Optimal Inapproximability of Promise Equations over Finite Groups

Silvia Butti
Department of Computer Science, University of Oxford, UK

Alberto Larrauri
Department of Computer Science, University of Oxford, UK

Stanislav Živný
Department of Computer Science, University of Oxford, UK

Abstract

A celebrated result of Håstad established that, for any constant $\varepsilon > 0$, it is NP-hard to find an assignment satisfying a $(1/|\mathcal{G}| + \varepsilon)$-fraction of the constraints of a given 3-LIN instance over an Abelian group $\mathcal{G}$ even if one is promised that an assignment satisfying a $(1 - \varepsilon)$-fraction of the constraints exists. Engebretsen, Holmerin, and Russell showed the same result for 3-LIN instances over any finite (not necessarily Abelian) group. In other words, for almost-satisfiable instances of 3-LIN the random assignment achieves an optimal approximation guarantee. We prove that the random assignment algorithm is still best possible under a stronger promise that the 3-LIN instance is almost satisfiable over an arbitrarily more restrictive group.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Theory of computation → Problems, reductions and completeness

Keywords and phrases promise constraint satisfaction, approximation, linear equations

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.13

Category Track A: Algorithms, Complexity and Games

Related Version *Full version*: `https://arxiv.org/abs/2411.01630`

Funding *Silvia Butti*: UKRI EP/X024431/1
Alberto Larrauri: UKRI EP/X024431/1
Stanislav Živný: UKRI EP/X024431/1

Acknowledgements We thank the anonymous reviewers for their feedback.

1 Introduction

The PCP theorem [3, 2, 20] is one of the jewels of computational complexity and theoretical computer science more broadly [1]. One of its equivalent statements is as follows: The maximum number of simultaneously satisfiable constraints of a Constraint Satisfaction Problem, or CSP for short, is NP-hard to approximate within some constant factor. That is, while NP-hardness of CSPs means that it is NP-hard to distinguish instances that are satisfiable from those that are unsatisfiable, the PCP theorem shows that there is an absolute constant $\alpha < 1$ such that it is NP-hard to distinguish satisfiable CSP instances from those in which strictly fewer than an α-fraction of the constraints can be simultaneously satisfied. Thus it is NP-hard to find an assignment that satisfies an α-fraction of the constraints even if one is promised that a satisfying assignment exists. For some CSPs, as we shall see shortly, the optimal value of α is known.

A classic example of a CSP is 3-SAT, the satisfiability problem of CNF-formulas in which each clause contains 3 literals. The random assignment gives a method to find an assignment that satisfies a 7/8-fraction of the clauses. Håstad famously showed that this is optimal in

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).
Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joel Ouaknine, and Gabriele Puppis; Article No. 13; pp. 13:1–13:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the following sense: For any constant $\varepsilon > 0$, it is NP-hard to find an assignment satisfying a $(7/8+\varepsilon)$-fraction of the clauses of a 3-SAT instance even if one is promised that a satisfying assignment exists [25].

Another classic CSP is 3-LIN, the problem of solving linear equations in 3 variables over the Boolean domain $\{0,1\}$. If all equations can be satisfied simultaneously then a satisfying assignment can be found in polynomial time by Gaussian elimination. What can be done if no satisfying assignment exists? As for 3-SAT, the random assignment gives a method to find a somewhat satisfying assignment, namely one that satisfies a 1/2-fraction of the constraints. As it turns out, this is best possible even for instances of 3-LIN that are almost satisfiable. In detail, Håstad showed that for any constant $\varepsilon > 0$, it is NP-hard to find an assignment satisfying a $(1/2+\varepsilon)$-fraction of the constraints of a 3-LIN instance even if one is promised that an assignment satisfying a $(1-\varepsilon)$-fraction of the constraints exists. In fact, Håstad established optimal inapproximability results for 3-LIN over any finite Abelian group, not just $\{0,1\}$. This result was later extended by Engebretsen, Holmerin, and Russell to all finite groups [23]. Since these foundational works, Guruswami and Raghavendra [24] showed NP-hardness of finding a barely satisfying assignment for a 3-LIN instance over the reals (and thus also over the integers) even if a nearly satisfying assignment is promised to exist over the integers. The same result was later established for 2-LIN for large enough cyclic groups [35]. Khot and Moshkovitz [28] studied inapproximability of 3-LIN over the reals.

In this work, we strengthen the optimal inapproximability results for 3-LIN over finite groups by establishing NP-hardness of beating the random assignment threshold even if the instance is almost satisfiable in an arbitrarily more restrictive setting. Formally, this is captured by fixing (not one but) two groups and a homomorphism between them, following the framework of promise CSPs [5, 8]. In detail, (decision) promise CSPs [8] can be seen as a qualitative form of approximation: Each constraint comes in two forms, a strong one and a weak one. The promise is that there is a solution satisfying all constraints in the strong form while the (potentially easier) goal is to find a solution satisfying all constraints in the weak form. An example of a strong vs. weak constraint on the same, say Boolean, domain is 1-in-3 vs NAE, where the former is

$$\{(0,0,1),(0,1,0),(1,0,0)\}$$

and the latter is

$$\{(0,0,1),(0,1,0),(1,0,0),(1,1,0),(1,0,1),(0,1,1)\}.$$

NAE is weaker as the relation contains more tuples. While these two constraint relations capture the well-known NP-hard problems of 1-in-3-SAT and Not-All-Equal-SAT respectively [38], finding an NAE-assignment turns out to be doable in polynomial time under the promise that a 1-in-3-assignment exists [15]![1]

For constraints on different domains, the notion of strong vs. weak constraint is captured by a homomorphism between the (sets of all) constraint relations; in the example above, the homomorphism is just the identity function. The exact solvability of 3-LIN in the promise setting was resolved in [32].

Recent work of Barto et al. [9] considered (quantitative) approximation of promise CSPs. In the context of 3-LIN, here are two simple examples captured by this framework. First,

[1] The algorithm involves solving an instance of 3-LIN over the integers and rounding positive integers to 1 and non-positive integers to 0, demonstrating the importance of 3-LIN among promise CSPs.

let $\mathcal{G}$ be a group and $\mathcal{H}$ be a subgroup of $\mathcal{G}$. Given an almost-satisfiable system over the subgroup $\mathcal{H}$, maximise the number of satisfied equations over $\mathcal{G}$. Our results imply that beating the random assignment over $\mathcal{H}$ is NP-hard. In the second example, consider a group $\mathcal{G}$, a normal subgroup $\mathcal{H}$, and an almost-satisfiable system over $\mathcal{G}$. The goal this time is to maximise the number of satisfied equations in the system over the quotient $\mathcal{G}/\mathcal{H}$. Our results show that doing better than the random assignment over $\mathcal{G}/\mathcal{H}$ is NP-hard. More generally, going beyond subgroups and quotients of a given group, we fix two groups $\mathcal{G}_1$ and $\mathcal{G}_2$ and a group homomorphism φ from a subgroup $\mathcal{H}_1$ of $\mathcal{G}_1$ to a subgroup $\mathcal{H}_2$ of $\mathcal{G}_2$ with the property that φ extends to a group homomorphism from $\mathcal{G}_1$ to $\mathcal{G}_2$. Given an almost-satisfiable system of equations over $\mathcal{G}_1$ with constants in $\mathcal{H}_1$, the goal is to maximise the number of satisfied equations over $\mathcal{G}_2$ where the constants are interpreted in $\mathcal{H}_2$ via φ. Our main result establishes that doing better than the random assignment over $\mathcal{H}_2$ is NP-hard, cf. Theorem 3. Thus we give an optimal inapproximability result for a natural and fundamental fragment of promise CSPs, systems of linear equations.

