness in all ℓ_p norms with $p > 1$ in [24]. More recently, Hair and Sahai [16] showed *deterministic* NP-hardness of approximation for SVP in ℓ_p norms for $p > 2$, and then showed hardness of approximation for SVP in the Euclidean norm [17].

More generally, quite a few works on the hardness of lattice problems have shown their strongest results in large ℓ_p norms (see, e.g., [1, 4, 5]).

1.3 Technical Overview

At a high level, our hardness reduction for $\mathrm{BinGapCRP}_p$ follows the hardness reduction for $\mathrm{LinDisc}_\infty$ in [29]. The reduction in [29] reduces from NAE-E3-SAT,[4] the constraint

[4] In fact, [29] reduces from the NAE-3-SAT problem (rather than the NAE-E3-SAT problem), in which constraints may have two literals (rather than always having three). The reduction simply doubles an arbitrary literal in any constraint with two literals, which leads to the "incidence" matrix $\mathbf{B}$ defined below potentially containing entries from $\{-2,2\}$ as well as $\{-1,0,1\}$. In this work, we reduce from NAE-E3-SAT.

satisfaction problem in which each constraint is a function of three literals (Boolean variables v_j or their negations $\neg v_j$), and a constraint is satisfied by an assignment to its variables if and only if the assignment neither satisfies none of its literals nor all of its literals.

Specifically, on input a NAE-E3-SAT formula $\phi(v_1, \ldots, v_n)$ with n Boolean variables and with m constraints, the NP-hardness reduction in [29] first computes the constraint-variable incidence matrix $\mathbf{B} = \mathbf{B}(\phi) \in \mathbb{Z}^{m \times n}$ with entries

$$B_{i,j} := \begin{cases} 1 & \text{if constraint } i \text{ contains } v_j \ , \\ -1 & \text{if constraint } i \text{ contains } \neg v_j \ , \\ 0 & \text{otherwise} \ . \end{cases}$$

It then outputs a LinDisc_∞ instance with matrix

$$\mathbf{A} := \begin{pmatrix} \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} \\ & \mathbf{G} \otimes \mathbf{I}_n & \end{pmatrix} \in \mathbb{Q}^{(m+3n) \times 3n} \ , \tag{3}$$

where $\mathbf{G}$ is the 3×3 gadget matrix

$$\mathbf{G} := \begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{pmatrix} \ ,$$

$\otimes$ is the Kronecker product and $\mathbf{I}_n$ is the $n \times n$ identity matrix.

Both the completeness and soundness analysis of the reduction in [29] crucially use properties of $\mathbf{G}$. For example, [29, Lemma 3] argues that $\mathbf{G}$ has low discrepancy by upper bounding the distance between $\mathbf{G}\boldsymbol{u}$ and $\mathbf{G} \cdot \{0,1\}^3$ for each $\boldsymbol{u} \in [0,1]^3$ via a case analysis on the ℓ_1 norm of $\boldsymbol{u}$. The soundness analysis shows that, if there exists a close vector in the set $\mathbf{A} \cdot \{0,1\}^{3n}$ to $\mathbf{A} \cdot (\frac{1}{2}\mathbf{1})$, then ϕ must be satisfiable. [29] also shows Π_2-hardness of $(9/8 - \varepsilon)$-LinDisc_∞ by outputting and analyzing a matrix $\mathbf{A}' = \mathbf{A}'(\phi)$ defined in terms of ϕ that is related to the matrix $\mathbf{A}$ defined in Equation (3) but more complicated. We use this same matrix $\mathbf{A}'$ (defined in Equation (27)) to prove Theorem 2.

At a technical level, our NP-hardness result adapts the machinery in [29] from ℓ_∞ to general ℓ_p norms, and shows that, if the NAE-E3-SAT formula ϕ is unsatisfiable, then not only is $\text{lindisc}(\mathbf{A})$ large, in fact $\mu(\mathcal{L}(\mathbf{A}))$ is large. That is, we show that if ϕ is unsatisfiable then there exists $\boldsymbol{w} \in [0,1]^{3n}$ (in fact, $\boldsymbol{w} := \frac{1}{2} \cdot \mathbf{1}$) such that $\|\mathbf{A}(\boldsymbol{w} - \boldsymbol{x})\|_p$ is large not only for all $\boldsymbol{x} \in \{0,1\}^{3n}$ but for all $\boldsymbol{x} \in \mathbb{Z}^{3n}$. Doing these things requires making a number of modifications to the hardness reduction and analysis used by [29], such as:

1. Using *hardness of approximation* for NAE-E3-SAT with soundness $15/16 + \varepsilon$ and perfect completeness; see Theorem 10.
2. Taking the Kronecker product of G with the "ℓ_p norm degree" diagonal matrix

   $$\mathbf{D}_p := \mathbf{D}_p(\phi) = \text{diag}(\deg(v_1)^{1/p}, \ldots, \deg(v_n)^{1/p}) \in \mathbb{R}^{n \times n}$$

   in the definition of $\mathbf{A}$ instead of just taking the Kronecker product with I_n. Here $\deg(v_i)$ counts the number of occurrences of the variable v_i and its negation $\neg v_i$ in the formula ϕ.
3. Extending the analysis from [29] for the gadget matrix $\mathbf{G}$. For example, we prove an ℓ_p analog of the fact that $\text{lindisc}(\mathbf{G})_\infty$ is small (and in fact, that something stronger is true) in Lemma 11, and we observe that $\mathbf{G} \cdot \mathbf{0} = \mathbf{0}$ and $\mathbf{G} \cdot \mathbf{1} = \mathbf{1}$ are the only closest vectors in $\mathcal{L}(\mathbf{G}) = \mathbf{G} \cdot \mathbb{Z}^3$ to $\mathbf{G} \cdot (\frac{1}{2}\mathbf{1})$ (and not just the only closest vectors in $\mathbf{G} \cdot \{0,1\}^3$ to $\mathbf{G} \cdot (\frac{1}{2}\mathbf{1})$) in Lemma 12.

For Item 1, we note that the soundness-to-completeness ratio of $15/16$ in Theorem 10 is *not* tight. In fact, recent work of Brakensiek, Huang, Potechin, and Zwick [6] shows tight upper and lower bounds for NAE-E3-SAT with a lower (i.e., better) soundness-to-completeness ratio of roughly $0.9089 < 15/16 = 0.9375$. However, Theorem 10 shows hardness of (δ, ε)-NAE-E3-SAT *with perfect completeness* (i.e., with $\varepsilon = 1$), whereas the result of [6] does not. This is important since our main hardness reduction for $\mathrm{BinGapCRP}_p$ requires reducing from (δ, ε)-NAE-E3-SAT with nearly perfect completeness. The hardness result in Theorem 10 appeared in work of Guruswami [14], which observed that it follows by combining a reduction due to Petrank [35] and a hardness result due to Håstad [18]. (We include a proof of this result in the full version.)

The reason for using $\mathbf{D}_p$ as in Item 2 is as follows. In our analysis, we analyze the quantity $\|\mathbf{A}\boldsymbol{y}\|_p^p$ for coefficient vectors $\boldsymbol{y} \in \mathbb{Z}^{3n}$ expressed as a sum indexed by the m constraints in ϕ (see Lemma 13). We wish to penalize vectors $\boldsymbol{y}$ that do not correspond to valid assignments to ϕ by having $\|\mathbf{A}\boldsymbol{y}\|_p^p$ be larger for such $\boldsymbol{y}$. Using the matrix $\mathbf{D}_p$ allows us to do this. Specifically, it allows us to balance the contributions of the top part $\left(\frac{1}{3}\mathbf{B} \quad \frac{1}{3}\mathbf{B} \quad \frac{1}{3}\mathbf{B}\right) \cdot \boldsymbol{y}$ and the bottom part $(\mathbf{G} \otimes \mathbf{D}_p) \cdot \boldsymbol{y}$ of $\mathbf{A}\boldsymbol{y}$ to each term in this sum. The top part's contribution to a summand will be large if $\boldsymbol{y}$ corresponds to a valid but unsatisfying assignment to ϕ, and the bottom part's contribution to a summand will be large if $\boldsymbol{y}$ corresponds to an invalid assignment to ϕ.

1.4 Open Questions

We again emphasize the central open question that we consider in this work: For which values of $\gamma = \gamma(p)$ is γ-GapCRP in the ℓ_p norm NP-hard, and for which is it Π_2-hard? In particular, we reiterate that the problem of showing NP-hardness of exact GapCRP in the ℓ_2 norm (the most important special case) remains open. We again note that the AM protocol of [15] for approximate GapCRP (conditionally) rules out Π_2-hardness for γ-GapCRP_p for $\gamma \geq 2$, but says nothing about NP-hardness of γ-GapCRP_p. Manurangsi [29] notes that no such analogous AM protocol is known for LinDisc, and asks if one exists.

2 Preliminaries

We will use boldface, lowercase letters to denote vectors. We will sometimes abuse notation when defining column vectors in terms of other column vectors by writing things like $\boldsymbol{x} := (\boldsymbol{y}, \boldsymbol{z})$ to mean $\boldsymbol{x} := (\boldsymbol{y}^T, \boldsymbol{z}^T)^T$. We will use $\mathbf{0}$ and $\mathbf{1}$ to denote the all-0s and all-1s vectors, respectively.

We define the predicate $\mathrm{sgndiff}(\boldsymbol{x}, \boldsymbol{y})$ on vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ to be true if and only if the Hadamard (entry-wise) product $\boldsymbol{x} \odot \boldsymbol{y} \in \mathbb{R}^n$ has at least one non-positive coordinate and at least one non-negative coordinate. We will use the simple fact that if $\mathrm{sgndiff}(\boldsymbol{x}, \boldsymbol{y})$ holds and if $\max_i |x_i y_i| \leq r$ for some $r \geq 0$, then $|\langle \boldsymbol{x}, \boldsymbol{y} \rangle| = |\sum_{i=1}^n x_i y_i| \leq r(n-1)$.

