Optimal Inapproximability of Satisfiable k -LIN over Non-Abelian Groups

Amey Bhangale

University of California, Riverside
Riverside, USA
ameyb@ucr.edu

Subhash Khot*

New York University
New York, USA
khot@cs.nyu.edu

ABSTRACT

A seminal result of Håstad (2001) shows that it is NP-hard to find an assignment that satisfies $\frac{1}{|G|} + \epsilon$ fraction of the constraints of a given k -LIN instance over an abelian group, even if there is an assignment that satisfies $(1 - \epsilon)$ fraction of the constraints, for any constant $\epsilon > 0$. Engebretsen, Holmerin and Russell (2004) later showed that the same hardness result holds for k -LIN instances over any finite non-abelian group.

Unlike the abelian case, where we can efficiently find a solution if the instance is satisfiable, in the non-abelian case, it is NP-complete to decide if a given system of linear equations is satisfiable or not, as shown by Goldmann and Russell (1999).

Surprisingly, for certain non-abelian groups G , given a satisfiable k -LIN instance over G , one can in fact do better than just outputting a random assignment using a simple but clever algorithm. The approximation factor achieved by this algorithm varies with the underlying group. In this paper, we show that this algorithm is *optimal* by proving a tight hardness of approximation of satisfiable k -LIN instance over *any* non-abelian G , assuming $\mathbf{P} \neq \mathbf{NP}$. As a corollary, we also get 3-query probabilistically checkable proofs with perfect completeness over large alphabets with improved soundness.

Our proof crucially uses the *quasi-random* properties of the non-abelian groups defined by Gowers (2008).

CCS CONCEPTS

• Theory of computation → Problems, reductions and completeness.

KEYWORDS

linear equations, non-abelian groups, hardness of approximation, constraint satisfaction problems

ACM Reference Format:

Amey Bhangale and Subhash Khot. 2021. Optimal Inapproximability of Satisfiable k -LIN over Non-Abelian Groups. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '21)*, June 21–25, 2021, Virtual, Italy.

*Research funded by NSF grant CCF-1813438, Simons Investigator Award, and Simons Collaboration on Algorithms and Geometry.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '21, June 21–25, 2021, Virtual, Italy

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8053-9/21/06...\$15.00

<https://doi.org/10.1145/3406325.3451003>

2021, Virtual, Italy. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3406325.3451003>

1 INTRODUCTION

Constraint satisfaction problems (CSPs) are the most fundamental problems in computer science. A simplest such CSP which we know how to solve is a system of k -LIN equations over some abelian group. More generally, an instance of Max- k -LIN over a group $G = (G, \cdot)$, not necessarily abelian, consists of a set of variables $x_1, x_2, \dots, x_n$ and a set of constraints $C_1, C_2, \dots, C_m$. Each C_i is a linear equation involving k variables, for example $a_1 \cdot x_{i_1} \cdot a_2 \cdot x_{i_2} \cdot \dots \cdot a_k \cdot x_{i_k} = b$, for some group elements $a_1, a_2, \dots, a_k, b \in G$. The task is to find an assignment to the variables that satisfies as many constraints as possible.

For *any* abelian group G , if there is a perfect solution to the given Max- k -LIN instance over G , then it can be found efficiently in polynomial time using Gaussian elimination. A given instance is *almost* satisfiable if there exists an assignment that satisfies $(1 - \epsilon)$ -fraction of the constraints for small constant $\epsilon > 0$. If the given instance of Max- k -LIN over an abelian group G is almost satisfiable, then Håstad [13] showed that it is NP-hard to even find an assignment that satisfies $\frac{1}{|G|} + \epsilon$ of the constraints for every constant $\epsilon > 0$. In other words, one cannot do significantly better than just outputting a random assignment.

The situation changes completely if the instance is a set of linear equations over a non-abelian group. In this case, Goldmann and Russell [11] showed that the problem of deciding if a given instance is satisfiable or not is NP-complete, for *every* non-abelian group.

An algorithm (folklore): It turns out, one can do much better than outputting a random assignment for some groups G , when the instance is satisfiable. Given an instance ϕ over G , consider an instance ϕ' over $H = G/[G, G]$ where $[G, G]$ is a commutator subgroup of G , i.e., the subgroup generated by the elements $\{g^{-1}h^{-1}gh \mid g, h \in G\}$. The instance ϕ' is same as ϕ except that all the group constants are replaced by their equivalence class in $G/[G, G]$. The important property of this quotient group H is that it is an abelian group. Since ϕ has a satisfying assignment over G , ϕ' has a satisfying assignment over H . Hence, we can find the satisfying assignment σ of ϕ' in polynomial time. The solution σ is an assignment of cosets of $[G, G]$ to the variables. We construct a random assignment to ϕ such that for every variable x , we select a random group element from $\sigma(x)$ and assign it to x . It is easy to see that each constraint is satisfied with probability equal to the $\frac{1}{|[G, G]|}$. Thus, this gives an assignment that satisfies $\frac{1}{|[G, G]|}$ fraction of the constraints in expectation. Therefore, if there exists a non-trivial

commutator subgroup of G , then we get an algorithm which does better than the random assignment threshold.

If the instance is almost satisfiable, then it is not clear how to modify the above algorithm for almost satisfiable instances to get better than $\frac{1}{|G|}$ approximation. In fact, for almost satisfiable instances over any non-abelian group, Engebretsen, Holmerin and Russell [9] showed that it is **NP**-hard to do better than outputting a random assignment.

This leaves an intriguing question of finding the correct approximation threshold for satisfiable instances over non-abelian groups. In this paper, we show that the above described algorithm for satisfiable instances over non-abelian groups is the best one can hope for. More specifically, we prove the following theorem.

THEOREM 1.1. *For any constant $\varepsilon > 0$, given a satisfiable instance of a Max-3-LIN over a finite non-abelian group G , it is **NP**-hard to find an assignment that satisfies $\frac{1}{|G|} + \varepsilon$ fraction of the constraints.*

The theorem can be extended to Max- k -LIN for any $k \geq 3$ to imply a similar hardness result for all Max- k -LIN problems over G .

If G is a *perfect group*, i.e., $[[G, G]] = |G|$, then Theorem 1.1 implies an NP-hardness of approximating satisfiable Max-3-CSP instances over an alphabet of size q to within a factor of $\frac{1}{q} + \varepsilon$, for every constant $\varepsilon > 0$. As a direct consequence, we get improved soundness of 3-query probabilistically checkable proofs (PCPs) with perfect completeness over large alphabets. Since PCPs are not the main focus of this paper, we refer interested readers to the book by Arora and Barak [1, Chapter 18] to see the relation between PCPs and CSPs.

COROLLARY 1.2. *For infinitely many $q \in \mathbb{Z}^+$, any language in **NP** is decided by a non-adaptive PCP with answers from a domain of size q that queries three positions in the proof, has perfect completeness and soundness $\frac{1}{q} + \varepsilon$ for any constant $\varepsilon > 0$.*

This improves a result by Engebretsen and Holmerin [8] where they constructed PCPs with soundness $\frac{1}{q} + \frac{1}{q^2} + \varepsilon$, and also a result by Tang [17] in which they showed a conditional result with soundness $\frac{1}{q} + \frac{1}{q^2} - \frac{1}{q^3} + \varepsilon$, for any constant $\varepsilon > 0$.¹

1.1 Techniques

We assume some familiarity with the Fourier analysis of functions over abelian groups (for instance, Chapter 8 of Ryan O'Donnell's book [15]). Throughout the section, $\varepsilon > 0$ is an arbitrarily small constant and $\delta(\varepsilon) > 0$ decays with ε . We only discuss 3-LIN here, however the argument is similar for k -LIN in general.

Given a well established field of hardness of approximation where the starting point is the Label Cover problem (see Definition 2.1), at the heart of these reductions are the *dictatorship tests*. A function $f : G^n \rightarrow G$ is a dictator function if it depends only on one variable, i.e., $f(x_1, x_2, \dots, x_n) = x_i$ for some $i \in [n]$. On the other hand, we have functions which are *far* from dictator functions.

To understand a notion of distance from a dictator function, which is useful for a reduction to work, define the influence of the i^{th} coordinate on the function to be the probability that on a random input, changing the i^{th} coordinate changes the values of

the function. In terms of the Fourier coefficients of f , this is equal to the following quantity:

$$\text{Inf}_i(f) := \sum_{\alpha: \alpha_i \neq 0} |\hat{f}(\alpha)|^2.$$

Thus, the i^{th} dictator function has $\text{Inf}_i(f) = 1$. At the first attempt, it might make sense to define functions which are close to dictators are the functions with a coordinate with large influence. However, note that there are linear functions $\ell_S = \sum_{i \in S} \beta_i x_i$ where $S \subseteq [n]$ such that ℓ_S has all the variables $i \in S$ with influences 1. We would like to isolate these functions with large $|S|$ from the dictator functions. This motivates to define a more refined notion of a low-degree influence of a variable i as follows:

$$\text{Inf}_i^{\leq d}(f) := \sum_{\alpha: \alpha_i \neq 0 \wedge |\alpha| \leq d} |\hat{f}(\alpha)|^2,$$

where $|\alpha|$ is the number of non-zero coordinates of α . Thus, for the i^{th} dictator function, its low-degree ($d = 1$) influence of the coordinate i is 1 (and rest of the influences are 0). A function is far from any dictator function if all the low-degree influences, for some $d = O(1)$ which is independent of n , of the function are small, say at most ε .