Other fragments of promise CSPs whose quantitative approximation has been studied includes almost approximate graph colouring [22, 21, 29, 26], approximate colouring [33], and approximate graph homomorphism [34]. Recent work of Brakensiek, Guruswami, and Sandeep [16] studied robust approximation of promise CSPs; in particular, they observed that Raghavendra's celebrated theorem on approximate CSPs [36] applies to promise CSPs, which in combination with the work of Brown-Cohen and Raghavendra [17] gives an alternative framework for studying quantitative approximation of promise CSPs.

The general approach for establishing inapproximability of systems of equations, going back to [25, 23], can be seen as a reduction from another CSP that is hard to approximate. In this reduction, one initially transforms an instance of the original CSP to a system of equations of the form $xyz = 1$. To guarantee the soundness of this reduction, one needs to show that any assignment that beats the random assignment in the target system of equations can be transformed into a "good" assignment of the original instance. To do this it is necessary to rule out vacuous assignments (e.g., the assignment that sends all variables to the group identity) through a procedure called folding, which introduces constants in the system of equations. Afterwards, the soundness bounds are shown by performing Fourier analysis on certain functions derived from the system. Our proof follows this general approach. The main obstacle to applying the techniques of [23] directly is the fact that in our setting the constants lie in a proper subgroup of the ambient group, which precludes us from applying classical folding over groups. Instead, we use a weaker notion of folding. This, however, implies that in the soundness analysis we have to take care of functions whose Fourier expansion has non-zero value for the trivial term. To tackle this issue, we consider the behaviour of irreducible group representations when they are restricted to the subgroup of constants via Frobenius Reciprocity. We note that, as in the non-promise setting, the proof in the Abelian case is much simpler, particularly using that the set of characters of an Abelian group corresponds to its Pontryagin dual. Thus, much of the complicated machinery in the paper is to obtain the main result for all groups.

Before formal description of our results, we mention other related work. First, extending the work from [25], Austrin, Brown-Cohen, and Håstad established optimal inapproximability of 3-LIN over Abelian groups with a universal factor graph [4]. Similarly, Bhangale and Stankovic established optimal inapproximability of 3-LIN over non-Abelian groups with a universal factor graph [14]. Second, unlike over Abelian groups, for 3-LIN over non-Abelian groups finding a satisfying assignment is NP-hard even under the promise that one exists. There is a folklore randomised algorithm for satisfiable 3-LIN instances over non-Abelian groups (whose approximation factor depends on the group $\mathcal{G}$ and is $1/|\mathcal{G}|$ if $\mathcal{G}$ is a so-called perfect group but can beat the naive random assignment for non-perfect groups). Bhangale and Khot showed that this algorithm is optimal [10], spawning a number of follow-up works going beyond 3-LIN, e.g. [11, 12, 13]. Third, going beyond 3-LIN, building on a long line of work Chan established optimal (up to a constant factor) NP-hardness for CSPs [19]. There are other works on various inapproximability notions for CSPs, e.g., [6, 30, 31]. We note that As in the non-promise setting, the proof in the Abelian case is much simpler, particularly using that the set of characters of an Abelian group corresponds to its Pontryagin dual. Thus, much of the complicated machinery in the paper could be avoided. Finally, we mention that Khot's influential Unique Games Conjecture [27] postulates, in one of its equivalent forms, NP-hardness of finding a barely satisfying solution to a 2-LIN instance given that an almost-satisfying assignment exists (for a large enough domain size).

1.1 Preliminaries and notation

We use $[\![\cdot]\!]$ to denote the Iverson bracket; i.e., $[\![P]\!]$ is 1 if P is true and 0 otherwise. As usual, $[n]$ denotes the set $\{1, 2, \ldots, n\}$.

We consider matrices whose sets of indices are arbitrary finite sets. Given two finite sets N and M, an $N \times M$ complex matrix A consists of a family of complex numbers $A_{i,j}$ indexed by pairs $i \in N$, $j \in M$. Algebraic notions such as matrix product, trace, and transpose are defined in the natural way. Given an $N_1 \times N_2$ complex matrix A, and an $M_1 \times M_2$ complex matrix B, the *tensor product* $A \otimes B$ is an $(N_1 \times M_1) \times (N_2 \times M_2)$ matrix, where $(A \otimes B)_{(i,s)(j,t)} = A_{i,j}B_{s,t}$ for each $i \in N_1, j \in N_2, s \in M_1, t \in M_2$. The group of invertible $N \times N$ complex matrices (equipped with matrix multiplication and matrix inversion) is denoted by $\mathrm{GL}(N)$, and the set of $N \times M$ complex matrices is denoted by $\mathbb{C}^{N \times M}$.

Let X and Y be sets. We identify tuples $\mathbf{y} \in Y^X$ with functions $\mathbf{y} : X \to Y$, where the x^{th} component of $\mathbf{y}$ is given by $\mathbf{y}(x)$. Composition is defined from left to right in a natural way, i.e., if $\mathbf{y} \in Y^X$ and $\mathbf{z} \in Z^Y$, then $\mathbf{z} \circ \mathbf{y} \in X^Z$ (also denoted just by $\mathbf{zy}$ when there is no ambiguity) is defined by $(\mathbf{z} \circ \mathbf{y})(x) = \mathbf{z}(\mathbf{y}(x))$ for each $x \in X$.

A subset $\mathcal{H} \subseteq \mathcal{G}$ of a group $\mathcal{G}$ is called a *subgroup* of $\mathcal{G}$, denoted by $\mathcal{H} \leq \mathcal{G}$, if $\mathcal{H}$ equipped with the group operation of $\mathcal{G}$ forms a group. Given a group $\mathcal{G}$, a subgroup $\mathcal{H}$ of $\mathcal{G}$, and an element $g \in \mathcal{G}$, the *right coset of* $\mathcal{H}$ *in* $\mathcal{G}$ *by* g is the set $\mathcal{H}g := \{hg \mid h \in \mathcal{H}\}$. The set of right cosets of $\mathcal{H}$ in $\mathcal{G}$ is denoted by $\mathcal{H}\backslash\mathcal{G}$. Let N be a finite set. The N^{th} *direct power* of $\mathcal{G}$, denoted by $\mathcal{G}^N$, is the group whose elements are N-tuples $\mathbf{g} \in \mathcal{G}^N$ of elements from $\mathcal{G}$, and where the group operation is taken component-wise, i.e., $\mathbf{g} \cdot \mathbf{h}(n) = \mathbf{g}(n) \cdot \mathbf{h}(n)$ for each $n \in N$. If $\mathcal{H} \leq \mathcal{G}$, we define $(h \cdot \mathbf{g})(n) = h \cdot \mathbf{g}(n)$ for each $h \in \mathcal{H}$ and $\mathbf{g} \in \mathcal{G}^N$. With this notation, the notion of coset extends to include the right cosets of $\mathcal{H}$ in $\mathcal{G}^N$ in a natural way.