2.1 Covering Radius and Linear Discrepancy

We call a matrix $\mathbf{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n) \in \mathbb{R}^{m \times n}$ with full column rank a *basis*. In particular, $m \geq n$ for a basis. The lattice $\mathcal{L}(\mathbf{B})$ generated by $\mathbf{B}$ is defined as

$$\mathcal{L}(\mathbf{B}) := \Big\{ \sum_{i=1}^{n} a_i \boldsymbol{b}_i : a_1, \ldots, a_n \in \mathbb{Z} \Big\} .$$

For $p \geq 1$, the ℓ_p *covering radius* of a lattice $\mathcal{L} \subset \mathbb{R}^m$ is defined as

$$\mu_p(\mathcal{L}) := \max_{\boldsymbol{t} \in \mathrm{span}(\mathcal{L})} \mathrm{dist}_p(\boldsymbol{t}, \mathcal{L}) \ ,$$

where for a discrete set S, we define $\mathrm{dist}_p(\boldsymbol{t}, S) := \min_{\boldsymbol{y} \in S} \|\boldsymbol{t} - \boldsymbol{y}\|_p$. The *span* $\mathrm{span}(\mathcal{L})$ of a lattice $\mathcal{L}$ is the linear span of the vectors in the lattice (or equivalently, the linear span $\mathrm{span}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ of the vectors in a basis $\mathbf{B} := (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ of $\mathcal{L}$). Because $\mathrm{dist}_p(\boldsymbol{t} + \boldsymbol{y}, \mathcal{L}) = \mathrm{dist}_p(\boldsymbol{t}, \mathcal{L})$ for any $\boldsymbol{t} \in \mathbb{R}^m$ and $\boldsymbol{y} \in \mathcal{L}$, we have that if $\mathcal{L} = \mathcal{L}(\mathbf{B})$ for a basis $\mathbf{B} \in \mathbb{R}^{m \times n}$,

$$\mu_p(\mathcal{L}) = \max_{\boldsymbol{w} \in \mathbb{R}^n} \min_{\boldsymbol{x} \in \mathbb{Z}^n} \|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p = \max_{\boldsymbol{w} \in [0,1]^n} \min_{\boldsymbol{x} \in \mathbb{Z}^n} \|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p \ . \tag{4}$$

That is, because the distance of a vector $\mathbf{B}\boldsymbol{w}$ to $\mathcal{L}(\mathbf{B})$ is invariant under shifts by lattice vectors (i.e., vectors $\mathbf{B}\boldsymbol{x}$ for $\boldsymbol{x} \in \mathbb{Z}^n$), it suffices to consider the maximum distance of vectors $\mathbf{B}\boldsymbol{w}$ in the fundamental parallelepiped $\mathcal{P}(\mathbf{B}) := \mathbf{B} \cdot [0,1]^n$ to the lattice.

We will also work with *linear discrepancy*, a quantity closely related to the covering radius that was first defined in [27]. For $p \geq 1$, the linear discrepancy $\mathrm{lindisc}(\mathbf{B})$ of a matrix $\mathbf{B}$ is defined as[5]

$$\mathrm{lindisc}_p(\mathbf{B}) := \max_{\boldsymbol{w} \in [0,1]^n} \min_{\boldsymbol{x} \in \{0,1\}^n} \|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p \ . \tag{5}$$

Covering radius versus linear discrepancy

We note that the equivalent formulation of $\mu_p(\mathcal{L})$ on the right-hand side of Equation (4) and the definition of $\mathrm{lindisc}(\mathbf{B})$ are syntactically very similar. The only difference is that the minimum in Equation (4) is taken over $\boldsymbol{x} \in \mathbb{Z}^n$ (which corresponds to the entire lattice $\mathcal{L}(\mathbf{B}) = \{\mathbf{B}\boldsymbol{x} : \boldsymbol{x} \in \mathbb{Z}^n\}$) whereas in Equation (5) the minimum is only taken over $\boldsymbol{x} \in \{0,1\}^n$ (which corresponds to *binary* linear combinations $\mathbf{B}\boldsymbol{x}$ of basis vectors).

Despite the similarity of the definitions of μ_p and $\mathrm{lindisc}_p$, we emphasize that they capture different quantities. In particular, the closest vector in the lattice $\mathcal{L}(\mathbf{B}) := \mathbf{B} \cdot \mathbb{Z}^n$ to a vector $\mathbf{B} \cdot \boldsymbol{w}$ for $\boldsymbol{w} \in [0,1]^n$ need not be a vector in $\mathbf{B} \cdot \{0,1\}^n$. We give a simple example of this for the lattice $\mathbb{Z}^2$ when the input basis $\mathbf{B}$ consists of rather long vectors.

▶ **Example 3.** Consider the basis

$$\mathbf{B} := \begin{pmatrix} 3 & 4 \\ 1 & 1 \end{pmatrix} \ ,$$

and note that $\mathcal{L}(\mathbf{B}) = \mathbb{Z}^2$ since $\det(\mathbf{B}) = -1$ (i.e., $\mathbf{B}$ is unimodular). Let $\boldsymbol{w} := (1, 1/2)^T$, and note that $\mathbf{B}\boldsymbol{w} = (5, 3/2)^T$. Then

$$\mathrm{dist}_p(\mathbf{B}\boldsymbol{w}, \mathcal{L}(\mathbf{B})) = \mathrm{dist}_p(\mathbf{B}\boldsymbol{w}, \mathbb{Z}^2) = \|\mathbf{B}((3,-1)^T - \boldsymbol{w})\|_p = 1/2 \ ,$$

but one can check that $\|\mathbf{B}(\boldsymbol{x} - \boldsymbol{w})\|_p > 1/2$ for all $\boldsymbol{x} \in \{0,1\}^2$.

2.2 Covering Radius and Linear Discrepancy as Computational Problems

We next define the approximate decisional Covering Radius Problem in the ℓ_p norm, γ-GapCRP, and the closely related approximate Linear Discrepancy Problem in the ℓ_p norm, γ-$\mathrm{LinDisc}_p$.

[5] The definition of $\mathrm{lindisc}(\mathbf{B})$ does not require $\mathbf{B}$ to be a basis. I.e., $\mathbf{B}$ is allowed to have linearly dependent columns.

▶ **Definition 4.** *Let $p \in [1, \infty]$ and let $\gamma = \gamma(n) \geq 1$. The γ-*approximate decisional Covering Radius Problem in the ℓ_p norm *(γ-GapCRP_p), is the problem where, on input a basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ and a value $r > 0$, the goal is to distinguish between the following two cases.*
- *(YES instance.)* $\mu_p(\mathcal{L}(\mathbf{B})) \leq r$.
- *(NO instance.)* $\mu_p(\mathcal{L}(\mathbf{B})) > \gamma r$.

▶ **Definition 5.** *Let $p \in [1, \infty]$ and let $\gamma = \gamma(n) \geq 1$. The γ-*approximate decisional Linear Disrepancy Problem in the ℓ_p norm *(γ-$\mathrm{LinDisc}_p$) is the problem where, on input a basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ and a value $r > 0$, the goal is to distinguish between the following two cases.*
- *(YES instance.)* $\mathrm{lindisc}_p(\mathbf{B}) \leq r$.
- *(NO instance.)* $\mathrm{lindisc}_p(\mathbf{B}) > \gamma r$.

Finally, we define the binary decisional Covering Radius Problem BinGapCRP, which has the YES instances of LinDisc and the NO instances of GapCRP.

▶ **Definition 6.** *Let $p \in [1, \infty]$ and let $\gamma = \gamma(n) \geq 1$. The γ-*approximate decisional Binary Covering Radius Problem in the ℓ_p norm *(γ-$\mathrm{BinGapCRP}_p$) is the problem where, on input a basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ and a value $r > 0$, the goal is to distinguish between the following two cases.*
- *(YES instance.)* $\mathrm{lindisc}_p(\mathbf{B}) \leq r$.
- *(NO instance.)* $\mu_p(\mathcal{L}(\mathbf{B})) > \gamma r$.

We note that $\mu_p(\mathcal{L}(\mathbf{B})) \leq \mathrm{lindisc}_p(\mathbf{B})$ for any $p \in [1, \infty]$ and any basis $\mathbf{B}$. Therefore, for any $p \in [1, \infty]$ and $\gamma \geq 1$, YES (respectively, NO) instances of γ-$\mathrm{BinGapCRP}_p$ are YES (respectively, NO) instances of both γ-GapCRP_p and γ-$\mathrm{LinDisc}_p$. We therefore immediately have the following observation.

▶ **Proposition 7.** *For any $\gamma \geq 1$ and $p \in [1, \infty]$, the identity mapping $(\mathbf{B}, r) \mapsto (\mathbf{B}, r)$ is a (polynomial-time) reduction from γ-$\mathrm{BinGapCRP}_p$ to γ-GapCRP_p and to γ-$\mathrm{LinDisc}_p$.*

2.3 Constraint Satisfaction Problems

A (Boolean) *constraint satisfaction problem* (CSP) is defined by a set of Boolean variables $X = \{x_1, \ldots, x_n\}$, and a set of functions called *constraints* $\mathcal{C} = \{C_1, \ldots, C_m\}$ with $C_i \colon \{0,1\}^{|C_i|} \to \{0,1\}$. An *assignment* of the variables is a map from $X \to \{0,1\}$. A constraint C_i is *satisfied* by an assignment if the output of the constraint on the assignment of the variables is 1. Let ϕ be a CSP and let ψ be an assignment of the variables in ϕ. We define $\mathrm{val}_\psi(\phi)$ to be the fraction of constraints in ϕ satisfied by an assignment ψ, and define the *value* of ϕ to be $\mathrm{val}(\phi) := \max_\psi \mathrm{val}_\psi(\phi)$. We then define the computational problem (δ, ε)-CSP as follows.

▶ **Definition 8** (Constraint Satisfaction Problems (CSP))**.** *For $\delta, \varepsilon \in \mathbb{R}$ be such that $0 < \delta \leq \varepsilon \leq 1$, (δ, ε)-CSP is the decision problem defined as follows. On input a collection of constraints $\mathcal{C} = \{C_1, \ldots, C_m\}$, the goal is to decide between the follows two cases.*
- *(YES instance.)* $\mathrm{val}(\phi) \geq \varepsilon$.
- *(NO instance.)* $\mathrm{val}(\phi) < \delta$.

We now introduce NAE-Ek-SAT, a k-ary CSP that will be important for our hardness reductions.

▶ **Definition 9** ((δ, ε)-NAE-Ek-SAT problem)**.** *A* NAE-Ek-SAT *formula ϕ consists of n Boolean variables $X = \{x_1, \ldots, x_n\}$ and m constraints $C_1, \ldots, C_m$, where each constraint C_i is a function of exactly k distinct literals (Boolean variables x_j or their negations $\neg x_j$). A constraint C_i is satisfied by an assignment if and only if not all k of its literals evaluate either to 1 or to 0.*

Let $\delta, \varepsilon \in \mathbb{R}$ be such that $0 < \delta \leq \varepsilon \leq 1$. The goal of the (δ, ε)-NAE-Ek-SAT *problem is, given a* NAE-Ek-SAT *formula as input, to decide between the following two cases.*

- *(YES instance.)* $\mathrm{val}(\phi) \geq \varepsilon$.
- *(NO instance.)* $\mathrm{val}(\phi) < \delta$.