Although the above definition is the correct definition for most reductions, in case of linear equations, we work with an even weaker notion of the distance. We consider the following definition. A function is far from dictator functions if for every α such that $|\alpha| = O(1)$ which is independent of n , $|\hat{f}(\alpha)|^2 \leq \varepsilon$. In other words, all the low-degree Fourier coefficients of f have small weights. Note that this notion still isolates ℓ_S with large $|S|$.

A (non-adaptive) dictatorship test queries the function f at a few locations and based on the values it sees, decides if the function is a dictator function or far from it. This choice of predicate is tightly connected to the specific constraint satisfaction problem (CSP) for which we want to show a NP-hardness result. Furthermore, the gap between the test passing probability in the *completeness case* (when f is a dictator function) and the *soundness case* (when f is far from any dictator function) translates into the inapproximability factor of NP-hardness of the CSP.²

1.1.1 Abelian Groups. Let us look at a candidate dictatorship test where the predicate is a linear equation in 3 variables over an abelian group $G \cong \mathbb{Z}_q$. Here, 0 is the identity element of G .

- Select $x, y \sim G^n$ uniformly at random.
- Set $z = x + y$.
- Check if $f(x) + f(y) = f(z)$.

It is clear that if f is an i^{th} dictator then the test passes with probability 1. The non-trivial thing is to analyze the test passing probability if f is far from any dictator. It is easy to see that there are functions which are far from dictators and still the test passes with probability 1 on them. A family of such functions are the linear functions of the form $\ell_S = \sum_{i \in S} \beta_i x_i$ where $S \subseteq [n]$ and $\beta_i \in G \setminus \{0\}$, with large $|S|$. It is not hard to see that these functions pass the test with probability 1. In fact, Blum, Luby and Rubinfeld [7]

¹The theorem in [8] holds for every $q \geq 3$, and the theorem in [17] holds for every $q \geq 4$. Our theorem holds for q such that there are perfect groups of cardinality q .

²This is not totally correct as one has to overcome other important issues when such a test is used in the actual reduction, starting with the Label Cover instance.

showed that this is a good test for the linear functions (instead of the dictator functions).

One must be able to design a test such that ℓ_S for large S passes with small probability. To design such a test, Håstad [13] introduced the so called *noise* to each coordinate. The modified test is as follows.

- Select $x, y \sim G^n$ uniformly at random.
- Set $z = x + y$.
- For each $i \in [n]$, resample (x_i, y_i, z_i) from G^3 uniformly at random, with probability ε .
- Check if $f(x) + f(y) = f(z)$.

This noise takes care of the earlier mentioned counterexamples, i.e., functions ℓ_S for large S now pass the test with probability roughly $\frac{1}{|G|}$. In general, Håstad [13] showed that if f is far from dictator functions then the test passes with probability at most $\frac{1}{|G|} + \delta$ for a small constant $\delta > 0$. The proof of this statement uses Fourier analysis over abelian groups. Note that this bound is optimal as even a random function passes the test with probability $\frac{1}{|G|}$.

However, now the guarantee in the completeness case is no longer the same. We only get that dictator functions pass this test with probability $1 - \varepsilon$ (instead of 1). This gap in the test passing probability is translated into the NP-hardness of 3-LIN over abelian group and coincidentally, in this abelian case, the NP-hardness result is *optimal*. More precisely, given a system of linear equations over an abelian group G , where each equation involves 3 variables, it is NP-hard to distinguish between the cases when there exists an assignment that satisfies at least $(1 - \varepsilon)$ -fraction of the constraints vs. no assignment can satisfy more than $\frac{1}{|G|} + \delta$ fraction of the constraints.

1.1.2 Non-Abelian Groups. We now look into the non-abelian case. Since we would like to design a test which passes with probability 1 (or $(1 - \varepsilon)$) in the completeness case, there is a natural generalization of the above mentioned tests to a non-abelian group G . Here we denote the group operation by the symbol $\cdot$ and the identity element of G by 1_G .

We first describe the test with completeness $(1 - \varepsilon)$, which is similar to the test over abelian group with noise we described earlier.

- Select $x, y \sim G^n$ uniformly at random.
- For each $i \in [n]$, set $z_i = y_i^{-1} \cdot x_i^{-1}$.
- For each $i \in [n]$, resample (x_i, y_i, z_i) from G^3 uniformly at random, with probability ε .
- Check if $f(x) \cdot f(y) \cdot f(z) = 1_G$.

The analysis of this test is implicit in the work of Engebretsen, Holmerin and Russell [9]. Firstly, it is easy to see that the dictator functions pass this test with probability $(1 - \varepsilon)$. The soundness of this test is analyzed in [9] where the authors show that in the soundness case, the test passes with probability at most $\frac{1}{|G|} + \delta$ for small constant $\delta > 0$. Their proof goes via Fourier analysis over non-abelian groups. As in the abelian case, in this case also, it can be shown that the noise takes care of *high degree* Fourier terms.³

³Although we have not formally defined a degree of a Fourier coefficient of a function in this non-abelian setting, think of it as the number of non-trivial irreducible representations in α (See Proposition 2.24).

This implies a NP-hardness result of approximating 3-LIN instances over non-abelian group, which is similar to the abelian case.

Although the proof of the soundness of the test in [9] uses representation theory and Fourier analysis of functions on non-abelian groups, the proof now follows from the more general statement called the *invariance principle* of Mossel [14]. The distribution on the tuple (x, y, z) is a product distribution $\mu^{\otimes n}$ where μ is a distribution on (x_i, y_i, z_i) (note that for all i it is the same distribution). Since we add noise to each coordinate with some non-zero probability, the distribution μ is *connected* and hence we can easily take care of high degree functions in the analysis. Furthermore, the distribution μ is pairwise independent. These two conditions imply that in the soundness, the test passes with probability at most $\frac{1}{|G|} + \delta$. This statement is implicit in the result of Austrin and Mossel [4].

We now state the (obvious) dictatorship test with perfect completeness.

- Select $x, y \sim G^n$ uniformly at random.
- For each $i \in [n]$, set $z_i = y_i^{-1} \cdot x_i^{-1}$.
- Check if $f(x) \cdot f(y) \cdot f(z) = 1_G$.

Our main contribution is the soundness analysis of the above dictatorship test over non-abelian groups *without noise*. Note that we cannot use the invariance principle based techniques in this case, as the distribution does not satisfy the condition of *connectedness*.

Proof overview. Our proof of the soundness analysis is inspired by the magic that was discovered by Gowers [12] to show that there are non-abelian groups where the size of any product free set is sublinear in $|G|$.⁴ Gowers' trick worked only for quasi-random groups⁵ as he was interested in $o(|G|)$ bound on the product free sets, whereas we are able to carry out our reduction for every non-abelian group.

The trick is elegantly captured by the following inequality by Babai, Nikolov, and Pyber [5]. For any functions $f, g : G \rightarrow \mathbb{C}$ with at least one of f, g having mean zero:

$$\|f * g\|_{L^2(G)} \leq \frac{1}{\sqrt{D}} \|f\|_{L^2(G)} \|g\|_{L^2(G)}, \quad (1)$$

where D is the smallest dimension of a non-trivial representation of G .⁶ In comparison, a trivial application of Cauchy–Schwartz inequality gives an upper bound of $\|f\|_{L^2(G)} \|g\|_{L^2(G)}$. Thus, Equation (1) has a multiplicative improvement of a factor $\frac{1}{\sqrt{D}}$ over a trivial upper bound.

Coming back to analyzing the soundness of our test, its analysis boils down to analyzing the following expression:

$$\begin{aligned} & \mathbb{E}[g_1(x)g_2(y)g_3(z)] \\ &= \mathbb{E}[(g_1 * g_2 * g_3)(1_{G^n})] \\ &\leq \sum_{\alpha \in \text{Irrep}(G^n)} \dim(\alpha) \cdot \|\hat{g}_1(\alpha)\|_{\text{HS}} \cdot \|\hat{g}_2(\alpha)\|_{\text{HS}} \cdot \|\hat{g}_3(\alpha)\|_{\text{HS}}, \quad (2) \end{aligned}$$

⁴unlike the abelian case where one can always find, in this case, a 'sum-free' set of size $\Omega(|G|)$.

⁵A group G is called *quasi-random* if the smallest dimension of any non-trivial representation of G is large.

⁶see Definition 2.10 and Definition 2.15 for the definitions of $\|\cdot\|_{L^2(G)}$ and the convolution operator $*$, and Section 2.2.1 for representation theory.

where g_i are bounded functions, derived from f , from $G^n \rightarrow \mathbb{C}$, i.e., $\|g_i\|_2 \leq 1$. $\hat{g}_i(\alpha)$ is the “fourier coefficient” of g_i corresponding to the irreducible representation α of G^n and $\|\cdot\|_{\text{HS}}$ is the Hilbert–Schmidt norm of a matrix. Here, the inequality follows by using the Fourier expansion of $(g_1 * g_2 * g_3)$ and a triangle inequality.