A *homomorphism* from a group $\mathcal{G}_1$ to a group $\mathcal{G}_2$ is a map $\varphi : \mathcal{G}_1 \to \mathcal{G}_2$ which satisfies that $\varphi(g \cdot h) = \varphi(g) \cdot \varphi(h)$ for every $g, h \in \mathcal{G}_1$. The domain and image of φ are denoted $\mathrm{Dom}(\varphi)$ and $\mathrm{Im}(\varphi)$ respectively. Let N be a finite set, $\mathcal{G}_i$ groups, $i \in [2]$, $\mathcal{H}_i \leq \mathcal{G}_i$, and $\varphi : \mathcal{H}_1 \to \mathcal{H}_2$ be a homomorphism. We say that a function $f : \mathcal{G}_1^N \to \mathcal{G}_2$ is *folded over* φ if $f(h\mathbf{g}) = \varphi(h)f(\mathbf{g})$ for all $h \in \mathcal{H}_1$ and $\mathbf{g} \in \mathcal{G}_1^N$. Given an arbitrary function $f : \mathcal{G}_1^N \to \mathcal{G}_2$ and a homomorphism between subgroups, there is a natural way to construct a folded function that resembles f. Fix an arbitrary representative from each right coset of $\mathcal{H}_1$ in $\mathcal{G}_1^N$. For each $\mathbf{g} \in \mathcal{G}^N$, denote by $\mathbf{g}^\dagger$ the representative of $\mathcal{H}_1\mathbf{g}$, and let $h_\mathbf{g} \in \mathcal{H}_1$ be such that $\mathbf{g}^\dagger = h_\mathbf{g}\mathbf{g}$. Then the *folding of* f *over* φ (with respect to this choice of representatives) is the map $f_\varphi : \mathcal{G}_1^N \to \mathcal{G}_2$ given by $f_\varphi(\mathbf{g}) = \varphi(h_\mathbf{g}^{-1})f(\mathbf{g}^\dagger)$.

Fix a pair of disjoint finite sets D, E, called the *label sets*, and a subset $\Pi \subseteq E^D$ of *labeling functions*. An instance of the *Label Cover* problem is a bipartite graph with vertex set $U \sqcup V$ and a labeling function $\pi_{uv} \in \Pi$ for each edge $\{u, v\}$ in the graph. The task is to decide whether there is a pair of assignments $h_D : U \to D$, $h_E : V \to E$ that satisfies all the constraints, i.e., such that $\pi_{uv}(h_D(u)) = h_E(v)$ for each edge $\{u, v\}$.

Given additionally a pair of rational constants $0 < s \le c \le 1$, the gap version of this problem, known as the *Gap Label Cover* problem with completeness c and soundness s and denoted $\mathrm{GLC}_{D,E}(c, s)$, is the problem of distinguishing instances where a c-fraction of the constraints can be satisfied from instances where not even an s-fraction of the constraints can be satisfied.

The hardness of Gap Label Cover with perfect completeness stated below is a consequence of the PCP theorem [2, 3] and the Parallel Repetition Theorem [37].

▶ **Theorem 1.** *For every $\alpha > 0$ there exist finite sets D, E such that $\mathrm{GLC}_{D,E}(1, \alpha)$ is NP-hard.*

Fourier Analysis

We follow closely [39] for our main definitions and preliminary results. A *representation* of a group $\mathcal{G}$ is a group homomorphism $\gamma : \mathcal{G} \to \mathrm{GL}(N_\gamma)$ for some finite set N_γ. We call $|N_\gamma|$ the *dimension* of γ and write $\dim_\gamma = |N_\gamma|$. Given a pair of indices $i, j \in N_\gamma^2$, $\gamma_{i,j}$ denotes the (i, j)-th entry of γ. The *character* of a representation γ, denoted by χ_γ, is its trace. The *trivial representation*, denoted 1, maps all group elements to the number one (i.e., the one-dimensional identity matrix). A representation γ is said to be *unitary* if its image contains only unitary matrices.

We say that two representations α and β of some group $\mathcal{G}$ are equivalent, written $\alpha \simeq \beta$, if there is an invertible $N_\beta \times N_\alpha$ complex matrix T such that $\alpha(g) = T^{-1}\beta(g)T$ for all $g \in \mathcal{G}$. In particular, $\dim_\alpha = \dim_\beta$. Similarly, the representation β is said to be a *sub-representation* of α if there is an invertible matrix T, such that $T^{-1}\alpha(g)T$ can be written as

$$\begin{pmatrix} \beta(g) & * \\ 0 & * \end{pmatrix}$$

for all $g \in \mathcal{G}$. The representation β is said to be *irreducible* if all its sub-representations are equivalent to itself. If β is irreducible, its *multiplicity* in α is the non-negative integer n satisfying that α is equivalent to a block diagonal representation with two diagonal blocks α_1, α_2, where (1) α_1 is another block-diagonal representation consisting of n diagonal blocks equal to β, and (2) α_2 does not have β as a sub-representation.

Given a group $\mathcal{G}$, we use $\widehat{\mathcal{G}}$ to denote some arbitrary and fixed complete set of inequivalent irreducible unitary representations of $\mathcal{G}$; such a set exists by, e.g., [39, Proposition 1].

The space $\mathcal{L}^2(\mathcal{G})$ is the vector space of complex-valued functions over $\mathcal{G}$, equipped with the following inner product:[2]

$$\langle F, H \rangle = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} F(g)\overline{H(g)}.$$

Let $\mathcal{G}$ be a group, and let $F : \mathcal{G} \to \mathbb{C}$ be a complex-valued function. Given $\gamma \in \widehat{\mathcal{G}}$ and $i, j \in N_\gamma$, the *Fourier coefficient* $\widehat{F}(\gamma_{i,j})$ is defined as the product $\langle F, \gamma_{i,j} \rangle$. The matrix entries of the representations $\gamma \in \widehat{\mathcal{G}}$ form an orthogonal basis of $\mathcal{L}^2(\mathcal{G})$, and allow us to perform Fourier analysis on this space, as stated in the following theorem [39, Theorem 2].

[2] Note the additional normalising factor of $\frac{1}{|\mathcal{G}|}$ compared to [39].

▶ **Theorem 2.** *Let $\mathcal{G}$ be a finite group. Then the set*

$$\{\gamma_{i,j} \mid \gamma \in \widehat{\mathcal{G}},\ i,j \in N_\gamma\}$$

is an orthogonal basis of $\mathcal{L}^2(\mathcal{G})$, and $\dim_\gamma \|\gamma_{i,j}\|^2 = 1$ for all $\gamma_{i,j}$. Moreover, the following hold:

1. Plancherel's Theorem: *Given $F \in \mathcal{L}^2(\mathcal{G})$,*

$$\|F\|^2 = \sum_{\gamma\in\widehat{\mathcal{G}}, i,j\in N_\gamma} \dim_\gamma |\widehat{F}(\gamma_{i,j})|^2.$$

2. Fourier Inversion: *Given $F \in \mathcal{L}^2(\mathcal{G})$,*

$$F(g) = \sum_{\gamma\in\widehat{\mathcal{G}}, i,j\in N_\gamma} \dim_\gamma \widehat{F}(\gamma_{i,j})\gamma_{i,j}(g) \qquad \text{for all } g \in \mathcal{G}.$$

We also consider Fourier transforms of matrix-valued functions $F : \mathcal{G} \to \mathbb{C}^{N_F\times N_F}$. Given $\gamma \in \widehat{\mathcal{G}}$ and indices $i, j \in N_\gamma$, we define the $N_F \times N_F$ matrix $\widehat{F}(\gamma_{i,j})$ as the one whose (s,t)-th entry is $\widehat{F_{s,t}}(\gamma_{i,j})$ for each $s, t \in N_F$. In other words,

$$\widehat{F}(\gamma_{i,j}) = \frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} F(g)\overline{\gamma_{i,j}(g)}.$$