We represent NAE-Ek-SAT constraints C as sets of k literals, and we use the notation $\mathrm{vars}(C)$ to denote the set of indices of variables involved in a NAE-E3-SAT constraint C. For example, if $k = 3$ and $C := \{v_1, \neg v_3, \neg v_4\}$ then $\mathrm{vars}(C) = \{1, 3, 4\}$. We define $\mathrm{sgn}(\ell) \in \{-1, 1\}$ and $\mathrm{var}(\ell)$ to be the sign and the index of the variable underlying a literal ℓ, respectively. For example, $\mathrm{sgn}(\neg v_4) = -1$ and $\mathrm{var}(\neg v_4) = 4$.

We next give a result showing NP-hardness of $(15/16+\varepsilon, 1)$-NAE-E3-SAT for any constant $\varepsilon > 0$ in Theorem 10. This hardness result appeared in work of Guruswami [14], which observed that it follows by combining a reduction due to Petrank [35] and a hardness result due to Håstad [18].

▶ **Theorem 10.** *For any constant $\varepsilon \in (0, 1/16)$, $(15/16 + \varepsilon, 1)$-NAE-E3-SAT is NP-hard.*

See our full version for a proof of Theorem 10.

3 Reduction from NAE-E3-SAT to BinaryGapCRP

3.1 The Reduction

We start by giving our reduction from (δ, ε)-NAE-E3-SAT to γ-$\mathrm{BinGapCRP}_p$ for certain values of $\delta, \varepsilon, \gamma, p$. In what follows, we will analyze the reduction and for which parameter values it is valid. We emphasize again that the reduction is similar to the NP-hardness reduction in [29] for $\mathrm{LinDisc}_\infty$.

Let ϕ be the input (δ, ε)-NAE-E3-SAT formula, and assume that ϕ has n variables $v_1, \dots, v_n$ and m constraints $C_1, \dots, C_m$. We define the incidence matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$ of ϕ as

$$B_{i,j} := B_{i,j}(\phi) = \begin{cases} 1 & \text{if constraint } C_i \text{ contains literal } v_j \ , \\ -1 & \text{if constraint } C_i \text{ contains literal } \neg v_j \ , \\ 0 & \text{otherwise} \ . \end{cases} \tag{6}$$

We will also use the following gadget matrix, defined in [29], which we analyze in Section 3.2.

$$\mathbf{G} := \begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{pmatrix} \ . \tag{7}$$

Let $\deg_\phi(v_j)$ denote the number of constraints C_i in ϕ containing either v_j or $\neg v_j$. When the formula ϕ is clear from context, we simply write $\deg(v_j)$. For $p \in [1, \infty]$, define the $n \times n$ diagonal matrix

$$\mathbf{D}_p := \mathbf{D}_p(\phi) = \mathrm{diag}(\deg(v_1)^{1/p}, \dots, \deg(v_n)^{1/p}) = \begin{pmatrix} \deg(v_1)^{1/p} & & \\ & \ddots & \\ & & \deg(v_n)^{1/p} \end{pmatrix} \ .$$

Finally, we output the γ-$\mathrm{BinGapCRP}_p$ instance consisting of the matrix

$$\mathbf{A} := \mathbf{A}_p(\phi) = \begin{pmatrix} \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} \\ & \mathbf{G} \otimes \mathbf{D}_p & \end{pmatrix} \in \mathbb{R}^{(m+3n) \times 3n} \ , \tag{8}$$

and distance threshold

$$r := (\varepsilon m \cdot (4/3)^p + (1 - \varepsilon) m \cdot 2^p + 3m(2 + (4/3)^p))^{1/p} > 0 \ . \tag{9}$$

3.2 Analysis of the matrix G

We begin by presenting some technical lemmas about $\mathbf{G}$ that will be useful in our hardness reductions. We start with Lemma 11, which adapts [29, Lemma 3] from the ℓ_∞ norm to ℓ_p norms.

▶ **Lemma 11.** *For any $\boldsymbol{u} \in [0,1]^3$ and $b \in \{-1,1\}$, there exists $\boldsymbol{z} \in \{0,1\}^3$ such that*

1. *(Low discrepancy w.r.t. $\mathbf{G}$.) $\|\mathbf{G}(\boldsymbol{u}-\boldsymbol{z})\|_p^p \le 2 + (4/3)^p$, where $\mathbf{G}$ is as defined in Equation (7).*
2. *(Sign agreement.) $b \cdot (\mathbf{1}^T(\boldsymbol{u}-\boldsymbol{z})) \ge 0$.*
3. *(Low discrepancy w.r.t. $\mathbf{1}^T$.) $|\mathbf{1}^T(\boldsymbol{u}-\boldsymbol{z})| \le 2$.*

Proof. The second and third properties are identical to the corresponding properties in [29, Lemma 3] (which are proved there), and so we only prove the first property. We follow the structure of the proof of [29, Lemma 3, Item 1] to prove the first property.

We assume without loss of generality that $0 \le u_1 \le u_2 \le u_3 \le 1$ throughout the proof. This assumption is justified by the easy-to-check fact that for all 3×3 permutation matrices $\mathbf{P}$ and $\boldsymbol{x} \in \mathbb{R}^3$, $\|\mathbf{G}\boldsymbol{x}\|_\infty = \|\mathbf{GP}\boldsymbol{x}\|_\infty$.

We now proceed to the proof of Item 1. We will consider three cases, and in each case we will bound each coordinate of $\mathbf{G}(\boldsymbol{u}-\boldsymbol{z})$ individually.

Case 1: $\mathbf{1}^T\boldsymbol{u} \ge 2$. In this case we set $\boldsymbol{z} := (0,1,1)^T$, and have

$$\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}) = \begin{pmatrix} u_1 + u_2 - u_3 \\ u_1 - u_2 + u_3 \\ -u_1 + u_2 + u_3 - 2 \end{pmatrix} .$$

We have following bounds:

- $-1 \le -u_3 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_1 = (\mathbf{G}\boldsymbol{u})_1 \le u_1 \le 1$.
- $-1 \le -u_2 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_2 = (\mathbf{G}\boldsymbol{u}))_2 \le u_3 \le 1$.
- $-4/3 \le u_3 - 2 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_3 = (\mathbf{G}\boldsymbol{u})_3 - 2 \le -u_1 \le 0$.

Case 2: $\mathbf{1}^T\boldsymbol{u} < 2$ and $(\mathbf{G}\boldsymbol{u})_3 \le 4/3$.[6] In this case we set $\boldsymbol{z} := \mathbf{0}$, and have

$$\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}) = \begin{pmatrix} u_1 + u_2 - u_3 \\ u_1 - u_2 + u_3 \\ -u_1 + u_2 + u_3 \end{pmatrix} .$$

We have the following bounds:

- $-1 \le -u_3 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_1 = (\mathbf{G}\boldsymbol{u})_1 \le u_1 \le 1$.
- $-1 \le -u_2 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_2 = (\mathbf{G}\boldsymbol{u})_2 \le u_3 \le 1$.
- $0 \le u_3 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_3 = (\mathbf{G}\boldsymbol{u})_3 \le 4/3$.

Case 3: $\mathbf{1}^T\boldsymbol{u} < 2$ and $(\mathbf{G}\boldsymbol{u})_3 > 4/3$. In this case we set $\boldsymbol{z} := (0,0,1)^T$, and have

$$\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}) = \begin{pmatrix} u_1 + u_2 - u_3 + 1 \\ u_1 - u_2 + u_3 - 1 \\ -u_1 + u_2 + u_3 - 1 \end{pmatrix} .$$

We have the following bounds:

[6] In fact, the assumption that $\mathbf{1}^T\boldsymbol{u} < 2$ is not even necessary here.

- $0 \le -u_3+1 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_1 = (\mathbf{G}\boldsymbol{u})_1+1 \le u_1+1 = \frac{1}{2}(\mathbf{1}^T\boldsymbol{u}-(\mathbf{G}\boldsymbol{u})_3)+1 < 1/3+1 = 4/3$.
- $-1 \le u_1 - 1 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_2 = (\mathbf{G}\boldsymbol{u})_2 - 1 \le u_3 - 1 \le 0$.
- $-1 \le u_3 - 1 \le (\mathbf{G}(\boldsymbol{u}-\boldsymbol{z}))_3 = (\mathbf{G}\boldsymbol{u})_3 - 1 \le \mathbf{1}^T\boldsymbol{u} - 1 < 1$.

In each case, we have $\|\mathbf{G}(\boldsymbol{u}-\boldsymbol{z})\|_p^p \le 2 + (4/3)^p$, as needed. ◀

Lemma 12 is the other part of the extension of [29, Lemma 3], and facilitates a reduction to BinGapCRP, rather than just LinDisc. We both adapt [29, Lemma 3] to ℓ_p norms for $p < \infty$ and show that $\mathbf{G}(1/2\cdot\mathbf{1}) = (1/2)\cdot\mathbf{1}$ is far from all vectors in $\mathcal{L}(\mathbf{G}) \setminus \{\mathbf{0},\mathbf{1}\} = \mathbf{G}\cdot(\mathbb{Z}^3 \setminus \{\mathbf{0},\mathbf{1}\})$ rather than just vectors in $\mathbf{G}\cdot(\{0,1\}^3 \setminus \{\mathbf{0},\mathbf{1}\})$.