Once we have this expression, similar to Equation (1), it is easy to bound the terms with large $\dim(\alpha)$: By applying Cauchy–Schwarz inequality and using Parseval’s identity, we can show the following:

$$\begin{aligned} & \sum_{\dim(\alpha) \geq D} \dim(\alpha) \cdot \|\hat{g}_1(\alpha)\|_{\text{HS}} \cdot \|\hat{g}_2(\alpha)\|_{\text{HS}} \cdot \|\hat{g}_3(\alpha)\|_{\text{HS}} \\ & \leq \frac{1}{\sqrt{D}} \|g_1\|_2 \cdot \|g_2\|_2 \cdot \|g_3\|_2 \leq \frac{1}{\sqrt{D}}. \end{aligned}$$

Thus, we can effectively bound the *higher dimension* terms in the expression. Therefore, if the original expectation is δ , then we essentially get

$$\sum_{\dim(\alpha) \leq D} \dim(\alpha) \cdot \|\hat{g}_1(\alpha)\|_{\text{HS}} \cdot \|\hat{g}_2(\alpha)\|_{\text{HS}} \cdot \|\hat{g}_3(\alpha)\|_{\text{HS}} \approx \delta,$$

for a small D . Taking the maximum $\|\hat{g}_1(\alpha)\|_{\text{HS}}$ out in the summation, the remaining sum can be upper bounded by 1. Therefore, we get that there exists an α such that the dimension of α is at most D and $\|\hat{g}_1(\alpha)\|_{\text{HS}} \approx \delta - \frac{1}{\sqrt{D}}$.

Our analysis shows that if the test passes with probability greater than $\frac{1}{\|G, G\|} + \delta$, then there exists an α such that $\dim(\alpha) \leq O_\delta(1)$ and $\|\hat{f}(\alpha)\|_{\text{HS}} \approx \delta$, for some function $\tilde{f}$ derived from f . Note that this conclusion is different than what we had aimed for, i.e., concluding that there exists a *low-degree* Fourier coefficient with large magnitude. However, we show that such a bound is enough to carry out the actual soundness analysis of the reduction.

Since we do not introduce noise to each coordinate, there are functions with the Fourier mass concentrated on large dimensional α which pass this test with probability $\frac{1}{\|G, G\|}$. Thus, our analysis of the test is also optimal.

Although the dictatorship test works, there are many complications that arise when we compose this test with the Label Cover instance. We briefly discuss three issues here:

- (1) As observed before, in the soundness analysis we conclude that there is a low dimension Fourier coefficient whose norm is large, if the test passes with non-trivial probability. However, there are terms with dimension 1 but with high degree, which are problematic for the final *decoding strategy* in our reduction. In [9], these problematic terms were handled by adding noise, a technique which is similar to the abelian case. In our case, we do not have the noise. However, we observe, in Lemma 2.25, a stronger property of the folded functions (A function f is called folded if $f(c \cdot x) = c \cdot f(x)$ for all $x \in G^n$ and $c \in G$). Namely, if f is folded, then the function $\rho(f(x))_{ij}$, where ρ is any irreducible representation of G of dimension at least 2, has all the Fourier coefficients with dimension 1 *zero*. Thus, we can just focus on terms with dimensions at least 2.
- (2) Our decoding strategy is different from the one in [9]. The decoding strategy in [9] is based on non-empty low-degree Fourier coefficients, which was similar to Håstad’s decoding

strategy. In our reduction, we slightly changed the decoding strategy — it is based on the Fourier coefficients whose dimension is at least 2, but can be of high degree. This condition is forced on us by the way we are handling the higher dimensions terms. Fortunately, the decoding strategy works without much trouble.

- (3) Because of the d -to-1 nature of the projection constraints, we have to take care of many potential scenarios in the actual reduction when the error term can be large. We handle this collectively, in Claim 9, by using a careful choice of permuting the columns of the matrices (i.e., the Fourier coefficients and the representation matrices) involved in the soundness analysis.

2 PRELIMINARIES

2.1 Label Cover

We start by defining the LABEL-COVER problem which we use as a starting point for our reduction.

Definition 2.1 (LABEL-COVER). An instance $\mathcal{H} = (\mathcal{U}, \mathcal{V}, E, [L], [R], \{\pi_e\}_{e \in E})$ of the LABEL-COVER constraint satisfaction problem consists of a bi-regular bipartite graph $(\mathcal{U}, \mathcal{V}, E)$, two sets of alphabets $[L]$ and $[R]$ and a surjective projection map $\pi_e : [R] \rightarrow [L]$ for every edge $e \in E$. Given a labeling $\ell : \mathcal{U} \rightarrow [L], \ell' : \mathcal{V} \rightarrow [R]$, an edge $e = (u, v)$ is said to be satisfied by ℓ if $\pi_e(\ell(v)) = \ell(u)$.

$\mathcal{H}$ is said to be *satisfiable* if there exists a labeling that satisfies all the edges. $\mathcal{H}$ is said to be *at most δ -satisfiable* if every labeling satisfies at most a δ fraction of the edges.

The hardness of LABEL-COVER stated below follows from the PCP Theorem [2, 3, 10] and Raz’s Parallel Repetition Theorem [16]. The additional structural property on the hard instances (item 2 below) is proved by Håstad [13, Lemma 6.9].

THEOREM 2.2 (HARDNESS OF LABEL-COVER). For every $r \in \mathbb{N}$, there is a deterministic $n^{O(r)}$ -time reduction from a 3-SAT instance of size n to an instance $\mathcal{H} = (\mathcal{U}, \mathcal{V}, E, [L], [R], \{\pi_e\}_{e \in E})$ of LABEL-COVER with the following properties:

- (1) $|\mathcal{U}|, |\mathcal{V}| \leq n^{O(r)}$; $L, R \leq 2^{O(r)}$; $\mathcal{H}$ is bi-regular with degrees bounded by $2^{O(r)}$.
- (2) (Smoothness) There exists a constant $d_0 \in (0, 1/3)$ such that for any $v \in \mathcal{V}$ and $\alpha \subseteq [R]$, for a random neighbor u ,

$$\mathbb{E}_u [|\pi_{uv}(\alpha)|^{-1}] \leq |\alpha|^{-2d_0},$$

where $\pi_{uv}(\alpha) := \{i \in [L] \mid \exists j \in \alpha \text{ s.t. } \pi_{uv}(j) = i\}$. This implies that

$$\forall v, \alpha, \quad \Pr_u \left[|\pi_{uv}(\alpha)| < |\alpha|^{d_0} \right] \leq \frac{1}{|\alpha|^{d_0}}.$$

- (3) There is a constant $s_0 \in (0, 1)$ such that,
 - YES Case : If the 3-SAT instance is satisfiable, then $\mathcal{H}$ is satisfiable.
 - NO Case : If the 3-SAT instance is unsatisfiable, then $\mathcal{H}$ is at most $2^{-s_0 r}$ -satisfiable.

2.2 Fourier Analysis

In this section, we give a brief overview of the representation theory of non-abelian group and Fourier analysis over non-abelian groups.

For more comprehensive understanding, we refer the reader to the book by Terras [18]. We state many propositions in the following subsection, and the proofs of these propositions can be found in the same book [18].

2.2.1 Representation Theory. In this paper, we only consider non-abelian groups which are *finite*. Let $G = (G, \cdot)$ be a finite non-abelian group. The identity element of a group is denoted by 1_G .

Definition 2.3. A representation (V, ρ) of G is a vector space V together with a group homomorphism $\rho : G \rightarrow \text{GL}(V)$ from G to the group $\text{GL}(V)$ of invertible $\mathbb{C}$ -linear transformations from V to V . The dimension of the vector space V is denoted by $\dim(\rho)$.

For convenience, we just use the letter ρ to denote a representation of G and use ρ_V to denote the underlying vector space. We view a representation $\rho(\cdot)$ as its corresponding matrix of the linear transformation. Thus $\rho(\cdot)_{ij}$ is used to denote the $(i, j)^{\text{th}}$ entry of that matrix. We always work with representations which are unitary. There is one representation which is obvious – just map everything to $1 \in \mathbb{C}$. This representation is called the *trivial representation* which has dimension 1. We will denote the trivial representation by $\{1\}$.

Definition 2.4. Let ρ and τ be representations of G . An isomorphism from ρ_V to τ_V is an invertible linear transformation $\phi : \rho_V \rightarrow \tau_V$ such that

$$\phi \circ \rho(g) = \tau(g) \circ \phi,$$

for all $g \in G$. We say that ρ_V and τ_V are isomorphic and write $\rho_V \cong \tau_V$ if there exists an isomorphism from ρ_V to τ_V .

Definition 2.5. Let ρ be a representation of G . A vector subspace $W \subset \rho_V$ is G -invariant if $\rho(g)w \in W$ for all $g \in G$ and $w \in W$.