Let N be a finite set. Given a pair of functions function $F, H : \mathcal{G} \to \mathbb{C}^{N\times N}$, we define their *convolution* $F * H$ by

$$(F * H)(g) := \frac{1}{|\mathcal{G}|}\sum_{h\in\mathcal{G}} F(h)H(h^{-1}g).$$

We will also need to perform Fourier analysis over powers of the form $\mathcal{G}^D$ for a given group $\mathcal{G}$ and finite set D. It is possible to identify $\widehat{\mathcal{G}^D}$ with $(\widehat{\mathcal{G}})^D$ [39]. This way, an element $\rho \in \widehat{\mathcal{G}^D}$ is given by a tuple $(\rho^d)_{d\in D}$ where $\rho^d \in \widehat{\mathcal{G}}$ for each $d \in D$ in such a way that

$$\rho(\mathbf{g}) = \bigotimes_{d\in D} \rho^d(\mathbf{g}(d))$$

for all $\mathbf{g} \in \mathcal{G}^D$. Observe we use superscripts for the "components" of the representation ρ on the power group $\mathcal{G}^D$, rather than subscripts, which we utilise to denote matrix entries. The *degree* of ρ, written $|\rho|$, is the number of indices $d \in D$ for which ρ^d is non-trivial.[3]

1.2 Results

Let $\mathcal{G}_1, \mathcal{G}_2$ be two groups and φ a group homomorphism with domain $\mathrm{Dom}(\varphi) \le \mathcal{G}_1$ and image $\mathrm{Im}(\varphi) \le \mathcal{G}_2$ that extends to a full homomorphism from $\mathcal{G}_1$ to $\mathcal{G}_2$. We shall refer to triples $(\mathcal{G}_1, \mathcal{G}_2, \varphi)$ of this kind as *templates*. Further, let $0 < s \le c \le 1$ be rational constants. We consider the problem 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, c, s)$ which asks, given a weighted system of linear equations with exactly three variables in each equation and constants in $\mathrm{Dom}(\varphi)$ that is c-satisfiable in $\mathcal{G}_1$, to decide whether there exists an s-approximation in $\mathcal{G}_2$, where the constants are interpreted through φ.

[3] This quantity is called "weight" in [23, 14].

To be more precise, an instance to 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, c, s)$ over a set of variables X is a weighted systems of linear equations where each equation is of the form

$$x^i y^j z^k = g$$

for some $x, y, z \in X$, $g \in \mathrm{Dom}(\varphi)$, $i, j, k \in \{-1, 1\}$, and each equation has a non-negative rational weight. Without loss of generality, we assume that the weights are normalised, i.e., sum up to 1. For $t \in [2]$, an assignment $f : X \to \mathcal{G}_t$ *satisfies* an equation $x^i y^j z^k = g$ in $\mathcal{G}_t$ if $f(x)^i f(y)^j f(z)^k = g$ for $t = 1$, and $f(x)^i f(y)^j f(z)^k = \varphi(g)$ for $t = 2$. The task then is to accept if there is an assignment that satisfies a c-fraction (i.e., a fraction of total weight c) of equations in $\mathcal{G}_1$, and to reject if there is no assignment that satisfies an s-fraction of the equations in $\mathcal{G}_2$. It is easy to verify that, if $(\mathcal{G}_1, \mathcal{G}_2, \varphi)$ is a template and $s \le c$, then the sets of accept and reject instances are, in fact, disjoint.[4]

3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, c, s)$ is trivially tractable when $\mathrm{Im}(\varphi) = \{1\}$, so we focus on the case where $|\mathrm{Im}(\varphi)| \ge 2$. The main result of this paper is that 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, 1-\epsilon, 1/|\mathrm{Im}(\varphi)| + \delta)$ is NP-hard for all $\epsilon, \delta > 0$ for which the problem is well-defined. This is achieved by a reduction from the Gap Label Cover problem with perfect completeness and soundness $\alpha = \delta^2/(4\kappa|\mathcal{G}_1|^\kappa|\mathcal{G}_2|^4)$, where $\kappa = \lceil(\log_2 \delta - 2)/(\log_2(1-\epsilon))\rceil$.

▶ **Theorem 3 (Main).** *Let ϵ, δ be positive constants satisfying $1 - \epsilon \ge 1/|\mathrm{Im}(\varphi)| + \delta$. Then,* 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, 1-\epsilon, 1/|\mathrm{Im}(\varphi)| + \delta)$ *is NP-hard.*

The hardness result in Theorem 3 is tight for many, but perhaps surprisingly not all, templates. We call a template $(\mathcal{G}_1, \mathcal{G}_2, \varphi)$ *cubic* if for every $h \in \mathrm{Im}(\varphi)$ there is an element $g \in \mathcal{G}_2$ satisfying $g^3 = h$. Theorem 3 is tight for cubic templates. Indeed, for these templates, the random assignment over $\mathrm{Im}(\varphi)$ achieves a $1/|\mathrm{Im}(\varphi)|$ expected fraction of satisfied equations (and this can be derandomised, e.g., by the method of conditional expectations).

▶ **Theorem 4.** *Let $(\mathcal{G}_1, \mathcal{G}_2, \varphi)$ be a cubic template and $0 < s \le c < 1$. Then the following holds:* 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, c, s)$ *is tractable if $s \le 1/|\mathrm{Im}(\varphi)|$ and NP-hard otherwise.*

Let us now turn to non-cubic templates. An equation is *unsatisfiable* if it is of the form $x^3 = h$ or $x^{-3} = h$ for some $h \in \mathrm{Dom}(\varphi)$ such that $g^3 \ne \varphi(h)$ for all $g \in \mathcal{G}_2$.[5] Note that a template has unsatisfiable equations if and only if it is non-cubic. Note that the naive random assignment cannot achieve a positive approximation factor in systems of equations over non-cubic templates since the system could consist exclusively of unsatisfiable equations. However, there is a simple algorithm for 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, c, c/|\mathrm{Im}(\varphi)|)$ that works even for non-cubic templates, which we describe next.

Given a weighted system of equations over $(\mathcal{G}_1, \mathcal{G}_2, \varphi)$, consider its set of unsatisfiable equations. Since φ extends to a full homomorphism, if the total weight of the set of unsatisfiable equations is more than $1 - c$, then the instance cannot be c-satisfiable in $\mathcal{G}_1$, hence, reject. Otherwise, the random assignment over $\mathrm{Im}(\varphi)$ satisfies at least a $1/|\mathrm{Im}(\varphi)|$-fraction of the satisfiable equations over $\mathcal{G}_2$, which is at least a $c/|\mathrm{Im}(\varphi)|$-fraction of the entire system. It is a simple corollary of Theorem 3 that this algorithm is optimal for non-cubic groups, leading to the following result. Details are deferred to the full version of this paper [18].

▶ **Theorem 5.** *Let $(\mathcal{G}_1, \mathcal{G}_2, \varphi)$ be a non-cubic template and $0 < s \le c < 1$. Then,* 3-LIN$(\mathcal{G}_1, \mathcal{G}_2, \varphi, c, s)$ *is tractable if $s/c \le 1/|\mathrm{Im}(\varphi)|$ and NP-hard otherwise.*

---

[4] 3-LIN can be alternatively phrased as a promise constraint satisfaction problem, cf. [18] for details.

[5] Note that, since φ extends to a homomorphism from $\mathcal{G}_1$ to $\mathcal{G}_2$, this also implies that $g^3 \ne h$ for all $g \in \mathcal{G}_1$.

The structure of the paper is as follows. The rest of this section gives a sketch of the main proof: In Section 1.3 we present the reduction from Gap Label Cover to 3-LIN($\mathcal{G}_1, \mathcal{G}_2, \varphi, 1-\epsilon, 1/|\mathrm{Im}(\varphi)|+\delta$), and in Section 1.4 we give an overview of the techniques used in the analysis of this reduction and of the main challenges that arise in extending previous work to the promise setting. The rest of the paper then gives the main ideas of the technical details. All details are deferred to the full version of this paper [18], which also relates our results to a recent theory of Barto et al. [9], who developed a systematic approach to study (in)approximability of promise CSPs, which includes approximability of promise linear equations, from the viewpoint of universal algebra. In particular, we show in [18] that the proof of Theorem 3 implies that the collection of symmetries[6] of 3-LIN($\mathcal{G}_1, \mathcal{G}_2, \varphi, 1-\epsilon, 1/|\mathrm{Im}(\varphi)|+\delta$) can be mapped homomorphically to the collection of symmetries of Gap Label Cover, a condition that, based on the algebraic theory from [9], is known to guarantee NP-hardness of the former problem.