▶ **Lemma 12.** *Let $\boldsymbol{z} \in \mathbb{Z}^3$ and $p \in [1,\infty]$. Then we have the following:*

1. *For $\boldsymbol{z} \in \{\mathbf{0},\mathbf{1}\}$, $\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{z}) \in \{-1/2,1/2\}^3$. Therefore, $\|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{z})\|_p^p = 3/2^p$ if $p<\infty$ and $\|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{z})\|_\infty = 1/2$.*
2. *For $\boldsymbol{z} \in \mathbb{Z}^3\setminus\{\mathbf{0},\mathbf{1}\}$, $\|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{z})\|_p^p \ge (2+3^p)/2^p$ if $p<\infty$ and $\|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{z})\|_\infty \ge 3/2$.*

Proof. We note that $\mathbf{G}((1/2)\cdot\mathbf{1}-\boldsymbol{z}) = (1/2)\cdot\mathbf{1} - \mathbf{G}\boldsymbol{z}$ so that[7]

$$\{(\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{z}) : \boldsymbol{z}\in\mathbb{Z}^3\} = \{\mathbf{G}\boldsymbol{z} : \boldsymbol{z}\in\mathbb{Z}^3\} + (1/2)\cdot\mathbf{1} \subseteq (\mathbb{Z}+1/2)^3 \ .$$

That is, for any $\boldsymbol{z}\in\mathbb{Z}^3$, each coordinate of $G((1/2\cdot\mathbf{1}-\boldsymbol{z})$ is contained in $\mathbb{Z}+1/2$. Fix such a coefficient vector $\boldsymbol{z}\in\mathbb{Z}^3$. Then for $p<\infty$, $\|\mathbf{G}((1/2)\cdot\mathbf{1}-\boldsymbol{z})\|_p^p = 3/2^p$ if $\mathbf{G}\boldsymbol{z}\in\{0,1\}^3$ and $\|\mathbf{G}((1/2)\cdot\mathbf{1}-\boldsymbol{z})\|_p^p \ge (2+3^p)/2^p$ otherwise. Similarly, $\|\mathbf{G}((1/2)\cdot\mathbf{1}-\boldsymbol{z})\|_\infty = 1/2$ if $\mathbf{G}\boldsymbol{z}\in\{0,1\}^3$ and $\|\mathbf{G}((1/2)\cdot\mathbf{1}-\boldsymbol{z})\|_\infty \ge 3/2$ otherwise. The lemma then follows by noting that $\mathcal{L}(\mathbf{G})\cap\{0,1\}^3 = \{\mathbf{0},\mathbf{1}\}$. ◀

3.3 Completeness Analysis

We start with an expression for the value of $\|\mathbf{A}(\boldsymbol{w}-\boldsymbol{x})\|_p^p$ for the matrix $\mathbf{A}=\mathbf{A}(\phi)$ defined in Equation (8) in terms of a NAE-E3-SAT formula ϕ and vectors $\boldsymbol{w},\boldsymbol{x}$.

▶ **Lemma 13.** *Let $p\ge1$, and for $i\in[3]$, let $\boldsymbol{y}^i\in\mathbb{R}^n$, let ϕ be an* NAE-E3-SAT *formula with constraints $C_1,\dots,C_m$, and let $\mathbf{A} := \mathbf{A}_p(\phi)$ as defined in Equation (8). Let $\boldsymbol{y} := (\boldsymbol{y}^1,\boldsymbol{y}^2,\boldsymbol{y}^3)\in\mathbb{R}^{3n}$, let $\boldsymbol{y}^{sum} := \boldsymbol{y}^1+\boldsymbol{y}^2+\boldsymbol{y}^3\in\mathbb{R}^n$, and for $j\in[n]$, let $\boldsymbol{y}_j := (y_j^1,y_j^2,y_j^3)\in\mathbb{R}^3$. Then*

$$\|\mathbf{A}\boldsymbol{y}\|_p^p = (1/3^p)\cdot\|\mathbf{B}\boldsymbol{y}^{sum}\|_p^p + \sum_{j=1}^n \deg(v_j)\cdot\left\|\mathbf{G}\boldsymbol{y}_j\right\|_p^p \ . \tag{10}$$

Let $\boldsymbol{x}^1,\boldsymbol{x}^2,\boldsymbol{x}^3\in\mathbb{R}^n$, let $\boldsymbol{x} := (\boldsymbol{x}^1,\boldsymbol{x}^2,\boldsymbol{x}^3)\in\mathbb{R}^{3n}$, and let $\boldsymbol{x}^{sum} := \boldsymbol{x}^1+\boldsymbol{x}^2+\boldsymbol{x}^3\in\mathbb{R}^n$. Furthermore, for $j\in[n]$, let $\boldsymbol{x}_j := (x_j^1,x_j^2,x_j^3)^T\in\mathbb{R}^3$, and let $\boldsymbol{b}_i$ denote the ith row of $\mathbf{B}$. Then in particular,

$$\left\|\mathbf{A}\left(\frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}\right)\right\|_p^p = \sum_{i=1}^m\left(\left|\frac{1}{3}\boldsymbol{b}_i\left(\frac{3}{2}\cdot\mathbf{1}-\boldsymbol{x}^{sum}\right)\right|^p + \sum_{j\in\mathrm{vars}(C_i)}\left\|\mathbf{G}\left(\frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}_j\right)\right\|_p^p\right) \ . \tag{11}$$

[7] This is equivalent to the fact that $\mathbf{G}$ has $\mathbf{1}$ as an eigenvector with corresponding eigenvalue 1.

Proof. For Equation (10),

$$\begin{aligned}\|\mathbf{A}(\boldsymbol{y})\|_p^p &= \left\|(1/3)\cdot(\mathbf{B}\boldsymbol{y}^1+\mathbf{B}\boldsymbol{y}^2+\mathbf{B}\boldsymbol{y}^3)\right\|_p^p + \|(\mathbf{G}\otimes\mathbf{D}_p)\cdot\boldsymbol{y}\|_p^p \\ &= \|(1/3)\cdot\mathbf{B}\boldsymbol{y}^{\mathrm{sum}}\|_p^p + \sum_{j=1}^n \deg(v_j)\cdot\left\|\mathbf{G}\boldsymbol{y}_j\right\|_p^p \ .\end{aligned}$$

For Equation (11), we set $\boldsymbol{y} := \frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}$ and continue by re-grouping some terms.

$$\begin{aligned}\left\|\mathbf{A}\left(\frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}\right)\right\|_p^p &= \frac{1}{3^p}\cdot\left\|\mathbf{B}\left(\frac{3}{2}\cdot\mathbf{1}-\boldsymbol{x}^{\mathrm{sum}}\right)\right\|_p^p + \sum_{j=1}^n \deg(v_j)\cdot\left\|\mathbf{G}\left(\frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}_j\right)\right\|_p^p \\ &= \frac{1}{3^p}\cdot\sum_{i=1}^m\left|\boldsymbol{b}_i\left(\frac{3}{2}\cdot\mathbf{1}-\boldsymbol{x}^{\mathrm{sum}}\right)\right|^p + \sum_{j=1}^n \deg(v_j)\cdot\left\|\mathbf{G}\left(\frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}_j\right)\right\|_p^p \\ &= \sum_{i=1}^m\left(\left|\frac{1}{3}\boldsymbol{b}_i\left(\frac{3}{2}\cdot\mathbf{1}-\boldsymbol{x}^{\mathrm{sum}}\right)\right|^p + \sum_{j\in\mathrm{vars}(C_i)}\left\|\mathbf{G}\left(\frac{1}{2}\cdot\mathbf{1}-\boldsymbol{x}_j\right)\right\|_p^p\right).\end{aligned}$$

The first equality follows by setting $\boldsymbol{w} := 1/2\cdot\mathbf{1}$ in Equation (10). The third equality follows by noting for any expression $K(j)$ depending on j,

$$\sum_{j=1}^n \deg(v_j)K(j) = \sum_{j=1}^n\ \sum_{\mathrm{vars}(C_i):x_j\in\mathrm{vars}(C_i)} K(j) = \sum_{i=1}^m\ \sum_{j:x_j\in\mathrm{vars}(C_i)} K(j)\ . \tag{12}$$

Indeed, one can view the two double sums as counting constraint-variable incidences in two different ways. ◀

We now state lemmas that will be used in the proof of Theorem 17. We begin with a lemma for completeness in which we upper bound $\mathrm{lindisc}_p(\mathbf{A})$.

▶ **Lemma 14.** *Let $p\in[1,\infty)$, let $\varepsilon\in[0,1]$, let ϕ be an* NAE-E3-SAT *instance with m constraints and $\mathrm{val}(\phi)\geq\varepsilon$, and let $\mathbf{A} := \mathbf{A}_p(\phi)$ be defined as in Equation* (8)*. Then*

$$\mathrm{lindisc}_p(\mathbf{A})^p \leq (\varepsilon\cdot(4/3)^p + (1-\varepsilon)\cdot 2^p + 3(2+(4/3)^p))\cdot m\ . \tag{13}$$

Proof. Let ψ be an assignment to the variables of ϕ such that $\mathrm{val}_\psi(\phi)\geq\varepsilon$. Let $\boldsymbol{w}^1,\boldsymbol{w}^2,\boldsymbol{w}^3\in[0,1]^n$, and let $\boldsymbol{w} := (\boldsymbol{w}^1,\boldsymbol{w}^2,\boldsymbol{w}^3)$. We define $\boldsymbol{x} := (\boldsymbol{x}^1,\boldsymbol{x}^2,\boldsymbol{x}^3)$ for $\boldsymbol{x}^1,\boldsymbol{x}^2,\boldsymbol{x}^3\in\{0,1\}^n$ as follows. For each $j\in[n]$, the vector $\boldsymbol{x}_j := (x_j^1,x_j^2,x_j^3)^T$ is the vector $\boldsymbol{z}$ guaranteed by Lemma 11 applied to the vector $\boldsymbol{u} := \boldsymbol{w}_j = (w_j^1,w_j^2,w_j^3)^T$ and sign $b := 1-2\psi(v_j)$.