If a representation (V, ρ) has a G -invariant subspace W other than $\{0\}$ and V itself, then the action on W itself is a representation of G . This leads to the following important definition of irreducible representations.

Definition 2.6. A representation ρ of G is irreducible if $\rho_V \neq \{0\}$ and ρ_V has no G -invariant subspaces other than $\{0\}$ and ρ_V .

We will denote the set of all irreducible representations of G up to isomorphism by $\text{Irrep}(G)$.

FACT 1. Let G be a group and H be any subgroup of G , if $\rho \in \text{Irrep}(G)$ then ρ restricted to H is also a (not necessarily irreducible) representation of H .

Definition 2.7. The tensor product of two representations ρ and τ of a group G is the representation $\rho \otimes \tau$ on $\rho_V \otimes \tau_V$ defined by the condition

$$(\rho \otimes \tau)(g)(v \otimes w) = \rho(g)(v) \otimes \tau(g)(w),$$

and extended to all vectors in $\rho_V \otimes \tau_V$ by linearity.

Definition 2.8. The direct sum of two representations ρ and τ is the space $\rho_V \oplus \tau_V$ with the block-diagonal action $\rho \oplus \tau$ of G .

If the representation is not irreducible, then by an appropriate change of basis ρ can be converted into a block diagonal matrix with blocks corresponding to the invariant subspaces. Thus, any

representation can be completely decomposed into a direct sum of irreducible representations of G , by applying an appropriate unitary transformation. Note that this decomposition is *unique*. We use the following notation to denote the decomposition of a reducible representation: If ρ is a reducible representation of G then $\rho \cong \oplus_i n_i \rho_i$, where each i we have *distinct* $\rho_i \in \text{Irrep}(G)$ and n_i denotes the multiplicity of ρ_i in the decomposition. It will be convenient to think of this representation as a block diagonal matrices with ρ_i as the blocks along the diagonal with multiplicity n_i .

The following proposition shows that matrix entries of irreducible representations are 'orthogonal' with respect to a *symmetric bilinear form*, unless they are conjugates of each other – in which case the corresponding product is the inverse of the dimension of the representation.

PROPOSITION 2.9. If ρ and τ are two non-isomorphic irreducible representations of G then for any i, j, k, ℓ we have

$$\langle (\rho)_{ij} \mid (\tau)_{k\ell} \rangle_G = 0, \quad (3)$$

where $\langle f_1 \mid f_2 \rangle_G := \frac{1}{|G|} \sum_{g \in G} f_1(g) f_2(g^{-1})$ (called a "symmetric bilinear form"). Also,

$$\langle (\rho)_{ij} \mid (\rho)_{k\ell} \rangle_G = \frac{\delta_{i\ell} \delta_{jk}}{\dim(\rho)}, \quad (4)$$

where δ_{ij} is the delta-function which is 1 if $i = j$ and 0 otherwise.

2.2.2 Fourier analysis on non-abelian group. In this paper, we will be interested in studying $L^2(G)$, the space of functions from a finite group G to the complex numbers $\mathbb{C}$.

Definition 2.10. Define the inner product $\langle \cdot, \cdot \rangle_{L^2(G)}$ on $L^2(G)$ by

$$\langle f, g \rangle_{L^2(G)} = \mathbf{E}_{x \in G} [f(x) \overline{g(x)}].$$

We can define a character for every representation of a group.

Definition 2.11. The character of a representation ρ is the function $\chi_\rho : G \rightarrow \mathbb{C}$ defined by $\chi_\rho(g) = \text{tr}(\rho(g))$.

The following proposition shows that the characters corresponding to the irreducible representations of a group are orthogonal to each other.

PROPOSITION 2.12 (ORTHOGONALITY OF CHARACTERS). For $\rho, \tau \in \text{Irrep}(G)$, we have

$$\frac{1}{|G|} \sum_{g \in G} \chi_\rho(g) \overline{\chi_\tau(g)} = \begin{cases} 1 & \rho_V \cong \tau_V, \\ 0 & \text{otherwise.} \end{cases}$$

We use Proposition 2.9 many times in the proof. For convenience, we note an important identity that follows from Proposition 2.9 (by setting τ to be the trivial map $\{1\}$).

PROPOSITION 2.13. If $\rho \in \text{Irrep}(G) \setminus \{1\}$, $\sum_{g \in G} \rho(g) = 0$.

We have a following proposition. It also shows that the maximum dimension of any irreducible representation of G is at most $\sqrt{|G|}$.

PROPOSITION 2.14.

$$\sum_{\rho \in \text{Irrep}(G)} \dim(\rho) \chi_\rho(g) = \begin{cases} |G| & g = 1_G, \\ 0 & \text{otherwise.} \end{cases}$$

This implies the following;

$$\sum_{\rho \in \text{Irrep}(G)} \dim(\rho)^2 = |G|.$$

Definition 2.15. For two functions $f, g \in L^2(G)$ their *convolution* $f * g \in L^2(G)$ is defined as

$$(f * g)(x) := \mathbb{E}_{y \in G} [f(y)g(y^{-1}x)].$$

For an abelian group, any function $f : G \rightarrow \mathbb{C}$ can be written as linear combinations of characters, i.e., the characters span the whole space $L^2(G)$. However, for non-abelian groups, characters form an orthonormal basis only for the set of *class functions* – maps which are constant on *conjugacy classes*. A conjugacy class in G is a nonempty subset H of G such that the following two conditions hold: Given any $x, y \in H$, there exists $g \in G$ such that $gxg^{-1} = y$, and if $x \in H$ and $g \in G$ then $gxg^{-1} \in H$. Since this is an equivalence class, any group is a collection of disjoint conjugacy classes.

As in the abelian case, we can understand operations like inner product, convolution etc., using the Fourier transform which is defined as follows:

Definition 2.16. For a function $f \in L^2(G)$, define the Fourier transform of f to be the element $\hat{f} \in \prod_{\rho \in \text{Irrep}(G)} \text{End } \rho_V$ given by

$$\hat{f}(\rho) = \mathbb{E}_{x \in G} [f(x)\rho(x)] \in \text{End } \rho_V.$$

Definition 2.17. Let V be a finite-dimensional complex inner product space. Define an inner product $\langle \cdot, \cdot \rangle_{\text{End } V}$ on $\text{End } V$ by

$$\langle A, B \rangle_{\text{End } V} = \text{tr}(AB^*).$$

We can now state the Fourier inversion theorem.

PROPOSITION 2.18 (FOURIER INVERSION THEOREM). For $f \in L^2(G)$ we have

$$f(x) = \sum_{\rho \in \text{Irrep}(G)} \dim(\rho) \cdot \langle \hat{f}(\rho), \rho(x) \rangle_{\text{End } \rho_V}.$$

We have the following simple identities (See [18] for the proofs).

PROPOSITION 2.19 (PLANCHEREL'S IDENTITY).

$$\langle f, g \rangle_{L^2(G)} = \sum_{\rho \in \text{Irrep}(G)} \dim(\rho) \cdot \langle \hat{f}(\rho), \hat{g}(\rho) \rangle_{\text{End } \rho_V}.$$

PROPOSITION 2.20 (PARSEVAL'S IDENTITY).

$$\mathbb{E}_{x \in G} [|f(x)|^2] = \sum_{\rho \in \text{Irrep}(G)} \dim(\rho) \cdot \|\hat{f}(\rho)\|_{\text{HS}}^2,$$

where $\|A\|_{\text{HS}} := \sqrt{\langle A, A \rangle_{\text{End } V}} = \sqrt{\text{tr}(AA^*)} = \sqrt{\sum_{ij} |A_{ij}|^2}$.

Note that the norm $\|\cdot\|_{\text{HS}}$ satisfies a triangle inequality.

CLAIM 1. $\|AB\|_{\text{HS}} \leq \|A\|_{\text{HS}} \cdot \|B\|_{\text{HS}}$.

PROOF. $\|AB\|_{\text{HS}}^2 = \sum_{ij} |(AB)_{ij}|^2 \leq \sum_{ij} \left(\sum_k |A_{ik}| |B_{kj}| \right)^2$. Using the Cauchy–Schwartz inequality on the inner sum,

$$\begin{aligned} \|AB\|_{\text{HS}}^2 &\leq \sum_{ij} \left(\sum_k |A_{ik}|^2 \right) \left(\sum_{\ell} |B_{\ell j}|^2 \right) = \sum_{ijk\ell} |A_{ik}|^2 |B_{\ell j}|^2 \\ &= \left(\sum_{ik} |A_{ik}|^2 \right) \left(\sum_{\ell j} |B_{\ell j}|^2 \right) = \|A\|_{\text{HS}}^2 \cdot \|B\|_{\text{HS}}^2. \end{aligned}$$

CLAIM 2. Let A be any matrix and U be any unitary matrix, then $\|UA\|_{\text{HS}} = \|A\|_{\text{HS}}$.