1.3 Reduction

For the rest of the section we outline the proof of our main result, Theorem 3. From now on we fix a template ($\mathcal{G}_1, \mathcal{G}_2, \varphi$), and positive constants $\delta, \epsilon > 0$ with $1/|\mathrm{Im}(\varphi)| + \delta \leq 1 - \epsilon$. We define $\mathcal{H}_1 = \mathrm{Dom}(\varphi) \leq \mathcal{G}_1$ and $\mathcal{H}_2 = \mathrm{Im}(\varphi) \leq \mathcal{G}_2$.

Our proof follows from a reduction from $\mathrm{GLC}_{D,E}(1, \alpha)$ where

$$\alpha = \frac{\delta^2}{4\kappa|\mathcal{G}_1|^\kappa|\mathcal{G}_2|^4}, \quad \kappa = \left\lceil \frac{\log_2 \delta - 2}{\log_2(1-\epsilon)} \right\rceil,$$

and D, E are chosen to be large enough so that $\mathrm{GLC}_{D,E}(1, \alpha)$ is NP-hard by the PCP theorem [2, 3, 37] (cf. Theorem 1). This reduction constructs an instance Φ_Σ of 3-LIN($\mathcal{G}_1, \mathcal{G}_2, \varphi, 1-\epsilon, 1/|\mathcal{H}_2|+\delta$) for any given instance Σ of Gap Label Cover as described below.

Let $U \sqcup V$ be the underlying vertex set of Σ, D, E be the disjoint sets of labels, and π_{uv} be the labeling functions. We fix representatives from each right coset in $\mathcal{H}_1 \backslash \mathcal{G}_1^D$ and $\mathcal{H}_1 \backslash \mathcal{G}_1^E$. Given a tuple $\mathbf{x}$ in either $\mathcal{G}_1^D$ or $\mathcal{G}_1^E$ we write $\mathbf{x}^\dagger$ for the representative of the coset $\mathcal{H}_1\mathbf{x}$. Let $X = \{u_\mathbf{b} \,|u \in U, \mathbf{b} \in \mathcal{G}_1^D\} \sqcup \{v_\mathbf{a} \,|v \in V, \mathbf{a} \in \mathcal{G}_1^E\}$. Then Φ_Σ is the weighted system of equations over X that contains the equation

$$v_{\mathbf{a}^\dagger} u_{\mathbf{b}^{s_1}}^{s_1} u_{\mathbf{c}^{s_2}}^{s_2} = g_\mathbf{a} \tag{1}$$

for each edge $\{u, v\}$ of Σ, $\mathbf{a} \in \mathcal{G}_1^E$, $\mathbf{b} \in \mathcal{G}_1^D$, $s_1, s_2 \in \{-1, 1\}$, where $\mathbf{c}$ stands for $\mathbf{b}^{-1}(\mathbf{a} \circ \pi_{uv})^{-1}\boldsymbol{\nu}$ and $\boldsymbol{\nu} \in \mathcal{G}_1^D$ is a small perturbation factor. The element $g_\mathbf{a}$ is chosen so that $\mathbf{a}^\dagger = g_\mathbf{a}\mathbf{a}$. The weight of this equation in Φ_Σ is the joint probability of the independent events described in Figure 1.

Let us describe assignments of Φ_Σ over $\mathcal{G}_i$ for $i = 1, 2$. Formally, an assignment of Φ_Σ over $\mathcal{G}_i$ is a map $h : X \to \mathcal{G}_i$. Such an assignment can be described by two families of maps $A = (A_v)_{v \in V}$ from $\mathcal{G}_1^E$ to $\mathcal{G}_i$ and $B = (B_u)_{u \in U}$ from $\mathcal{G}_1^D$ to $\mathcal{G}_i$ by letting $A_v(\mathbf{a}) = h(v_\mathbf{a})$ for all $v \in V, \mathbf{a} \in \mathcal{G}_1^E$, and $B_u(\mathbf{b}) = h(u_\mathbf{b})$ for all $u \in U, \mathbf{b} \in \mathcal{G}_1^D$. It will be more convenient to talk about the pair (A, B) rather than the map h itself, so we will write $\Phi_\Sigma^{\mathcal{G}_i}(A, B)$ to refer to the proportion of equations satisfied by the assignment h. Let us give a more useful expression for $\Phi_\Sigma^{\mathcal{G}_i}(A, B)$. When $i = 1$, we can write

$$\Phi_\Sigma^{\mathcal{G}_1}(A, B) = \mathbb{E}_{\substack{uv,\mathbf{a},\mathbf{b},\\ \boldsymbol{\nu},s_1,s_2}} \left[[\![A_v(\mathbf{a}^\dagger) B_u(\mathbf{b}^{s_1})^{s_1} B_u((\mathbf{b}^{-1}(\mathbf{a} \circ \pi_{uv})^{-1}\boldsymbol{\nu})^{s_2})^{s_2} = g_\mathbf{a}]\!] \right],$$

[6] called the *valued minion of plurimorphisms* in [9].

(1)	The edge $\{u, v\}$ is chosen uniformly at random among all edges of Σ.
(2)	The elements $\mathbf{a}$ and $\mathbf{b}$ are chosen uniformly at random from $\mathcal{G}_1^E$ and $\mathcal{G}_1^D$ respectively.
(3)	The element $\boldsymbol{\nu} \in \mathcal{G}_1^D$ is chosen so that for each $d \in D$, independently, $\boldsymbol{\nu}(d) = 1_{\mathcal{G}_1}$ with probability $1 - \epsilon$, and $\boldsymbol{\nu}(d)$ is selected uniformly at random from $\mathcal{G}_1$ with probability ϵ.
(4)	The signs s_1, s_2 are chosen uniformly at random from $\{-1, 1\}$.

Figure 1 The sampling procedure for Φ_Σ.

where the expectation is taken over the probabilities described in Figure 1, and we use uv as a shorthand for an edge $\{u, v\}$. Folding the assignments A_v over the identity on $\mathcal{H}_1$ and using the fact that $(A_v)_{\mathrm{id}_{\mathcal{H}_1}}(\mathbf{a}) = g_{\mathbf{a}}^{-1} A_v(\mathbf{a}^\dagger)$, we obtain

$$\Phi_\Sigma^{\mathcal{G}_1}(A, B) = \mathbb{E}_{\substack{uv,\mathbf{a},\mathbf{b},\\ \boldsymbol{\nu},s_1,s_2}} \left[[\![(A_v)_{\mathrm{id}_{\mathcal{H}_1}}(\mathbf{a}) B_u(\mathbf{b}^{s_1})^{s_1} B_u((\mathbf{b}^{-1}(\mathbf{a} \circ \pi_{uv})^{-1}\boldsymbol{\nu})^{s_2})^{s_2} = 1_{\mathcal{G}_1}]\!] \right]. \tag{2}$$

Analogously, when $i = 2$ and A_v, B_u are families of maps to $\mathcal{G}_2$, we obtain a similar expression for $\Phi_\Sigma^{\mathcal{G}_2}(A, B)$:

$$\Phi_\Sigma^{\mathcal{G}_2}(A, B) = \mathbb{E}_{\substack{uv,\mathbf{a},\mathbf{b},\\ \boldsymbol{\nu},s_1,s_2}} \left[[\![(A_v)_{\varphi}(\mathbf{a}) B_u(\mathbf{b}^{s_1})^{s_1} B_u((\mathbf{b}^{-1}(\mathbf{a} \circ \pi_{uv})^{-1}\boldsymbol{\nu})^{s_2})^{s_2} = 1_{\mathcal{G}_2}]\!] \right]. \tag{3}$$

That is, a pair of assignments (A, B) satisfies an equation in Φ_Σ if and only if the corresponding pair of assignments obtained by folding A (over $\mathrm{id}_{\mathcal{H}_1}$ and φ respectively) maps the equation to the group identity (respectively, in $\mathcal{G}_1$ and $\mathcal{G}_2$). Thus, folding allows us to focus exclusively on the identity terms in these expectations, which will be useful in the analysis of the reduction.