We begin by upper bounding $\|\mathbf{A}(\boldsymbol{w}-\boldsymbol{x})\|_p^p$, which we do by upper bounding the terms on the right-hand side of Equation (10) individually with $\boldsymbol{y} := \boldsymbol{w}-\boldsymbol{x}\in\mathbb{R}^{3n}$. From Lemma 11, Item 1 and the fact that ϕ has total degree $\sum_{j=1}^n\deg(v_j) = 3m$, we have

$$\sum_{j=1}^n\deg(v_j)\cdot\|\mathbf{G}(\boldsymbol{w}_j-\boldsymbol{x}_j)\|_p^p \leq \sum_{j=1}^n\deg(v_j)\cdot(2+(4/3)^p) = 3m(2+(4/3)^p). \tag{14}$$

To upper bound $\left\|\frac{1}{3}\mathbf{B}(\boldsymbol{w}^{\mathrm{sum}}-\boldsymbol{x}^{\mathrm{sum}})\right\|_p^p$, we upper bound the magnitude of each coordinate i of $\mathbf{B}(\boldsymbol{w}^{\mathrm{sum}}-\boldsymbol{x}^{\mathrm{sum}})\in\mathbb{R}^m$. Suppose that constraint $C_i = \{\ell_{i,1},\ell_{i,2},\ell_{i,3}\}$ for some literals $\ell_{i,1},\ell_{i,2},\ell_{i,3}$. Then

$$(\mathbf{B}(\boldsymbol{w}^{\mathrm{sum}}-\boldsymbol{x}^{\mathrm{sum}}))_i = \sum_{j=1}^3\mathrm{sgn}(\ell_{i,j})\cdot(w^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})} - x^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})})\ . \tag{15}$$

By Lemma 11, Item 3, we have that $|\mathbf{1}^T(\boldsymbol{w}_j - \boldsymbol{x}_j)| = |w_j^{\mathrm{sum}} - x_j^{\mathrm{sum}}| \leq 2$ for all $j \in [n]$. So, the expression in Equation (15) always has magnitude at most 6 by triangle inequality and has magnitude at most 4 as long as the signs of its summands are not all the same (i.e., as long as $\mathrm{sgndiff}(\boldsymbol{b}_i, \boldsymbol{w}^{\mathrm{sum}} - \boldsymbol{x}^{\mathrm{sum}})$ holds, where $\boldsymbol{b}_i$ is the ith row of $\mathbf{B}$). We show that this is the case if ψ satisfies constraint C_i. Assume without loss of generality that $w^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})} \neq x^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})}$ for $j \in [3]$ (since otherwise the expression clearly has magnitude at most 4). By Lemma 11, Item 2 and our choice of the sign $1 - 2\psi(v_j)$ when applying Lemma 11 to $\boldsymbol{w}_j$, we have that for all $j \in [n]$, $(1 - 2\psi(v_j)) \cdot (w_j^{\mathrm{sum}} - x_j^{\mathrm{sum}}) \geq 0$, which implies that $(1 - 2\psi(v_j)) = \mathrm{sgn}(w_j^{\mathrm{sum}} - x_j^{\mathrm{sum}})$. So, for $j \in [3]$,

$$\mathrm{sgn}(\ell_{i,j}) \cdot \mathrm{sgn}(w^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})} - x^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})}) = \mathrm{sgn}(\ell_{i,j}) \cdot (1 - 2\psi(v_j)) = \begin{cases} -1 & \text{if } \psi \text{ satisfies literal } \ell_{i,j} \ , \\ 1 & \text{otherwise} \ . \end{cases} \tag{16}$$

By definition, a NAE-E3-SAT constraint C_i is satisfied if and only if at least one of its literals is satisfied and at least one is not satisfied. So, if ψ satisfies C_i, $\mathrm{sgn}(\ell_{i,j}) \cdot \mathrm{sgn}(w^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})} - x^{\mathrm{sum}}_{\mathrm{var}(\ell_{i,j})})$ is equal to 1 for some $j \in [3]$ and is equal to -1 for some $j \in [3]$. It follows that the expression in Equation (15) has magnitude at most 4. Because $\mathrm{val}(\phi) \geq \varepsilon$ by assumption,

$$\|(1/3) \cdot \mathbf{B}(\boldsymbol{w}^{\mathrm{sum}} - \boldsymbol{x}^{\mathrm{sum}})\|_p^p \leq \frac{1}{3^p} \cdot (\varepsilon m \cdot 4^p + (1-\varepsilon) m \cdot 6^p) = (\varepsilon \cdot (4/3)^p + (1-\varepsilon) \cdot 2^p) \cdot m \ . \tag{17}$$

Combining the upper bounds of Equations (14) and (17) to the terms in Equation (10) (with $\boldsymbol{y} := \boldsymbol{w} - \boldsymbol{x}$) implies Equation (13), as needed. ◀

3.4 Soundness Analysis

We now give two lemmas for analyzing the soundness of our reduction. The first lemma establishes that vectors of the form $\mathbf{A}\boldsymbol{x}$ where $\boldsymbol{x} \in \{0,1\}^{3n}$ are closest vectors in $\mathcal{L}(\mathbf{A})$ to $\mathbf{A} \cdot (1/2 \cdot \mathbf{1})$.

▶ **Lemma 15.** *Let $p \in [1, \infty]$, let ϕ be an* NAE-E3-SAT *instance, and let $\mathbf{A} := \mathbf{A}(\phi)$ be as defined in Equation* (8)*. Then for every $\boldsymbol{y} \in \mathbb{Z}^{3n}$, there exists $\boldsymbol{x} := (\boldsymbol{x}', \boldsymbol{x}', \boldsymbol{x}') \in \{0,1\}^{3n}$ with $\boldsymbol{x}' \in \{0,1\}^n$ such that $\|\mathbf{A}(1/2 \cdot \mathbf{1} - \boldsymbol{x})\|_p \leq \|\mathbf{A}(1/2 \cdot \mathbf{1} - \boldsymbol{y})\|_p$.*

Proof. Let $\boldsymbol{y}^1, \boldsymbol{y}^2, \boldsymbol{y}^3 \in \mathbb{Z}^n$ be such that $\boldsymbol{y} = (\boldsymbol{y}^1, \boldsymbol{y}^2, \boldsymbol{y}^3) \in \mathbb{Z}^{3n}$, let $\boldsymbol{y}^{\mathrm{sum}} := \boldsymbol{y}^1 + \boldsymbol{y}^2 + \boldsymbol{y}^3 \in \mathbb{Z}^n$, and for each $j \in [n]$, let $\boldsymbol{y}_j := (y_j^1, y_j^2, y_j^3)$. For each $j \in [n]$, define

$$x'_j := \begin{cases} 0 & \text{if } y_j^{\mathrm{sum}} \leq 1 \ , \\ 1 & \text{if } y_j^{\mathrm{sum}} \geq 2 \ . \end{cases} \tag{18}$$

Furthermore, define $\boldsymbol{x}' := (x'_1, \ldots, x'_n)^T \in \{0,1\}^n$, $\boldsymbol{x} := (\boldsymbol{x}', \boldsymbol{x}', \boldsymbol{x}') \in \{0,1\}^{3n}$, and for each $j \in [n]$, $\boldsymbol{x}_j := (x'_j, x'_j, x'_j) \in \{0,1\}^3$.

We continue the proof by upper and lower bounding the ith terms in the sum in Equation (11) for $\|\mathbf{A}(1/2 \cdot \mathbf{1} - \boldsymbol{x})\|_p^p$ and for $\|\mathbf{A}(1/2 \cdot \mathbf{1} - \boldsymbol{y})\|_p^p$, respectively. Specifically, we show that the value of the ith term in the former sum is at most the value of the ith term in the latter sum for every $i \in [m]$. By the definition of $\boldsymbol{x}$ in terms of $\boldsymbol{y}$, these ith terms are the same if $\boldsymbol{y}_j \in \{\mathbf{0}, \mathbf{1}\}$ for all $j \in \mathrm{vars}(C_i)$, where C_i is the ith constraint of ϕ. So, we assume without loss of generality that this is not the case, i.e., that for every $i \in [m]$ there exists $j \in \mathrm{var}(C_i)$ such that $\boldsymbol{y}_j \in \mathbb{Z}^3 \setminus \{\mathbf{0}, \mathbf{1}\}$.

We begin by upper bounding the terms in $\|\mathbf{A}(1/2\cdot\mathbf{1}-\boldsymbol{x})\|_p^p$. Fix some $i \in [m]$ corresponding to a term of Equation (11). Because $\boldsymbol{x}_j \in \{\mathbf{0},\mathbf{1}\}$ for all $j \in [n]$ by definition, by Lemma 12, we have the equality

$$\sum_{j\in \mathrm{vars}(C_i)} \|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{x}_j)\|_p^p = \sum_{j\in \mathrm{vars}(C_i)} 3/2^p = 3\cdot(3/2^p) = 9/2^p \ . \tag{19}$$

Let $\boldsymbol{b}_i$ be the ith row of the matrix $\mathbf{B}$ in Equation (6), as in Equation (11). Again using the fact that each $\boldsymbol{x}_j \in \{\mathbf{0},\mathbf{1}\}$, we have

$$|\langle \boldsymbol{b}_i, (1/2)\cdot\mathbf{1}-\boldsymbol{x}'\rangle| = \Big|\sum_{j\in \mathrm{vars}(C_i)} B_{i,j}\cdot(1/2-x'_j)\Big| = \begin{cases} 1/2 & \text{if } \mathrm{sgndiff}(\boldsymbol{b}_i, 1/2\cdot\mathbf{1}-\boldsymbol{x}') \ , \\ 3/2 & \text{otherwise} \ . \end{cases} \tag{20}$$

Combining Equations (19) and (20), we have

$$|\langle \boldsymbol{b}_i, (1/2)\cdot\mathbf{1}-\boldsymbol{x}'\rangle|^p + \sum_{j\in C_i} \|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{x}_j)'\|_p^p = \begin{cases} 10/2^p & \text{if } \mathrm{sgndiff}(\boldsymbol{b}_i, 1/2\cdot\mathbf{1}-\boldsymbol{x}') \ , \\ (9+3^p)/2^p & \text{otherwise} \ . \end{cases} \tag{21}$$

We now lower bound $\|\mathbf{A}(1/2\cdot\mathbf{1}-\boldsymbol{y})\|_p^p$. We have that

$$\sum_{j\in \mathrm{vars}(C_i)} \|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{y}_j)\|_p^p \geq (3/2)^p + 8/2^p \ . \tag{22}$$

Indeed, by Lemma 12, the terms for which $\boldsymbol{y}_j \in \{\mathbf{0},\mathbf{1}\}$ are equal to $3/2^p$ and the terms in the sum for which $\boldsymbol{y}_j \notin \{\mathbf{0},\mathbf{1}\}$ must have value at least $(3/2)^p + 2/2^p$. And, since we are only considering constraints C_i for which there exists j such that $\boldsymbol{y}_j \notin \{\mathbf{0},\mathbf{1}\}$, at least one term must have value at least $(3/2)^p + 2/2^p$.