PROOF. Let V be a unitary matrix which converts U to the identity matrix, i.e., $VUV^* = I$. Since the change of basis does not change the $\|\cdot\|_{\text{HS}}$, we have

$$\begin{aligned} \|UA\|_{\text{HS}} &= \|VUAV^*\|_{\text{HS}} = \|VUV^*VAV^*\|_{\text{HS}} = \|IVAV^*\|_{\text{HS}} \\ &= \|A\|_{\text{HS}}. \end{aligned}$$

PROPOSITION 2.21 (CONVOLUTION THEOREM). For $f, g \in L^2(G)$ we have

$$f * g(\rho) = \hat{f}(\rho)\hat{g}(\rho).$$

2.3 Important Claims

In this section, we prove a few statements that will be used in the soundness analysis. The following claim shows that the character functions always come in 'pairs' with respect to the complex conjugation.

CLAIM 3. Let G be any non abelian group. For every $\rho \in \text{Irrep}(G)$, such that $\dim(\rho) = 1$, there exists $\bar{\rho} \in \text{Irrep}(G)$ with $\dim(\bar{\rho}) = 1$ such that

$$\chi_{\rho}(g) = \overline{\chi_{\bar{\rho}}(g)}, \quad \forall g \in G.$$

PROOF. We claim that the set of characters corresponding to dimension 1 irreducible representations of G forms a group under point-wise multiplication. This will be enough to show the claim.

Let $G' = G/[G, G]$ be the abelian quotient group. Assume ρ is a degree 1 representation of G . Then it satisfies $\rho(a)\rho(b) = \rho(ab)$ for all $a, b \in G$. Define a map $\Gamma_{\rho} : G' \rightarrow \mathbb{C}$ as $\Gamma_{\rho}(g') = \rho(g)$ where $g' = g[G, G]$. This is a well defined map as

$$\begin{aligned} \rho(aba^{-1}b^{-1}) &= \rho(a)\rho(b)\rho(a^{-1})\rho(b^{-1}) \\ &= \rho(a)\rho(a^{-1})\rho(b)\rho(b^{-1}) = 1. \end{aligned}$$

Thus, the map ρ is constant on every coset of $[G, G]$ and hence Γ_{ρ} is well defined. The set of all $\{\Gamma_{\rho} \mid \rho \in \text{Irrep}(G), \dim(\rho) = 1\}$ is the set of all the multiplicative characters of the abelian group G' and hence form a group under coordinate-wise multiplication. There is a one-to-one correspondence between the coordinate wise multiplicative action of Γ_{ρ} 's and ρ 's. Thus, $\{\chi_{\rho} \mid \rho \in \text{Irrep}(G), \dim(\rho) = 1\}$ form a group under point-wise multiplication. $\square$

The following lemma shows that the direct sum decomposition of tensors of large dimension irreducible representations cannot contain overwhelming copies of a single dimension 1 representation. For the proof see the full version of this paper [6].

LEMMA 2.22. Let $\rho = \otimes_{k=1}^t \rho_{i_k}$ be a representation of G where each $\rho_{i_k} \in \text{Irrep}(G)$ and $\dim(\rho_{i_k}) \geq 2$ for all $k \in [t]$. Suppose following is the decomposition of ρ into its irreducible components

$$\otimes_{k=1}^t \rho_{i_k} \cong \oplus_{\ell=1}^r n_{j_{\ell}} \rho_{j_{\ell}},$$

where $\rho_{j_{\ell}}$ and $\rho_{j_{\ell'}}$ are distinct for every $\ell \neq \ell'$. Then for all $\ell \in [r]$, $n_{j_{\ell}} \leq \left(1 - \frac{1}{|G|}\right) \dim(\rho)$.

We have a following corollary that follows from the previous lemma.

COROLLARY 2.23. *Let $\rho = \otimes_{k=1}^t \rho_{i_k}$ be a representation of G where each $\rho_{i_k} \in \text{Irrep}(G)$ for all $k \in [t]$, and $\dim(\rho) \geq 2$. Suppose following is the decomposition of ρ into its irreducible components*

$$\otimes_{k=1}^t \rho_{i_k} \cong \oplus_{\ell=1}^r n_{j_\ell} \rho_{j_\ell},$$

where ρ_{j_ℓ} and $\rho_{j_{\ell'}}$ are distinct for every $\ell \neq \ell'$. Then for all $\ell \in [r]$, $n_{j_\ell} \leq \left(1 - \frac{1}{|G|}\right) \dim(\rho)$.

PROOF. Assume without loss of generality that the first t' terms are all the dimension 1 representations in the tensor product ρ . Now, the (tensor) product of dimension 1 representations is also a dimension 1 representation of G . Suppose $\tau = (\otimes_{k=1}^{t'} \rho_{i_k})$ where $\dim(\tau) = 1$. We can write ρ as:

$$\begin{aligned} \rho &= \otimes_{k=1}^t \rho_{i_k} = (\otimes_{k=1}^{t'} \rho_{i_k}) \otimes \rho_{i_{t'+1}} \otimes (\otimes_{k=t'+2}^t \rho_{i_k}) \\ &= (\tau \rho_{i_{t'+1}}) \otimes (\otimes_{k=t'+2}^t \rho_{i_k}). \end{aligned}$$

Now, $\tau \rho_{i_{t'+1}}$ itself is an irreducible representation of G of dimension at least 2. Therefore, the conclusion follows from Lemma 2.22. $\square$

2.4 Functions on G^n

For any non-abelian group G and $n \geq 1$, we have a group G^n where the group operation is defined coordinate wise. The irreducible representations of G^n are precisely those representations obtained by taking tensor products of n irreducible representations of G .

PROPOSITION 2.24 ([18]). *The set of irreducible representations of G^n is given by*

$$\text{Irrep}(G^n) = \{\alpha \mid \alpha = \otimes_{i \in [n]} \rho_i \text{ where } \rho_i \in \text{Irrep}(G)\}.$$

We denote α by the corresponding tuple $(\rho_1, \rho_2, \dots, \rho_n)$. We define the weight of a representation $\alpha = (\rho_1, \rho_2, \dots, \rho_n)$ (denoted by $|\alpha|$) to be the number of non-trivial representations in $(\rho_1, \rho_2, \dots, \rho_n)$.

We will be working with functions $f : G^n \rightarrow G$ which are *folded*. f is said to be folded if $f(c\mathbf{x}) = cf(\mathbf{x})$ for all $c \in G$ and $\mathbf{x} \in G^n$. The following claim shows that for all functions $g(\mathbf{x}) := \rho(f(\mathbf{x}))_{ij}$ where $\dim(\rho) \geq 2$ and $1 \leq i, j \leq \dim(\rho)$, all the Fourier coefficients corresponding to representations of dimension 1 are zero, if f is folded.

LEMMA 2.25. *Let $f : G^n \rightarrow G$ be any folded function and $g(\mathbf{x}) := \rho(f(\mathbf{x}))_{ij}$ where $\rho \in \text{Irrep}(G)$, $\dim(\rho) \geq 2$ and $1 \leq i, j \leq \dim(\rho)$. Let α be any representation of G^n such that $\dim(\alpha) = 1$, then $\hat{g}(\alpha) = 0$.*

PROOF. Recall, for any $\mathbf{x} \in G^n$, $\alpha(\mathbf{x})$ is a scalar as $\dim(\alpha) = 1$. f is folded which means that $f(c\mathbf{x}) = cf(\mathbf{x})$ for all $c \in G$ and $\mathbf{x} \in G^n$. For $c \in G$, let $\mathbf{c} \in G^n$ be such that every coordinate of $\mathbf{c}$ is c . Since $\rho(\cdot)$ has dimension at least 2, in the following analysis, we use $[\rho(\cdot)]$

to denote that matrix of linear transformation for clarity.

$$\begin{aligned} \hat{g}(\alpha) &= \mathbb{E}_{\mathbf{x}} [g(\mathbf{x}) \alpha(\mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x}} [\rho(f(\mathbf{x}))_{ij} \cdot \alpha(\mathbf{x})] \\ &= \frac{1}{|G|} \mathbb{E}_{\mathbf{x}} \left[\sum_{c \in G} [\rho(f(c\mathbf{x}))]_{ij} \cdot \alpha(c\mathbf{x}) \right] \\ &= \frac{1}{|G|} \mathbb{E}_{\mathbf{x}} \left[\left(\sum_{c \in G} \alpha(c) [\rho(c)] \right) \cdot [\rho(f(\mathbf{x}))]_{ij} \alpha(\mathbf{x}) \right]. \end{aligned}$$

Now, for $\alpha \in \text{Irrep}(G^n)$, let $\tilde{\alpha} \in \text{Irrep}(G^n)$ be the dimension 1 representation satisfying the condition in Claim 3. We have:

$$\begin{aligned} \sum_{c \in G} \alpha(c) [\rho(c)] &= \sum_{c \in G} \overline{\tilde{\alpha}(c)} \cdot [\rho(c)] \\ &= \sum_{c \in G} \tilde{\alpha}(c^{-1}) \cdot [\rho(c)] \\ &= \sum_{c \in G} (\otimes_{i=1}^n \tilde{\alpha}_i(c^{-1})) \cdot [\rho(c)] \\ &= \sum_{c \in G} \tau(c^{-1}) \cdot [\rho(c)] \quad \dim(\tau) = 1 \\ &= 0, \quad (\text{Using Proposition 2.13}) \end{aligned}$$

where in the second last step, we used the fact that the product of dimension 1 representations $(\otimes_{i=1}^n \tilde{\alpha}_i)$ of G is itself a dimension 1 representation (τ) of G . Therefore, $\hat{g}(\alpha) = 0$. $\square$

Fix any surjective projection map $\pi : [R] \rightarrow [L]$ for some $R \geq L$. Consider the following subgroup of G^R given by the elements

$$\{(x \circ \pi) \in G^R \mid x \in G^L\},$$

where $(x \circ \pi)_i = x_{\pi(i)}$. Let us denote this group by $\pi(G^R)$. Note that this group is isomorphic to G^L . Thus, any representation $\alpha \in \text{Irrep}(G^R)$ (which is a representation of G^L using Fact 1), can be decomposed into irreducible representations of G^L .