Theorem 3 follows from our completeness and soundness bounds for Φ_Σ, stated in the next results, using the fact that by Theorem 1, there are finite sets D, E such that $\mathrm{GLC}_{D,E}(1, \alpha)$ is NP-hard for the value of α chosen in Theorem 7 below. The proofs of the completeness and soundness bounds can be found in the full version [18].

▶ **Theorem 6** (Completeness). *Let Σ be a Gap Label Cover instance and Φ_Σ be the system defined in* (1)*. Suppose that Σ is 1-satisfiable. Then Φ_Σ is $(1 - \epsilon)$-satisfiable in $\mathcal{G}_1$.*

▶ **Theorem 7** (Soundness). *Let Σ be a Gap Label Cover instance and Φ_Σ be the system defined in* (1)*. Suppose that Φ_Σ is $(1/|\mathcal{H}_2| + \delta)$-satisfiable in $\mathcal{G}_2$. Then Σ is α-satisfiable, where $\alpha = \delta^2/(4\kappa|\mathcal{G}_1|^\kappa|\mathcal{G}_2|^4)$ and $\kappa = \lceil(\log_2 \delta - 2)/(\log_2(1 - \epsilon)\rceil$.*

1.4 Proof Outline

The main difficulty in proving the correctness of our reduction lies in showing the soundness bound (Theorem 7). The completeness result (Theorem 6) is relatively straightforward and follows as in [23]. In summary, suppose the Gap Label Cover instance Σ is satisfied by a pair of assignments $h_D : U \to D$, $h_E : V \to E$. Then we find families A, B such that $\Phi_\Sigma^{\mathcal{G}_1}(A, B) \geq 1 - \epsilon$ by letting A_v be the $h_E(v)$-th projection and B_u be the $h_D(u)$-th projection for each $v \in V, u \in U$. As usual, the noise introduced by the perturbation factor $\boldsymbol{\nu}$ is what forces us to give up perfect completeness.

The idea behind our soundness analysis has appeared many times in the literature (e.g., [25, 23, 10]), but the approach taken in [23] is the most similar to ours. Suppose that there are assignments A, B, satisfying

$$\Phi_{\Sigma}^{\mathcal{G}_2}(A,B) \geq \frac{1}{|\mathcal{H}_2|} + \delta. \tag{4}$$

In view of (3), this inequality can be understood as a lower bound for the success probability of the following 3-query dictatorship test: Sample all parameters according to the distribution shown in Figure 1, and then query the values $(A_v)_\varphi(\mathbf{a})$, $B_u(\mathbf{b}^{s_1})^{s_1}$, and $B_u((\mathbf{b}^{-1}(\mathbf{a}\circ\pi_{uv})^{-1}\boldsymbol{\nu})^{s_2})^{s_2}$. The test is passed if the product of the three values is the group identity, and failed otherwise. The soundness proof consists in showing that (4) implies that the functions $(A_v)_\varphi : \mathcal{G}_1^E \to \mathcal{G}_2$ and $B_u : \mathcal{G}_1^D \to \mathcal{G}_2$ are "close" to dictators (i.e., projections) for each $v \in V$, $u \in U$. Then, this fact allows us to find a good solution to the starting Gap Label Cover instance Σ. Indeed, suppose that for each $v \in V$ the map $(A_v)_\varphi$ is the projection on the e_v-th coordinate, and for each $u \in U$, the map B_u is the projection on the d_u-th coordinate. Then the assignment mapping v to e_v and u to d_u for each $v \in V, u \in U$ is a good solution for Σ. However, it is not clear how to extend this simple idea to the case where the maps $(A_v)_\varphi, B_u$ are not projections.

In order to find a good solution for Σ in this general case, we first find suitable maps $\gamma_1, \gamma_2 : \mathcal{G}_2 \to \mathbb{C}$ and analyse $\gamma_1 \circ (A_v)_\varphi$, $\gamma_2 \circ B_u$. Now, using the fact that $(A_v)_\varphi$ and B_u are close to projections, we can prove that choosing the labels e, d for the vertices v, u according to the "low-degree influence" of the e-th coordinate in $\gamma_1 \circ (A_v)_\varphi$ and the d-th coordinate in $\gamma_2 \circ B_u$ yields a good randomised assignment of Σ.

This overview so far also applies to the soundness analysis of [23]. Let us give more detail and highlight the main differences that sets our work apart. The first important difference has to do with the choice of γ_1, γ_2. We define $\gamma_1 = \omega_{x,y}$, and $\gamma_2 = \omega_{y,z}$, where ω is some irreducible representation of $\mathcal{G}_2$, and x, y, z are suitable indices in N_ω. In [23], the representation ω is a non-trivial representation chosen so that

$$\left|\mathbb{E}\left[\chi_\omega\left((A_v)_\varphi(\mathbf{a})B_u(\mathbf{b}^{s_1})^{s_1}B_u((\mathbf{b}^{-1}(\mathbf{a}\circ\pi)^{-1}\boldsymbol{\nu})^{s_2})^{s_2}\right)\right]\right| \geq \dim_\omega \delta.$$

Here the expectation is taken over the probability space described in Figure 1, and the dependence of π on the edge $\{u, v\}$ is left implicit. In our case, rather than using the Fourier characters for choosing ω, we consider "penalized characters" $\widetilde{\chi_\omega}$. We define $\widetilde{\chi_\omega} : \mathcal{G}_2 \to \mathbb{C}$ as the map $\chi_\omega - \eta_\omega$, where the penalty η_ω is the multiplicity of the trivial representation in the restriction $\omega|_{\mathcal{H}_2}$. This way, we pick $\omega \in \widehat{\mathcal{G}_2}$ so that the previous inequality holds after replacing χ_ω with $\widetilde{\chi_\omega}$. Equivalently, we find ω satisfying

$$\left|\mathbb{E}\left[\chi_\omega\left((A_v)_\varphi(\mathbf{a})B_u(\mathbf{b}^{s_1})^{s_1}B_u((\mathbf{b}^{-1}(\mathbf{a}\circ\pi)^{-1}\boldsymbol{\nu})^{s_2})^{s_2}\right)\right]\right| \geq \dim_\omega \delta + \eta_\omega. \tag{5}$$

The fact that such ω exists is a consequence of (4) together with

$$\sum_{\omega \in \widehat{\mathcal{G}_2}} \dim_\omega \eta_\omega = |\mathcal{G}_2|/|\mathcal{H}_2|,$$

which follows from the Frobenius Reciprocity Theorem, as shown in detail in [18]. This additional factor of η_ω is crucial to our soundness analysis, as we will see.