We continue by giving a lower bound for $|(1/3)\cdot\boldsymbol{b}_i(3/2\cdot\mathbf{1}-\boldsymbol{y}^{\mathrm{sum}})|$. Here, we observe that $(1/2)-(y_j^{\mathrm{sum}}/3) \in \{1/6, 3/6, 5/6\} + \mathbb{Z}$. This implies that

$$\begin{aligned} \langle \boldsymbol{b}_i, (1/2)\cdot\mathbf{1}-\boldsymbol{y}^{\mathrm{sum}}/3\rangle &= \Big|\sum_{j\in \mathrm{vars}(C_i)} B_{i,j}\cdot(1/2-y_j^{\mathrm{sum}}/3)\Big| \\ &\geq \begin{cases} 1/6 & \text{if } \mathrm{sgndiff}(\boldsymbol{b}_i, 1/2\cdot\mathbf{1}-\boldsymbol{y}^{\mathrm{sum}}/3) \\ 1/2 & \text{otherwise} \end{cases} , \end{aligned} \tag{23}$$

where the inequality follows by noting that because all three summands in the middle expression are elements of $\{1/6, 3/6, 5/6\} + \mathbb{Z}$. If all signs of the summands are the same, one obtains the lower bound $3/6 = 1/2$, whereas if the signs of the summands differ, one only has the lower bound of $1/6$.

By combining Equations (22) and (23), we have the lower bound

$$\begin{aligned} |\langle \boldsymbol{b}_i, (1/2)\cdot\mathbf{1}-\boldsymbol{y}^{\mathrm{sum}}/3\rangle|^p + &\sum_{j\in \mathrm{vars}(C_i)} \|\mathbf{G}(1/2\cdot\mathbf{1}-\boldsymbol{y}_j)\|_p^p \\ &\geq \begin{cases} 1/6^p + (3/2)^p + 8/2^p & \text{if } \mathrm{sgndiff}(\boldsymbol{b}_i, 1/2\cdot\mathbf{1}-\boldsymbol{y}^{\mathrm{sum}}/3) \\ (9+3^p)/2^p & \text{otherwise} \end{cases} . \end{aligned} \tag{24}$$

Finally, by definition of $\boldsymbol{x}'$ (Equation (18)), we have that for all $j \in [n]$, $\mathrm{sgn}(x'_j - 1/2) = \mathrm{sgn}(y^{\mathrm{sum}}_j/3 - 1/2)$. Then, it follows that for all $i \in [m]$ that $\mathrm{sgndiff}(\boldsymbol{b}_i, 1/2 \cdot \mathbf{1} - \boldsymbol{x}')$ if and only if $\mathrm{sgndiff}(\boldsymbol{b}_i, 1/2 \cdot \mathbf{1} - \boldsymbol{y}^{\mathrm{sum}}/3)$. By noting that for all $p \geq 1$, $2/2^p \leq 1/6^p + (3/2)^p$, we combine Equations (21) and (24) to obtain that for all $i \in [m]$,

$$\begin{aligned} &|\langle \boldsymbol{b}_i, (1/2) \cdot \mathbf{1} - \boldsymbol{x}' \rangle|^p + \sum_{j \in \mathrm{vars}(C_i)} \|\mathbf{G}(1/2 \cdot \mathbf{1} - \boldsymbol{x}_j)\|_p^p \\ &\leq |\langle \boldsymbol{b}_i, (1/2) \cdot \mathbf{1} - \boldsymbol{y}^{\mathrm{sum}}/3 \rangle|^p + \sum_{j \in \mathrm{vars}(C_i)} \|\mathbf{G}(1/2 \cdot \mathbf{1} - \boldsymbol{y}_j)\|_p^p \ , \end{aligned}$$

where we note that the left- and right-hand sides of the inequality are equal to the ith term in the sum in Equation (11). The result follows. ◀

The following lemma shows soundness for our reduction. Specifically, it shows how to construct an assignment satisfying a large fraction of constraints in an NAE-E3-SAT formula ϕ from a close vector $\mathbf{A}\boldsymbol{x}$ for $\boldsymbol{x} \in \{0,1\}^{3n}$ to $\mathbf{A}(1/2 \cdot \mathbf{1}_{3n})$, where $\mathbf{A} = \mathbf{A}_p(\phi)$.

► **Lemma 16.** *Let $p \in [1, \infty)$, let ϕ be an* NAE-E3-SAT *instance with n variables and m constraints, let $\mathbf{A} := \mathbf{A}_p(\phi)$ be as defined in Equation* (8)*, and let $\delta \in [0,1]$. Suppose that there exists some vector $\boldsymbol{x} := (\boldsymbol{x}^1, \boldsymbol{x}^2, \boldsymbol{x}^3) \in \mathbb{Z}^{3n}$ such that*

$$\|\mathbf{A}((1/2) \cdot \mathbf{1} - \boldsymbol{x})\|_p^p \leq \delta m \cdot (10/2^p) + (1-\delta) m \cdot (9/2^p + (3/2)^p) \ . \tag{25}$$

Then $\mathrm{val}(\phi) \geq \delta$*.*

Proof. Suppose that there exists $\boldsymbol{x} \in \mathbb{Z}^{3n}$ satisfying Equation (25). Then by Lemma 15, we may assume without loss of generality that $\boldsymbol{x} = (\boldsymbol{x}', \boldsymbol{x}', \boldsymbol{x}')$ for $\boldsymbol{x}' \in \{0,1\}^n$ and let $\boldsymbol{x}_j := (x'_j, x'_j, x'_j)$ for each $j \in [n]$. Therefore, for each $i \in [m]$, both Equations (19) and (20) hold for $\boldsymbol{x}$.

First, we obtain the bound

$$\begin{aligned} \frac{1}{3^p} \cdot \|\mathbf{B}((3/2) \cdot \mathbf{1} - \boldsymbol{x}^{\mathrm{sum}})\|_p^p &= \|\mathbf{A}((1/2) \cdot \mathbf{1} - \boldsymbol{x})\|_p^p - \sum_{j=1}^n \deg(v_j) \cdot \|\mathbf{G}((1/2) \cdot \mathbf{1} - \boldsymbol{x}_j)\|_p^p \\ &= \|\mathbf{A}((1/2) \cdot \mathbf{1} - \boldsymbol{x})\|_p^p - m \cdot (9/2^p) \\ &\leq \delta m \cdot (1/2^p) + (1-\delta) m \cdot (3/2)^p \ , \end{aligned}$$

where the first equality uses Equation (10) with $\boldsymbol{y} := (1/2) \cdot \mathbf{1} - \boldsymbol{x}$, the second equality uses Equation (19), and the inequality is by assumption (Equation (25)).

It then follows from Equation (20) that at least an δ fraction of indices $i \in [m]$ satisfy $\mathrm{sgndiff}(\boldsymbol{b}_i, 1/2 \cdot \mathbf{1} - \boldsymbol{x}')$. Define an assignment ψ to the variables $v_1, \ldots, v_n$ of ϕ by $\psi(v_j) = x'_j$. We claim that if $\mathrm{sgndiff}(\boldsymbol{b}_i, 1/2 \cdot \mathbf{1} - \boldsymbol{x}')$ then ψ satisfies constraint C_i, which combined with the fact that at least an δ fraction of indices $i \in [m]$ satisfy $\mathrm{sgndiff}(\boldsymbol{b}_i, 1/2 \cdot \mathbf{1} - \boldsymbol{x}')$, would imply the lemma.

It remains to prove the claim. Observe that if constraint $C_i = \{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\}$ for literals $\ell_{i,1}, \ell_{i,2}, \ell_{i,3}$ then, as in Equation (15),

$$(\mathbf{B}(1/2 \cdot \mathbf{1} - \boldsymbol{x}'))_i = \sum_{j=1}^3 \mathrm{sgn}(\ell_{i,j}) \cdot (1/2 - x_{\mathrm{var}(\ell_{i,j})}) \ .$$

Furthermore, as in Equation (16), for $j \in [3]$,

$$\mathrm{sgn}(\ell_{i,j}) \cdot \mathrm{sgn}(1/2 - x_{\mathrm{var}(\ell_{i,j})}) = \begin{cases} -1 & \text{if } \psi \text{ satisfies literal } \ell_{i,j} \ , \\ 1 & \text{otherwise} \ . \end{cases}$$

The predicate $\mathrm{sgndiff}(\boldsymbol{b}_i, 1/2 \cdot \mathbf{1} - \boldsymbol{x}')$ holds exactly when $\mathrm{sgn}(\ell_{i,j}) \cdot (1/2 - x_{\mathrm{var}(\ell_{i,j})}) = 1$ for some $j \in [3]$ and $\mathrm{sgn}(\ell_{i,j}) \cdot (1/2 - x_{\mathrm{var}(\ell_{i,j})}) = -1$ for some $j \in [3]$, meaning that ψ satisfies at least one literal in C_i and does not satisfy at least one literal in C_i. This in turn implies that ψ satisfies C_i, as needed. ◀

3.5 The Main Theorem

We now state and prove the main technical theorem of the paper.

▶ **Theorem 17.** *Let $\delta, \varepsilon \in \mathbb{R}$ be such that $0 < \delta \le \varepsilon \le 1$ and let $p \in [1, \infty)$. There is a polynomial-time mapping reduction from (δ, ε)-NAE-E3-SAT to γ-$\mathrm{BinGapCRP}_p$ for*

$$\gamma = \gamma(\delta, \varepsilon, p) := \Big(\frac{\delta \cdot 10/2^p + (1-\delta) \cdot ((3/2)^p + 9/2^p)}{\varepsilon \cdot (4/3)^p + (1-\varepsilon)2^p + 3(2 + (4/3)^p)}\Big)^{1/p} . \tag{26}$$

Proof. Let ϕ be the input instance of (δ, ε)-NAE-E3-SAT. Assume that ϕ has m constraints. The reduction outputs the $\mathrm{BinGapCRP}_p$ instance consisting of the matrix $\mathbf{A} = \mathbf{A}_p(\phi)$ defined in Equation (8) and the distance threshold $r := (\varepsilon m \cdot (4/3)^p + (1-\varepsilon)m \cdot 2^p + 3m(2 + (4/3)^p))^{1/p}$ defined in Equation (9).

It is clear that the reduction runs in polynomial time, and it remains to show correctness. We start with completeness. If ϕ is a YES instance of (δ, ε)-NAE-E3-SAT, then $\mathrm{lindisc}_p(\mathbf{A}) \le r$ by Lemma 14 and so the output is a YES instance of $\mathrm{BinGapCRP}_p$.

For soundness, we show the contrapositive. Suppose that

$$\mu_p(\mathcal{L}(\mathbf{A})) \le \gamma r = (\delta m \cdot 10/2^p + (1-\delta)m \cdot ((3/2)^p + 9/2^p))^{1/p} .$$

Then there exists $\boldsymbol{y} \in \mathbb{Z}^{3n}$ such that $\|\mathbf{A}((1/2 \cdot \mathbf{1}) - \boldsymbol{y})\|_p \le \gamma r$, and so $\mathrm{val}(\phi) \ge \delta$ by Lemma 16. The result follows. ◀

We conclude by obtaining Theorem 1, which asserts that γ-$\mathrm{BinGapCRP}_p$ is NP-hard for any constant $1 \le \gamma < \gamma(p)$, as a corollary of Theorem 17.