The following lemma says that if α satisfies certain property, then for each irreducible representation occurring in the decomposition, either its dimension is large or its multiplicity is small. For the proof see the full version of this paper [6].

LEMMA 2.26. *Let $\pi : [R] \rightarrow [L]$ be any surjective projection map. Let $\varepsilon_0 \in (0, \frac{1}{2}]$ and $c \geq 10|G| \log(\frac{1}{\varepsilon_0})$. Suppose $\alpha \in \text{Irrep}(G^R)$,*

$$\alpha = \bigotimes_{i=1}^R \rho_i = \bigotimes_{\ell=1}^L \underbrace{\left(\bigotimes_{j \in \pi_{\ell}^{-1}(\ell)} \rho_j \right)}_{=: B_\ell}$$

such that number of ℓ with $\dim(B_\ell) \geq 2$ is at least c . If $\alpha \cong \oplus_m n_m \beta_m$ is the decomposition of α into irreducible representations of $\pi(G^L) \cong G^L$, then for every m either $\dim(\beta_m) \geq c$ or $n_m \leq \varepsilon_0^2 \cdot \dim(\alpha)$.

2.5 Notations

Whenever possible, we use the notation α, β to denote the representations of group G^n and ρ, τ for group G . Also, we use bold letters $\mathbf{x}, \mathbf{c}$ to denote the elements of G^n .

For a representation $\alpha \in \text{Irrep}(G^n)$ where $\alpha = \otimes_{i=1}^n \rho_i$, we use the notation $\dim_{\geq k}(\alpha)$ to denote the number of $i \in [n]$ such that $\dim(\rho_i) \geq k$.

3 WARM-UP: DICTATORSHIP TEST

In this section, we analyze the dictatorship test where the test involves checking some linear equation over a non-abelian group. The analysis will highlight a few important differences between our test and the linearity test over abelian groups.

Fix a non-abelian group G . Let $f : G^n \rightarrow G$ be a function. A function is called a dictator function if it is of the form $f(\mathbf{x}) = x_i$ for some $i \in [n]$. We use $\bullet$ to denote the group operation. Consider the following 3-query dictatorship test for f :

- (1) Sample $\mathbf{a} = (a_1, a_2, \dots, a_n)$ from G^n uniformly at random.
- (2) Sample $\mathbf{b} = (b_1, b_2, \dots, b_n)$ from G^n uniformly at random.
- (3) Calculate $\mathbf{c} = (c_1, c_2, \dots, c_n)$ such that $c_i = b_i^{-1} a_i^{-1}$.
- (4) Check if $f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c}) = 1_G$.

Completeness is trivial: If f is an i^{th} dictator function, in other words, $f(x_1, x_2, \dots, x_n) = x_i$, then the test passes with probability 1. This is because we are essentially checking if $a_i \cdot b_i \cdot c_i = 1_G$, which is always true by the definition of c_i .

We analyze the soundness of the test. The following lemma says that if the test passes with some non-trivial probability then it must be the case that f (or a minor variation of f) has a low dimension Fourier coefficient whose Hilbert-Schmidt norm is large. The actual conclusion is somewhat stronger than this. In the next section, we will show that such a conclusion can be used to analyze the soundness of the final reduction (which is also presented in next section).

LEMMA 3.1. *Assume f is folded. For all $\varepsilon > 0$ and $\delta > 0$, if f passes the test with probability $\frac{1}{|G, G|} + \varepsilon$, then there exist $\rho \in \text{Irrep}(G)$ and $1 \leq i, j \leq \dim(\rho)$ such that for $h(\mathbf{x}) := \rho(f(\mathbf{x}))_{ij}$,*

$$\max_{\substack{\alpha, \\ \dim(\alpha) \geq 2, \\ \dim_{\geq 2}(\alpha) < \frac{1}{2\delta^2}}} \|\hat{h}(\alpha)\|_{\text{HS}} \geq \frac{\varepsilon}{|G|} - \delta.$$

PROOF. Using Proposition 2.14, we can express the probability that the test passes as follows:

$$\begin{aligned} \Pr[\text{Test passes}] &= \frac{1}{|G|} \sum_{\rho \in \text{Irrep}(G)} \dim(\rho) \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c}))] \\ &= \frac{1}{|G|} \sum_{\substack{\rho \in \text{Irrep}(G), \\ \dim(\rho)=1}} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c}))] \\ &\quad + \frac{1}{|G|} \sum_{\substack{\rho \in \text{Irrep}(G), \\ \dim(\rho) \geq 2}} \dim(\rho) \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c}))]. \end{aligned}$$

In the first summation, for any $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) = 1$, using the multiplicativity of the characters, we have

$$\begin{aligned} \chi_\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c})) &= \chi_\rho(f(\mathbf{a})) \chi_\rho(f(\mathbf{b})) \chi_\rho(f(\mathbf{c})) \\ &\leq |\chi_\rho(f(\mathbf{a}))| \cdot |\chi_\rho(f(\mathbf{b}))| \cdot |\chi_\rho(f(\mathbf{c}))| \\ &= 1. \quad (\text{unitary representations}) \end{aligned}$$

As the number of dimension 1 representations of a group G is equal to the size of the quotient $G/[G, G]$, we get

$$\begin{aligned} \Pr[\text{Test passes}] &\leq \frac{1}{|[G, G]|} + \frac{1}{|G|} \sum_{\substack{\rho \in \text{Irrep}(G), \\ \dim(\rho) \geq 2}} \dim(\rho) \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c}))]. \end{aligned} \quad (5)$$

Now, fix any $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) \geq 2$. For $1 \leq i, j \leq \dim(\rho)$, let $g_{ij} : G^n \rightarrow \mathbb{C}$ be defined as $g_{ij}(\mathbf{x}) := \rho(f(\mathbf{x}))_{ij}$. Using the definition of characters, we have

$$\begin{aligned} &\mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c}))] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\text{tr}(\rho(f(\mathbf{a}) \cdot f(\mathbf{b}) \cdot f(\mathbf{c})))] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\text{tr}(\rho(f(\mathbf{a})) \cdot \rho(f(\mathbf{b})) \cdot \rho(f(\mathbf{c}))) \quad (\rho \text{ is a homomorphism}) \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} \left[\sum_{1 \leq i, j, k \leq \dim(\rho)} \rho(f(\mathbf{a}))_{ij} \cdot \rho(f(\mathbf{b}))_{jk} \cdot \rho(f(\mathbf{c}))_{ki} \right] \\ &= \sum_{1 \leq i, j, k \leq \dim(\rho)} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [g_{ij}(\mathbf{a}) g_{jk}(\mathbf{b}) g_{ki}(\mathbf{c})] \\ &= \sum_{1 \leq i, j, k \leq \dim(\rho)} (g_{ij} * g_{jk} * g_{ki})(1_{G^n}). \end{aligned} \quad (6)$$

Since we assume that the test passes with probability $\frac{1}{|[G, G]|} + \varepsilon$, from Equation (5) and Equation (6) (and using $\dim(\rho) \leq \sqrt{|G|}$), we conclude that there exists ρ and $1 \leq i, j, k \leq \dim(\rho)$ such that

$$|(g_{ij} * g_{jk} * g_{ki})(1_{G^n})| \geq \frac{\varepsilon}{|G|}.$$

We now analyze the term $(g_{ij} * g_{jk} * g_{ki})(1_{G^n})$ for a fixed (i, j, k) . For the ease of notation, we write $h_1 := g_{ij}$, $h_2 := g_{jk}$ and $h_3 := g_{ki}$.

$$\begin{aligned} \frac{\varepsilon}{|G|} &\leq |(g_{ij} * g_{jk} * g_{ki})(1_{G^n})| \\ &= |(h_1 * h_2 * h_3)(1_{G^n})| \\ &= \left| \sum_{\alpha \in \text{Irrep}(G^n)} \dim(\alpha) \cdot \text{tr}(h_1 * h_2 * h_3(\alpha)) \right| \\ &\leq \sum_{\alpha \in \text{Irrep}(G^n)} \dim(\alpha) \cdot |\text{tr}(\hat{h}_1(\alpha) \hat{h}_2(\alpha) \hat{h}_3(\alpha))| \\ &= \sum_{\alpha \in \text{Irrep}(G^n)} \dim(\alpha) \cdot \left| \langle \hat{h}_1(\alpha) \hat{h}_2(\alpha), \hat{h}_3(\alpha)^* \rangle_{\text{End } \mathcal{V}} \right| \\ &\leq \sum_{\alpha \in \text{Irrep}(G^n)} \dim(\alpha) \cdot \|\hat{h}_1(\alpha) \hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)^*\|_{\text{HS}}. \end{aligned}$$