Define the map $\mathcal{A} = \omega \circ (A_v)_\varphi$ and the map $\mathcal{B} : \mathcal{G}_1^D \to \mathcal{G}_2$ given by $\mathcal{B}(\mathbf{b}) = \mathbb{E}_{s\in\{-1,1\}}\omega \circ B_u(\mathbf{b}^s)^s$, where $s \in \{-1, 1\}$ is distributed uniformly.[7] To show the soundness bound we consider the Fourier expansions of $\mathcal{A}$ and $\mathcal{B} * \mathcal{B}$ in the expression

$$\left|\operatorname{tr}\mathbb{E}\left[\mathcal{A}(\mathbf{a})(\mathcal{B}*\mathcal{B})((\mathbf{a}\circ\pi)^{-1}\boldsymbol{\nu})\right]\right|,$$

[7] Observe that the maps $\mathcal{A}$ and $\mathcal{B}$ depend on the hidden parameters v and u respectively.

which is just a rearrangement of the left-hand-side in the previous inequality. More precisely, we look at the equivalent expression

$$\left|\operatorname{tr}\mathbb{E}\left[\left(\sum_{\tau\in\widehat{\mathcal{G}^E},s,t\in N_\tau}\dim_\tau \widehat{\mathcal{A}}(\tau_{s,t})\tau_{s,t}(\mathbf{a})\right)\right.\right. \tag{6}$$

$$\left.\left.\times\left(\sum_{\rho\in\widehat{\mathcal{G}^D},i,j\in N_\rho}\dim_\rho\widehat{(\mathcal{B}*\mathcal{B})}(\rho_{i,j})\rho_{i,j}((\mathbf{a}\circ\pi)^{-1}\boldsymbol{\nu})\right)\right]\right|.$$

Our goal is to find a bound κ, independent of $|D|, |E|$, satisfying that the contribution to this expression of non-trivial representations τ, ρ of degree less than κ is at least $\dim_\omega \delta/2$. This is achieved by controlling the contribution of the trivial term and the contribution of high-degree terms. The second main difference of our soundness analysis compared to [23] is our handling of the trivial term. In the full version [18] we prove that

$$\left|\operatorname{tr}\mathbb{E}\left[\widehat{\mathcal{A}}(1)\left(\sum_{\rho\in\widehat{\mathcal{G}^D},i,j\in N_\rho}\dim_\rho\widehat{(\mathcal{B}*\mathcal{B})}(\rho_{i,j})\rho_{i,j}((\mathbf{a}\circ\pi)^{-1}\boldsymbol{\nu})\right)\right]\right|\leq\eta_\omega.$$

In the non-promise setting, this bound is not necessary. Roughly, under the stronger notion of folding used in [23], it is possible to show that $\widehat{\mathcal{A}}(1)$ vanishes. Our weaker notion of folding does not allow us to prove the same result, but we are still able to leverage folding to obtain the above bound. This mismatch with [23] is the reason why the extra η_ω term was required in (5). The key insight in the proof of the inequality above is that if $F : \mathcal{G}_1^E \to \mathcal{G}_2$ is folded over φ, then the trace of $\widehat{(\omega\circ F)}(1)$ is at most η_ω in absolute value.

Our analysis of high-degree terms is in the same spirit as previous works that show hardness of approximation in the imperfect completeness setting. In [18] we prove that

$$\left|\operatorname{tr}\mathbb{E}\left[\left(\sum_{\tau\in\widehat{\mathcal{G}_1^E},\tau\neq 1}\sum_{s,t\in N_\tau}\dim_\tau \widehat{\mathcal{A}}(\tau_{s,t})\tau_{s,t}(\mathbf{a})\right)\times\right.\right.$$

$$\left.\left.\left(\sum_{\rho\in\widehat{\mathcal{G}_1^D},|\rho|\geq\kappa}\sum_{i,j\in N_\rho}\dim_\rho\widehat{(\mathcal{B}*\mathcal{B})}(\rho_{i,j})\rho_{i,j}((\mathbf{a}\circ\pi)^{-1}\boldsymbol{\nu})\right)\right]\right|\leq(\dim_\omega\delta)/2$$

for all $\kappa \geq (\log_2\delta - 2)/\log_2(1-\epsilon)$. The essential idea is that the "noise vector" $\boldsymbol{\nu}$ has a smoothing effect that limits the contribution of high-degree terms in (6).

Finally, having established that the contribution of non-trivial terms of degree less than κ in (6) is at least $\dim_\omega \delta/2$, in [18] we give a good randomised strategy to solve Σ. This strategy assigns the label $e \in E$ to $v \in V$ and the label $d \in D$ to $u \in U$ with probabilities

$$\Pr(v\mapsto e) = \sum_{\tau\in\widehat{\mathcal{G}_1^E},\tau^e\neq 1}\sum_{s,t\in N_\tau}\dim_\tau\frac{\left|\widehat{\mathcal{A}_{x,y}}(\tau_{s,t})\right|^2}{|\tau|}$$

and

$$\Pr(u\mapsto d) = \sum_{\rho\in\widehat{\mathcal{G}_1^D},\rho^d\neq 1}\sum_{i,j\in N_\rho}\dim_\rho\frac{\left|\widehat{\mathcal{B}_{y,z}}(\rho_{i,j})\right|^2}{|\rho|},$$

where $x, y, z \in N_\omega$ are suitable indices [18]. These probabilities are supposed to capture the influence of the e-th and d-th coordinates on $\mathcal{A}_{x,y} = \omega_{x,y} \circ (A_v)_\varphi$ and $\mathcal{B}_{y,z} = \omega_{y,z} \circ \mathbb{E}_s B_u(\,\cdot^s)^s$ respectively.[8] (At this point it may be helpful to recall that B_u is a function from $\mathcal{G}_1^D$ to $\mathcal{G}_2$ and s is a sign sampled uniformly from $\{-1, 1\}$. Thus, $B_u(\,\cdot^s)^s$ takes an element $b \in \mathcal{G}_1^D$ and returns $(B_u(b^s))^s$.) This turns out to be a good randomised assignment for Σ. That is,

$$\mathbb{E}_{uv}\left[\sum_{d\in D} \Pr(v \mapsto \pi_{uv}(d)) \Pr(u \mapsto d)\right] \geq \alpha, \tag{7}$$

where the expectation is taken uniformly over the edges $\{u, v\}$ of Σ, and α is the soundness constant appearing in Theorem 7. We are being informal with the usage of the word "probability" here: the quantities $\Pr(v \mapsto e)$ and $\Pr(u \mapsto d)$ may add up to less than 1, but this is easily fixed by normalising, or by letting our strategy default to the uniform assignment with some positive probability.

Let us give some more detail. More precisely, in the full version [18] we show that truncating our assignment probabilities to terms of degree less than κ is enough to satisfy this last inequality. Let $\ell \geq 0$. The probabilities $\Pr^{<\ell}(v \mapsto e)$, $\Pr^{<\ell}(u \mapsto d)$ are defined the same way as $\Pr(v \mapsto e)$ and $\Pr(u \mapsto d)$ but considering only representations τ, ρ of degree less than ℓ. These modified probabilities can be understood as the "low-degree influences" of each coordinate in $\mathcal{A}_{x,y}$ and $\mathcal{B}_{y,z}$. With this notation, in [18] we prove that (7) holds after replacing each assignment probability Pr with its truncated variant $\Pr^{<\kappa}$. In other words, we prove that

$$\mathbb{E}_{uv}\left[\sum_{d\in D} \sum_{\substack{\rho \in \widehat{\mathcal{G}_1^D}, \rho^d \neq 1 \\ |\rho| < \kappa,\, i,j \in N_\rho}} \sum_{\substack{\tau \in \widehat{\mathcal{G}_1^E}, \tau^{\pi_{uv}(d)} \neq 1 \\ |\tau| < \kappa,\, s,t \in N_\tau}} \frac{\dim_\tau \left|\widehat{\mathcal{A}_{x,y}}(\tau_{s,t})\right|^2}{|\tau|} \frac{\dim_\rho \left|\widehat{\mathcal{B}_{y,z}}(\rho_{i,j})\right|^2}{|\rho|}\right] \geq \alpha.$$

This shows that our proposed strategy produces a good randomised assignment for Σ and completes the soundness proof.