Proof of Theorem 1. Let $\gamma(p) := \gamma(15/16, 1, p)$, where $\gamma(\cdot, \cdot, \cdot)$ is as defined in Equation (26). Then

$$\begin{aligned}\gamma(p) &= \Big(\frac{(15/16) \cdot 10/2^p + (1/16) \cdot ((3/2)^p + 9/2^p)}{(4/3)^p + 3(2 + (4/3)^p)}\Big)^{1/p} \\ &= \Big(\frac{9^p + 159 \cdot 3^p}{8^{p+2} + 96 \cdot 6^p}\Big)^{1/p} ,\end{aligned}$$

By Theorem 10, $(15/16 + \varepsilon', 1)$-NAE-E3-SAT is NP-hard for every constant $\varepsilon' > 0$. So, by Theorem 17, $\gamma(15/16 + \varepsilon', 1, p)$-BinGapCRP is NP-hard for every such ε'. The theorem follows by noting that for every constant $\varepsilon > 0$ there exists a constant $\varepsilon' > 0$ such that $\gamma(15/16 + \varepsilon', 1, p) \ge \gamma(p) - \varepsilon$. ◀

4 Π_2-hardness of approximating BinaryGapCRP in the ℓ_∞ norm

In this section, we show that the Π_2-hardness of approximation result for $\mathrm{LinDisc}_\infty$ given in Manurangsi [29] can be extended to $\mathrm{BinGapCRP}_\infty$. We first give the definition of ∀∃-NAE-E3-SAT, the Π_2-hard problem that we reduce to approximate BinGapCRP. The problem ∀∃-NAE-3-SAT (in which constraints may have either two or three literals) was originally shown to be Π_2-hard in [12], and (a very special case of) ∀∃-NAE-E3-SAT was shown to be Π_2-hard in [10, Theorem 4.7].

▶ **Definition 18** ($\forall\exists$-NAE-E3-SAT). *The* $\forall\exists$-NAE-E3-SAT *problem is the decision problem where, on input a* NAE-E3-SAT *formula* $\phi(V_A, V_E)$ *defined over the disjoint union* $V_A \sqcup V_E$ *of variable sets* V_A *and* V_E*, the goal is to distinguish between the following two cases.*

- *(YES instance.) For every assignment* $\psi_A : V_A \to \{0,1\}$*, there exists an assignment* $\psi_E : V_E \to \{0,1\}$ *such that* (ψ_A, ψ_E) *satisfies* ϕ*.*
- *(NO instance.) There exists an assignment* $\psi_A : V_A \to \{0,1\}$ *such that for all assignments* $\psi_E : V_E \to \{0,1\}$*),* (ψ_A, ψ_E) *does not satisfy* ϕ*.*

Let $\phi(V_A, V_E)$ be a NAE-E3-SAT formula with $n' = |V_A|$ and $n - n' = |V_E|$. We define the matrix $\mathbf{A}' = \mathbf{A}'(\phi) \in \mathbb{R}^{(m+3n+2n')\times(3n+n')}$ (originally defined in [29]) as

$$\mathbf{A}' := \left(\begin{array}{cccccc|c} \multicolumn{6}{c|}{\mathbf{A}_\infty(\phi)} & \mathbf{0}_{n'\times n'} \\ \hline \frac{2}{3}\mathbf{I}_{n'} & \mathbf{0}_{n'\times(n-n')} & \frac{2}{3}\mathbf{I}_{n'} & \mathbf{0}_{n'\times(n-n')} & \frac{2}{3}\mathbf{I}_{n'} & \mathbf{0}_{n'\times(n-n')} & -2\mathbf{I}_{n'} \\ \hline \multicolumn{6}{c|}{\mathbf{0}_{n'\times 3n}} & \frac{8}{3}\mathbf{I}_{n'} \end{array}\right), \tag{27}$$

where $\mathbf{A}_\infty(\phi)$ is as defined in Equation (8) and where we have used the notation $\mathbf{0}_{m\times n}$ to denote an $m \times n$ matrix of 0s.

We next give a simple expression for $\|\mathbf{A}'(\boldsymbol{y}')\|_\infty$ for a vector $\boldsymbol{y}'$ in terms of the blocks of $\mathbf{A}'$.

▶ **Lemma 19** ([29, Equation 4]). *Let* $\phi(V_A, V_E)$ *be an* $\forall\exists$-NAE-E3-SAT *instance with* n' *universal variables (variables in* V_A*) and* $n - n'$ *existential variables (variables in* V_E*), and for each* $i \in [3]$*, let* $\boldsymbol{y}^i \in \mathbb{R}^n$*, and let* $\boldsymbol{y}^* \in \mathbb{R}^{n'}$*. Furthermore, let* $\boldsymbol{y}' := (\boldsymbol{y}^1, \boldsymbol{y}^2, \boldsymbol{y}^3, \boldsymbol{y}^*) = (\boldsymbol{y}, \boldsymbol{y}^*) \in \mathbb{R}^{3n+n'}$*. Finally, let* $\boldsymbol{y}^{sum} = \boldsymbol{y}^1 + \boldsymbol{y}^2 + \boldsymbol{y}^3$*. Then* $\mathbf{A}' = \mathbf{A}'(\phi)$ *satisfies*

$$\|\mathbf{A}'\boldsymbol{y}'\|_\infty = \max\{\|\mathbf{A}\boldsymbol{y}\|_\infty, \max_{i\in[n']}|(2/3)\cdot y_i^{sum} - 2y_i^*|, (8/3)\cdot \max_{i\in[n']}|y_i^*|\} \tag{28}$$

Proof. The claim follows by inspection of the blocks of $\mathbf{A}'\boldsymbol{y}'$. ◀

We next state the completeness and soundness results from the reduction in [29] for showing Π_2-hardness of $\mathrm{lindisc}_\infty(A)$.

▶ **Lemma 20** (Completeness for Π_2-hardness reduction, [29]). *Let* ϕ *be an* $\forall\exists$-NAE-E3-SAT *instance, and* $\mathbf{A}' := \mathbf{A}'(\phi)$ *as defined in Equation* (27)*. Suppose that* ϕ *is a YES instance. Then* $\mathrm{lindisc}_\infty(\mathbf{A}') \le 4/3$*.*

Proof. The claim follows from the completeness argument in [29, Proof of Lemma 5]. ◀

▶ **Lemma 21** (Soundness for Π_2-hardness reduction, [29]). *Let* $\phi(V_A, V_E)$ *be an instance of* $\forall\exists$-NAE-E3-SAT *with* $n' := |V_A|$ *universally quantified variables and* $n := |V_A| + |V_E|$ *variables, and let* $\mathbf{A}' = \mathbf{A}'(\phi)$ *be as defined in Equation* (27)*. Suppose that for every* $\boldsymbol{w}' := (\boldsymbol{w}, \boldsymbol{w}^*)$ *for* $\boldsymbol{w} := 1/2 \cdot \mathbf{1} \in \mathbb{R}^{3n}$ *and* $\boldsymbol{w}^* \in \{1/3, 2/3\}^n$ *there exists a vector* $\boldsymbol{x}' \in \{0,1\}^{3n+n'}$ *such that* $\|\mathbf{A}'(\boldsymbol{w}' - \boldsymbol{x}')\|_\infty < 3/2$*. Then* ϕ *is a YES instance of* $\forall\exists$-NAE-E3-SAT*.*

Proof. The claim follows from the soundness argument in [29, Proof of Lemma 5]. ◀

We now state a lemma that guarantees that if a vector $\boldsymbol{w}'$ of the form in Lemma 21 and $\boldsymbol{x}' \in \mathbb{Z}^{3n+n'}$ satisfy $\|\mathbf{A}'(\boldsymbol{w}' - \boldsymbol{x}')\|_\infty < 3/2$, then it must be the case that $\boldsymbol{x}' \in \{0,1\}^{3n+n'}$.

▶ **Lemma 22.** *Let* $\phi(V_A, V_E)$ *be an* $\forall\exists$-NAE-E3-SAT *instance with* $n' := |V_A|$ *universally quantified variables and* $n := |V_A| + |V_E|$ *variables, and let* $\mathbf{A}' = \mathbf{A}'(\phi)$ *be as defined in Equation* (27)*. Furthermore, let* $\boldsymbol{w} := 1/2 \cdot \mathbf{1} \in \mathbb{R}^{3n}$*, let* $\boldsymbol{w}^* \in \{1/3, 2/3\}^{n'}$*, let* $\boldsymbol{x} \in \mathbb{Z}^{3n}$*, and let* $\boldsymbol{x}^* \in \mathbb{Z}^{n'}$*. Let* $\boldsymbol{w}' := (\boldsymbol{w}, \boldsymbol{w}^*) \in \mathbb{R}^{3n+n'}$ *and let* $\boldsymbol{x}' := (\boldsymbol{x}, \boldsymbol{x}^*) \in \mathbb{Z}^{3n+n'}$*. Assume that* $\|\mathbf{A}'(\boldsymbol{w}' - \boldsymbol{x}')\|_\infty < 3/2$*. Then* $\boldsymbol{x}' \in \{0,1\}^{3n+n'}$*.*

Proof. We have that

$$\max_{j\in[n]}\|\mathbf{G}(\boldsymbol{w}_j-\boldsymbol{x}_j)\|_\infty \le \|\mathbf{A}(\boldsymbol{w}-\boldsymbol{x})\|_\infty \le \|\mathbf{A}'(\boldsymbol{w}'-\boldsymbol{x}')\|_\infty < 3/2 \ , \tag{29}$$

where the first inequality follows from the definition of $\mathbf{A}$ (Equation (8)), the second inequality follows from Lemma 19, and the third inequality is an assumption. Let $\boldsymbol{x}^1, \boldsymbol{x}^2, \boldsymbol{x}^3 \in \mathbb{Z}^n$ be such that $\boldsymbol{x} := (\boldsymbol{x}^2, \boldsymbol{x}^2, \boldsymbol{x}^3)$, and for $j \in [n]$, let $\boldsymbol{x}_j := (x_j^1, x_j^2, x_j^3) \in \mathbb{Z}^3$. By Equation (29) we then have that for all $j \in [n]$, $\|\mathbf{G}(1/2 \cdot \mathbf{1} - \boldsymbol{x}_j)\|_\infty < 3/2$. Therefore, by Lemma 12 we have have that for all $j \in [n]$, $\boldsymbol{x}_j \in \{\mathbf{0}, \mathbf{1}\}$. In particular, $\boldsymbol{x} \in \{0,1\}^{3n}$.