We now use Lemma 2.25 to conclude that for all $1 \leq i \leq 3$, $\hat{h}_i(\alpha) = 0$ if $\dim(\alpha) = 1$. Using this, we continue as follows:

$$\begin{aligned} &|(h_1 * h_2 * h_3)(1_G)| \\ &\leq \sum_{\substack{\alpha \in \text{Irrep}(G^n), \\ \dim(\alpha) \geq 2}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha) \hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)^*\|_{\text{HS}} \end{aligned}$$

$$= \sum_{\substack{\alpha, \\ \dim(\alpha) \geq 2}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}}.$$

Let $D := \frac{1}{2\delta^2}$. Now, we split the sum into two parts $|(h_1 * h_2 * h_3)(1_{G^n})| \leq \Theta_{\text{low}} + \Theta_{\text{high}}$ where

$$\Theta_{\text{low}} = \sum_{\substack{\alpha, \\ \dim(\alpha) \geq 2, \\ \dim_{\geq 2}(\alpha) < D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}},$$

and

$$\Theta_{\text{high}} = \sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}}.$$

3.1 Bounding Higher Order Terms

In this section, we show that the high degree terms can be upper bounded by a small constant, even though the *three* queries are perfectly correlated.

We bound Θ_{high} as follows:

$$\begin{aligned} \Theta_{\text{high}} &= \sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}} \\ &\leq \sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}} \|\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}} \quad (\text{Claim 1}) \\ &\leq \frac{1}{\sqrt{2D}} \sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha)^{3/2} \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}} \|\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}}. \end{aligned}$$

Here, we used that fact that all the representations α of G with $\dim_{\geq 2}(\alpha) \geq D$ have dimensions at least $2D$. At this point, we would like to point out the main source of effectively bounding the higher order terms. It is the size of $\dim(\alpha)$ in the summation. In Gowers' [12] proof, a similar expression appears in the analysis, with the same condition that all the representations in the summation have large dimension. It is in some sense the main difference between the abelian and the non-abelian setting (both in this work and Gowers'), similar to the Equation (1) mentioned in the introduction.

Now, using Cauchy–Schwartz inequality,

$$\begin{aligned} \Theta_{\text{high}} &\leq \frac{1}{\sqrt{2D}} \sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha)^{3/2} \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}} \|\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}} \\ &\leq \frac{1}{\sqrt{2D}} \left(\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}}^2 \right)^{1/2} \\ &\quad \left(\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha)^2 \cdot \|\hat{h}_2(\alpha)\|_{\text{HS}}^2 \|\hat{h}_3(\alpha)\|_{\text{HS}}^2 \right)^{1/2} \\ &\leq \frac{1}{\sqrt{2D}} \left(\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}}^2 \right)^{1/2}. \end{aligned}$$

$$\begin{aligned} &\left(\left(\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_2(\alpha)\|_{\text{HS}}^2 \right) \cdot \left(\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) \geq D}} \dim(\alpha) \cdot \|\hat{h}_3(\alpha)\|_{\text{HS}}^2 \right) \right)^{1/2} \\ &\leq \frac{1}{\sqrt{2D}} \cdot \|h_1\|_{L^2(G^n)} \cdot \|h_2\|_{L^2(G^n)} \cdot \|h_3\|_{L^2(G^n)}. \end{aligned}$$

(Using Proposition 2.20)

Since h_1 was defined as $h_1(\mathbf{x}) = \rho(f(\mathbf{x}))_{ij}$ where $\rho \in \text{Irrep}(G)$, $|h_1(\mathbf{x})| \leq 1$. Same is true for h_2 and h_3 , and hence all the norms are bounded by 1. Therefore,

$$\Theta_{\text{high}} \leq \frac{1}{\sqrt{2D}} = \delta.$$

3.2 Bounding Lower Order Terms

It remain to show that Θ_{low} is related to the Fourier mass of h_3 on the low dimension representations.

$$\begin{aligned} \Theta_{\text{low}} &= \sum_{\substack{\alpha, \\ \dim(\alpha) \geq 2, \\ \dim_{\geq 2}(\alpha) < D}} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \|\hat{h}_3(\alpha)\|_{\text{HS}} \\ &\leq \max_{\substack{\alpha, \\ \dim(\alpha) \geq 2, \\ \dim_{\geq 2}(\alpha) < D}} \|\hat{h}_3(\alpha)\|_{\text{HS}} \cdot \left(\sum_{\alpha} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \right). \end{aligned}$$

Using the Cauchy–Schwartz inequality, we can upper bound the summation by 1 as follows:

$$\begin{aligned} &\sum_{\alpha} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\hat{h}_2(\alpha)\|_{\text{HS}} \\ &\leq \sum_{\alpha} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}} \|\hat{h}_2(\alpha)\|_{\text{HS}} \\ &\leq \left(\sum_{\alpha} \dim(\alpha) \cdot \|\hat{h}_1(\alpha)\|_{\text{HS}}^2 \right)^{1/2} \cdot \left(\sum_{\alpha} \dim(\alpha) \cdot \|\hat{h}_2(\alpha)\|_{\text{HS}}^2 \right)^{1/2} \\ &= \|h_1\|_2 \cdot \|h_2\|_2 \quad (\text{Proposition 2.20}) \\ &\leq 1, \end{aligned}$$

where the last inequality uses the fact that $|h_1(\mathbf{x})|, |h_2(\mathbf{x})| \leq 1$ for all $\mathbf{x} \in G^n$. Using the upper bound on Θ_{high} , we have $\Theta_{\text{low}} \geq \frac{\varepsilon}{|G|} - \delta$.

Therefore, we get

$$\max_{\substack{\alpha, \\ \dim(\alpha) \geq 2, \\ \dim_{\geq 2}(\alpha) < D}} \|\hat{h}_3(\alpha)\|_{\text{HS}} \geq \left(\frac{\varepsilon}{|G|} - \delta \right).$$

□

4 MAIN REDUCTION

In this section, we prove Theorem 1.1. We give a reduction from an instance of a LABEL-COVER, $\mathcal{H} = (\mathcal{U}, \mathcal{V}, E, [L], [R], \{\pi_e\}_{e \in E})$ as in Definition 2.1, to a 3-LIN instance $\mathcal{I}$ over a non-abelian group G .

The set of variables of $\mathcal{I}$ is $(\mathcal{U} \times G^L) \cup (\mathcal{V} \times G^R)$. Any assignment to the instance $\mathcal{I}$ is given by a set of functions $f_u : G^L \rightarrow G$ and $f_v : G^R \rightarrow G$ for each $u \in \mathcal{U}$ and $v \in \mathcal{V}$. We further assume that these functions are *folded*.

The distribution of the 3-LIN constraints in $\mathcal{I}$ is given by the following test:

- (1) Choose an edge $e(u, v) \in E$ of $\mathcal{H}$ uniformly at random.
- (2) Sample $\mathbf{a} = (a_1, a_2, \dots, a_R)$ from G^R uniformly at random.
- (3) Sample $\mathbf{b} = (b_1, b_2, \dots, b_L)$ from G^L uniformly at random.
- (4) Let $\mathbf{c} = (c_1, c_2, \dots, c_R)$ be such that $c_i = (b \circ \pi_{uv})_i^{-1} \cdot a_i^{-1}$, here $\mathbf{x} \circ \pi \in G^R$ is the string defined as $(\mathbf{x} \circ \pi)_i := x_{\pi(i)}$ for $i \in [R]$.
- (5) Test if $f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c}) = 1_G$.

The value of the instance $\text{val}(\mathcal{I})$ is the maximum probability that the above test is satisfied, where the maximum is over all folded functions $\{f_v\}_{v \in \mathcal{V}}, \{f_u\}_{u \in \mathcal{U}}$.