References

1 Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach.* Cambridge University Press, 2009.

2 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. doi:10.1145/278298.278306.

3 Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998. doi:10.1145/273865.273901.

4 Per Austrin, Jonah Brown-Cohen, and Johan Håstad. Optimal inapproximability with universal factor graphs. *ACM Trans. Algorithms*, 2023. doi:10.1145/3631119.

5 Per Austrin, Venkatesan Guruswami, and Johan Håstad. (2+ϵ)-Sat is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017. doi:10.1137/15M1006507.

6 Per Austrin and Johan Håstad. On the usefulness of predicates. *ACM Trans. Comput. Theory*, 5(1):1:1–1:24, 2013. doi:10.1145/2462896.2462897.

[8] This is similar to the notion of influence in [10, 7].

7 Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009. doi:10.1007/s00037-009-0272-6.

8 Libor Barto, Jakub Bulín, Andrei A. Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–28:66, 2021. doi:10.1145/3457606.

9 Libor Barto, Silvia Butti, Alexandr Kazda, Caterina Viola, and Stanislav Živný. Algebraic approach to approximation. In *Proc. 39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'24)*, pages 10:1–10:14. ACM, 2024. doi:10.1145/3661814.3662076.

10 Amey Bhangale and Subhash Khot. Optimal Inapproximability of Satisfiable k-LIN over Non-Abelian Groups. In *Proc. 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC'21)*, pages 1615–1628. ACM, 2021. doi:10.1145/3406325.3451003.

11 Amey Bhangale, Subhash Khot, and Dor Minzer. On Approximability of Satisfiable k-CSPs: I. In *Proc. 54th Annual ACM Symposium on Theory of Computing (STOC'22)*, pages 976–988. ACM, 2022. doi:10.1145/3519935.3520028.

12 Amey Bhangale, Subhash Khot, and Dor Minzer. On Approximability of Satisfiable k-CSPs: II. In *Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23)*, pages 632–642. ACM, 2023. doi:10.1145/3564246.3585120.

13 Amey Bhangale, Subhash Khot, and Dor Minzer. On Approximability of Satisfiable k-CSPs: III. In *Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23)*, pages 643–655. ACM, 2023. doi:10.1145/3564246.3585121.

14 Amey Bhangale and Aleksa Stankovic. Max-3-Lin Over Non-abelian Groups with Universal Factor Graphs. *Algorithmica*, 85(9):2693–2734, 2023. doi:10.1007/S00453-023-01115-1.

15 Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric Boolean dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021. doi:10.1137/19M128212X.

16 Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. SDPs and robust satisfiability of promise CSP. In *Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23)*, pages 609–622, 2023. arXiv:2211.08373, doi:10.1145/3564246.3585180.

17 Jonah Brown-Cohen and Prasad Raghavendra. Combinatorial optimization algorithms via polymorphisms. Technical report, 2015. arXiv:1501.01598.

18 Silvia Butti, Alberto Larrauri, and Stanislav Živný. Optimal inapproximability of promise equations over finite groups. Technical report, 2024. arXiv:2411.01630.

19 Siu On Chan. Approximation resistance from pairwise-independent subgroups. *J. ACM*, 63(3):27:1–27:32, 2016. doi:10.1145/2873054.

20 Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007. doi:10.1145/1236457.1236459.

21 Irit Dinur, Subhash Khot, Will Perkins, and Muli Safra. Hardness of finding independent sets in almost 3-colorable graphs. In *Proc. 51th Annual IEEE Symposium on Foundations of Computer Science (FOCS'10)*, pages 212–221. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.84.

22 Lars Engebretsen and Jonas Holmerin. More efficient queries in PCPs for NP and improved approximation hardness of maximum CSP. *Random Struct. Algorithms*, 33(4):497–514, 2008. doi:10.1002/RSA.20226.

23 Lars Engebretsen, Jonas Holmerin, and Alexander Russell. Inapproximability results for equations over finite groups. *Theor. Comput. Sci.*, 312(1):17–45, 2004. doi:10.1016/S0304-3975(03)00401-8.

24 Venkatesan Guruswami and Prasad Raghavendra. Hardness of Solving Sparse Overdetermined Linear Systems: A 3-Query PCP over Integers. *ACM Trans. Comput. Theory*, 1(2):6:1–6:20, 2009. doi:10.1145/1595391.1595393.

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

26 Yahli Hecht, Dor Minzer, and Muli Safra. NP-Hardness of Almost Coloring Almost 3-Colorable Graphs. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms*

and Techniques (APPROX/RANDOM'23), volume 275 of *LIPIcs*, pages 51:1–51:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.APPROX/RANDOM.2023.51.

27 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC'02)*, pages 767–775, 2002. doi:10.1145/509907.510017.

28 Subhash Khot and Dana Moshkovitz. *NP*-Hardness of Approximately Solving Linear Equations over Reals. *SIAM J. Comput.*, 42(3):752–791, 2013. doi:10.1137/110846415.

29 Subhash Khot and Rishi Saket. Hardness of finding independent sets in almost q-colorable graphs. In *Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'12)*, pages 380–389. IEEE Computer Society, 2012. doi:10.1109/FOCS.2012.75.

30 Subhash Khot, Madhur Tulsiani, and Pratik Worah. A characterization of strong approximation resistance. In *Proc. 46th Symposium on Theory of Computing (STOC'14)*, pages 634–643. ACM, 2014. doi:10.1145/2591796.2591817.

31 Subhash Khot, Madhur Tulsiani, and Pratik Worah. The complexity of somewhat approximation resistant predicates. In *Proc. 41st International Colloquium on Automata, Languages, and Programming (ICALP'14)*, volume 8572 of *Lecture Notes in Computer Science*, pages 689–700. Springer, 2014. doi:10.1007/978-3-662-43948-7_57.

32 Alberto Larrauri and Stanislav Živný. Solving promise equations over monoids and groups. In *Proc. 51st International Colloquium on Automata, Languages, and Programming (ICALP'24)*, volume 297 of *LIPIcs*, pages 146:1–146:18, 2024. doi:10.4230/LIPIcs.ICALP.2024.146.

33 Tamio-Vesa Nakajima and Stanislav Živný. Maximum k- vs. ℓ-colourings of graphs. Technical report, 2023. arXiv:2311.00440.

34 Tamio-Vesa Nakajima and Stanislav Živný. Maximum bipartite vs. triangle-free subgraph. In *Proc. 52nd International Colloquium on Automata, Languages, and Programming (ICALP'25)*, 2025. arXiv:2406.20069.

35 Ryan O'Donnell, Yi Wu, and Yuan Zhou. Hardness of Max-2Lin and Max-3Lin over Integers, Reals, and Large Cyclic Groups. *ACM Trans. Comput. Theory*, 7(2):9:1–9:16, 2015. doi:10.1145/2751322.

36 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 245–254, 2008. doi:10.1145/1374376.1374414.

37 Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. doi:10.1137/S0097539795280895.

38 Thomas Schaefer. The complexity of satisfiability problems. In *Proc. 10th Annual ACM Symposium on the Theory of Computing (STOC'78)*, pages 216–226, 1978. doi:10.1145/800133.804350.

39 Audrey Terras. *Fourier Transform and Representations of Finite Groups*, page 240–266. London Mathematical Society Student Texts. Cambridge University Press, 1999.