Again, by setting $\boldsymbol{y}' := \boldsymbol{w}' - \boldsymbol{x}'$ in Lemma 19, we have that $(8/3) \cdot \max_{i\in[n']}|w_i^* - x_i^*| < 3/2$, or equivalently, that for all $i \in [n']$, $|w_i^* - x_i^*| < 9/16 < 2/3$. It follows that if $w_i^* = 1/3$ then $x_i^* = 0$, and if $w_i^* = 2/3$ then $x_i^* = 1$. In particular, $\boldsymbol{x}^* \in \{0,1\}^{n'}$, as needed. ◀

Finally, we obtain the following theorem that adapts [29, Lemma 5] to the $\mathrm{BinGapCRP}_\infty$ problem.

▶ **Theorem 23.** *For any constant $\varepsilon > 0$, there is a polynomial-time mapping reduction from* ∀∃-NAE-E3-SAT *to* $(9/8 - \varepsilon)$-$\mathrm{BinGapCRP}_\infty$.

Proof. Let ϕ be an instance of the ∀∃-NAE-E3-SAT problem. The reduction outputs the BinGapCRP instance with matrix $\mathbf{A}' := \mathbf{A}'(\phi)$, as defined in Equation (27), and distance threshold $r := 4/3$. The reduction is efficient by inspection. If ϕ is a YES instance, then $\mathrm{lindisc}_\infty(\mathbf{A}') \le 4/3$ by Lemma 20.

On the other hand, if $\mu_\infty(\mathcal{L}(\mathbf{A}')) < 3/2$, then by definition of the covering radius and Lemma 22 we have that for all $\boldsymbol{w}' := (\boldsymbol{w}, \boldsymbol{w}^*) \in \mathbb{R}^{3n+n'}$ with $\boldsymbol{w} := 1/2 \cdot \mathbf{1}^{3n}$ and $\boldsymbol{w}^* \in \{1/3, 2/3\}^{n'}$, there exists $\boldsymbol{x}' \in \{0,1\}^{3n+n'}$ such that $\|\mathbf{A}'(\boldsymbol{w}' - \boldsymbol{x}')\|_\infty < 3/2$. Therefore, ϕ is a YES instance by Lemma 21. The result follows. ◀

The proof of Theorem 2 now follows easily.

Proof of Theorem 2. The result follows from the Π_2-hardness of ∀∃-NAE-E3-SAT (shown in [10, Theorem 4.7]) and the reduction in Theorem 23. ◀

References

1. Divesh Aggarwal and Noah Stephens-Davidowitz. (gap/s)eth hardness of SVP. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 228–238. ACM, 2018. `doi:10.1145/3188745.3188840`. 5
2. Miklós Ajtai. The shortest vector problem in L_2 is *NP*-hard for randomized reductions (extended abstract). In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 10–19. ACM, 1998. `doi:10.1145/276698.276705`. 2
3. Huck Bennett. The complexity of the Shortest Vector Problem. *SIGACT News*, 54(1):37–61, 2023. `doi:10.1145/3586165.3586172`. 3
4. Huck Bennett and Chris Peikert. Hardness of bounded distance decoding on lattices in p norms. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, Saarbrücken, Germany (Virtual Conference), July 28-31, 2020*, volume 169 of *LIPIcs*, pages 36:1–36:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.CCC.2020.36`. 5

5 Huck Bennett, Chris Peikert, and Yi Tang. Improved hardness of BDD and SVP under gap-(s)eth. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, Berkeley, CA, USA, January 31 - February 3, 2022*, volume 215 of *LIPIcs*, pages 19:1–19:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ITCS.2022.19. 5

6 Joshua Brakensiek, Neng Huang, Aaron Potechin, and Uri Zwick. On the mysteries of MAX NAE-SAT. *SIAM J. Discret. Math.*, 39(1):267–313, 2025. doi:10.1137/23m1591578. 7

7 William Chen, Anand Srivastav, and Giancarlo Travaglini. *A Panorama of Discrepancy Theory*. Springer International Publishing, 2014. doi:10.1007/978-3-319-04696-9. 5

8 Daniel Dadush. On approximating the covering radius and finding dense lattice subspaces. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 1021–1026. ACM, 2019. doi:10.1145/3313276.3316397. 1

9 Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009. Preliminary version in STOC 2006. doi:10.1137/07068062X. 4

10 Janosch Döcker, Britta Dorn, Simone Linz, and Charles Semple. Placing quantified variants of 3-sat and not-all-equal 3-sat in the polynomial hierarchy. *Theor. Comput. Sci.*, 822:72–91, 2020. doi:10.1016/J.TCS.2020.04.003. 17, 19

11 Friedrich Eisenbrand, Dömötör Pálvölgyi, and Thomas Rothvoß. Bin packing via discrepancy of permutations. *ACM Trans. Algorithms*, 9(3):24:1–24:15, 2013. doi:10.1145/2483699.2483704. 5

12 Thomas Eiter and Georg Gottlob. Note on the complexity of some eigenvector problems. Technical Report CD-TR 95/89, Christian Doppler Laboratory for Expert Systems, TU Vienna, 1995. 17

13 Aparna Gupte and Vinod Vaikuntanathan. The fine-grained hardness of sparse linear regression. *CoRR*, abs/2106.03131, 2021. arXiv:2106.03131. 3

14 Venkatesan Guruswami. Inapproximability results for set splitting and satisfiabilityproblems with no mixed clauses. *Algorithmica*, 38(3):451–469, March 2004. doi:10.1007/s00453-003-1072-z. 7, 10

15 Venkatesan Guruswami, Daniele Micciancio, and Oded Regev. The complexity of the covering radius problem on lattices and codes. In *19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA*, pages 161–173. IEEE Computer Society, 2004. doi:10.1109/CCC.2004.1313831. 2, 3, 4, 7

16 Isaac M Hair and Amit Sahai. SVP_p is deterministically NP-hard for all $p > 2$, even to approximate within a factor of $2^{\log^{1-\varepsilon} n}$. Cryptology ePrint Archive, Paper 2025/2181, 2025. URL: https://eprint.iacr.org/2025/2181. 2, 5

17 Isaac M Hair and Amit Sahai. Deterministic hardness of approximation for SVP in all finite ℓ_p norms, 2026. doi:10.48550/arXiv.2604.01451. 2, 5

18 Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. doi:10.1145/502090.502098. 7, 10

19 Ishay Haviv, Vadim Lyubashevsky, and Oded Regev. A note on the distribution of the distance from a lattice. *Discret. Comput. Geom.*, 41(1):162–176, 2009. doi:10.1007/S00454-008-9123-5. 4

20 Ishay Haviv and Oded Regev. Hardness of the covering radius problem on lattices. *Chic. J. Theor. Comput. Sci.*, 2012, 2012. URL: http://cjtcs.cs.uchicago.edu/articles/2012/4/contents.html. 2, 3, 4

21 Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory Comput.*, 8(1):513–531, 2012. doi:10.4086/TOC.2012.V008A023. 2

22 Yahli Hecht and Muli Safra. Deterministic hardness of approximation of unique-svp and gapsvp in ℓ_p norms for $p > 2$, 2025. arXiv:2510.16991. 2

23 Rebecca Hoberg and Thomas Rothvoss. A logarithmic additive integrality gap for bin packing. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 2616–2625. SIAM, 2017. doi:10.1137/1.9781611974782.172. 5

24 Subhash Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005. Preliminary version in FOCS 2004. doi:10.1145/1089023.1089027. 2, 5

25 Subhash Khot. Hardness of approximating the shortest vector problem in high l_p norms. *J. Comput. Syst. Sci.*, 72(2):206–219, 2006. Preliminary version in FOCS 2003. doi:10.1016/J.JCSS.2005.07.002. 2, 5

26 Lily Li and Aleksandar Nikolov. On the Computational Complexity of Linear Discrepancy. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms (ESA 2020)*, volume 173 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 69:1–69:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2020.69. 5

27 L. Lovász, J. Spencer, and K. Vesztergombi. Discrepancy of set-systems and matrices. *European Journal of Combinatorics*, 7(2):151–160, 1986. doi:10.1016/S0195-6698(86)80041-5. 5, 8

28 Alexander Magazinov. A proof of a conjecture by haviv, lyubashevsky and regev on the second moment of a lattice voronoi cell. *Advances in Geometry*, 20(1):117–120, January 2020. doi:10.1515/advgeom-2018-0018. 2, 4

29 Pasin Manurangsi. Linear discrepancy is Π_2-hard to approximate. *Information Processing Letters*, 172:106164, 2021. doi:10.1016/j.ipl.2021.106164. 2, 3, 5, 6, 7, 10, 11, 12, 17, 18, 19

30 Aileen M. McLoughlin. The complexity of computing the covering radius of a code. *IEEE Trans. Inf. Theory*, 30(6):800–804, 1984. doi:10.1109/TIT.1984.1056978. 3

31 Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035, 2000. Preliminary version in FOCS 1998. doi:10.1137/S0097539700373039. 2

32 Daniele Micciancio. Almost perfect lattices, the covering radius problem, and applications to ajtai's connection factor. *SIAM J. Comput.*, 34(1):118–169, 2004. doi:10.1137/S0097539703433511. 1, 4

33 Daniele Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory Comput.*, 8(1):487–512, 2012. doi:10.4086/TOC.2012.V008A022. 2

34 S. Muthukrishnan and Aleksandar Nikolov. Optimal private halfspace counting via discrepancy. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 1285–1292, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2213977.2214090. 5

35 Erez Petrank. The hardness of approximation: Gap location. *computational complexity*, 4(2):133–157, June 1994. doi:10.1007/BF01202286. 7, 10

36 Oded Regev and Ricky Rosen. Lattice problems and norm embeddings. In Jon M. Kleinberg, editor, *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 447–456. ACM, 2006. doi:10.1145/1132516.1132581. 2

37 Thomas Rothvoss. Better bin packing approximations via discrepancy theory. *SIAM J. Comput.*, 45(3):930–946, 2016. doi:10.1137/140955367. 5

38 Daqing Wan. Np-hardness of svp in euclidean space, 2026. doi:10.48550/arXiv.2603.27398. 2