4.1 Analysis

LEMMA 4.1 (COMPLETENESS). *If $\mathcal{H}$ is a satisfiable instance of LABEL-COVER, then $\text{val}(\mathcal{I}) = 1$.*

PROOF. Fix a satisfying assignment $\ell : \mathcal{U} \rightarrow [L], \ell : \mathcal{V} \rightarrow [R]$ of $\mathcal{H}$. Consider the long code encoding of the labeling $\ell : f_v(\mathbf{x}) = x_{\ell(v)}$ and $f_u(\mathbf{x}) = x_{\ell(u)}$, for every $v \in \mathcal{V}$ and $u \in \mathcal{U}$. We show that this assignment to $\mathcal{I}$ satisfies all the constraints.

$$\begin{aligned} f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c}) &= a_{\ell(v)} \cdot b_{\ell(u)} \cdot c_{\ell(v)} \\ &= a_{\ell(v)} \cdot b_{\ell(u)} \cdot (b \circ \pi_{uv})_{\ell(v)}^{-1} \cdot a_{\ell(v)}^{-1} \\ &= a_{\ell(v)} \cdot b_{\ell(u)} \cdot b_{\pi_{uv}(\ell(v))}^{-1} \cdot a_{\ell(v)}^{-1} \\ &= a_{\ell(v)} \cdot b_{\ell(u)} \cdot b_{\ell(u)}^{-1} \cdot a_{\ell(v)}^{-1} \\ &\quad (\pi_{uv}(\ell(v)) = \ell(u)) \\ &= 1_G. \end{aligned}$$

□

We now prove the main soundness lemma. Note that Lemma 4.1 and Lemma 4.2 along with the NP-hardness of LABEL-COVER from Theorem 2.2 for large enough r imply our main theorem Theorem 1.1 for any constant $\varepsilon > 0$.

LEMMA 4.2 (SOUNDNESS). *Let $\delta \in (0, 1)$. Let C be a constant such that $C^{-d_0/2} \leq \frac{\delta^2}{12|G|^6}$, where d_0 is the constant from Theorem 2.2. If $\mathcal{H}$ is at most $\frac{\delta^2}{10|G|^{10C}}$ -satisfiable, then $\text{val}(\mathcal{I}) \leq \frac{1}{|G|} + \delta$.*

PROOF. Fix any assignment to the instance $\mathcal{I}$ given by a set of functions $f_u : G^L \rightarrow G$ and $f_v : G^R \rightarrow G$ for each $u \in \mathcal{U}$ and $v \in \mathcal{V}$. The value of the instance for this assignment is given by:

$$\begin{aligned} \text{val}(\mathcal{I}) &= \mathbb{E}_{e(u,v) \in E} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} \left[\frac{1}{|G|} \sum_{\rho \in \text{Irrep}(G)} \dim(\rho) \chi_\rho(f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c})) \right] \end{aligned}$$

For any $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) = 1$, we have

$$\begin{aligned} \chi_\rho(f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c})) &= \chi_\rho(f_v(\mathbf{a})) \cdot \chi_\rho(f_u(\mathbf{b})) \cdot \chi_\rho(f_v(\mathbf{c})) \\ &\leq |\chi_\rho(f_v(\mathbf{a}))| \cdot |\chi_\rho(f_u(\mathbf{b}))| \cdot |\chi_\rho(f_v(\mathbf{c}))| \\ &= 1. \quad (\text{unitary representations}) \end{aligned}$$

As the number of dimension 1 representations of a group G is equal to the size of the quotient $G/[G, G]$, we get

$$\begin{aligned} \text{val}(\mathcal{I}) &\leq \frac{1}{|G|} + \sum_{\substack{\rho \in \text{Irrep}(G), \\ \dim(\rho) \geq 2}} \left| \mathbb{E}_{e(u,v) \in E} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c}))] \right|. \end{aligned}$$

The lemma follows from the following Claim 4. □

CLAIM 4. *If $\mathcal{H}$ is at most $\frac{\delta^2}{10|G|^{10C}}$ -satisfiable, then for every $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) \geq 2$,*

$$\left| \mathbb{E}_{e(u,v) \in E} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c}))] \right| \leq \frac{\delta}{|G|}.$$

PROOF. Fix any $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) \geq 2$. Let

$$\Theta := \mathbb{E}_{e(u,v) \in E} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c}))].$$

We first look at the inner expectation. For $1 \leq p, q \in \dim(\rho)$, let $g_{pq} : G^R \rightarrow \mathbb{C}$ be defined as $g_{pq}(\mathbf{x}) := \rho(f_v(\mathbf{x}))_{pq}$. Also, let $h_{pq} : G^L \rightarrow \mathbb{C}$ be defined as $h_{pq}(\mathbf{y}) := \rho(f_u(\mathbf{y}))_{pq}$. We have

$$\begin{aligned} &\mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\chi_\rho(f_v(\mathbf{a}) f_u(\mathbf{b}) f_v(\mathbf{c}))] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\text{tr}(\rho(f_v(\mathbf{a}) \cdot f_u(\mathbf{b}) \cdot f_v(\mathbf{c})))] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [\text{tr}(\rho(f_v(\mathbf{a})) \cdot \rho(f_u(\mathbf{b})) \cdot \rho(f_v(\mathbf{c}))) \quad (\rho \text{ is a homomorphism}) \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} \left[\sum_{1 \leq p, q, r \leq \dim(\rho)} \rho(f_v(\mathbf{a}))_{pq} \cdot \rho(f_u(\mathbf{b}))_{qr} \cdot \rho(f_v(\mathbf{c}))_{rp} \right] \\ &= \sum_{1 \leq p, q, r \leq \dim(\rho)} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [g_{pq}(\mathbf{a}) \cdot h_{qr}(\mathbf{b}) \cdot g_{rp}(\mathbf{c})]. \end{aligned}$$

We now analyze the term $\Theta_{p,q,r}^e := \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [g_{pq}(\mathbf{a}) h_{qr}(\mathbf{b}) g_{rp}(\mathbf{c})]$ for a fixed (p, q, r) . For the ease of notations, we write $g' := g_{pq}$, $h := h_{qr}$ and $g := g_{rp}$. Also, we use π for π_{uv} .

$$\mathbb{E}_{\mathbf{a}, \mathbf{b}} [g'(\mathbf{a}) \cdot g((\mathbf{b} \circ \pi)^{-1} \cdot \mathbf{a}^{-1})] = \mathbb{E}_{\mathbf{b}} [(g * g')((\mathbf{b} \circ \pi)^{-1})].$$

We now bound the expectation as follows.

$$\begin{aligned} &\mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [g_{pq}(\mathbf{a}) \cdot h_{qr}(\mathbf{b}) \cdot g_{rp}(\mathbf{c})] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{c}} [g'(\mathbf{a}) \cdot h(\mathbf{b}) \cdot g((\mathbf{b} \circ \pi)^{-1} \cdot \mathbf{a}^{-1})] \\ &= \mathbb{E}_{\mathbf{b}} [(g * g')(\mathbf{b}^{-1} \circ \pi) \cdot h(\mathbf{b})] \\ &= \mathbb{E}_{\mathbf{b}} [(g * g')(\mathbf{b} \circ \pi) \cdot h(\mathbf{b}^{-1})] \\ &= \mathbb{E}_{\mathbf{b}} \left[\frac{(\sum_{\alpha} \dim(\alpha) \text{tr}(\hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi))) \cdot (\sum_{\beta} \dim(\beta) \text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1})))}{(\sum_{\alpha} \dim(\alpha) \text{tr}(\hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi))) \cdot (\sum_{\beta} \dim(\beta) \text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1})))} \right] \\ &= \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2}} \text{Term}^e(\alpha, \beta), \end{aligned}$$

where

$$\text{Term}^e(\alpha, \beta) := \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\text{tr}(\hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi)) \cdot \text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1})) \right].$$

The last step uses the fact that the functions g, g' and h satisfy the condition of Lemma 2.25 and hence $\hat{g}(\alpha) = 0$ if $\dim(\alpha) = 1$ (same for $\hat{g}'(\alpha)$ and $\hat{h}(\beta)$).

We now break the sum into two parts:

$$\begin{aligned} \Theta_{p,q,r}^e(\text{low}) &:= \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) \leq C}} \text{Term}^e(\alpha, \beta), \\ \Theta_{p,q,r}^e(\text{high}) &:= \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) > C}} \text{Term}^e(\alpha, \beta). \end{aligned}$$

Recall, $\dim_{\geq 2}(\alpha)$ denotes the number of representations in $\alpha = (\rho_1, \rho_2, \dots, \rho_R)$ which are of dimensions at least 2. With these notations, we have

$$\Theta := \sum_{p,q,r} \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{low})] + \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{high})].$$

The upper bound on Θ follows from Claim 5 and Claim 6 and triangle inequality (and also noting that p, q and r take at most $\sqrt{G}$ distinct values). $\square$

CLAIM 5. If $\mathcal{H}$ is at most $\frac{\delta^2}{10|G|^{10C}}$ -satisfiable, then for every $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) \geq 2$, and every $1 \leq p, q, r \leq \dim(\rho)$,

$$\left| \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{low})] \right| \leq \frac{\delta}{2|G|^3}.$$

CLAIM 6. Let C be a constant such that $C^{-d_0/2} \leq \frac{\delta^2}{12|G|^6}$, where d_0 is the constant from Theorem 2.2. For every $1 \leq p, q, r \leq \dim(\rho)$,

$$\left| \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{high})] \right| \leq \frac{\delta}{2|G|^3}.$$

4.1.1 Bounding the low terms.

CLAIM 7. (Restatement of Claim 5) If $\mathcal{H}$ is at most $\frac{\delta^2}{10|G|^{10C}}$ -satisfiable, then for every $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) \geq 2$, and every $1 \leq p, q, r \leq \dim(\rho)$,

$$\left| \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{low})] \right| \leq \frac{\delta}{2|G|^3}.$$

PROOF. Fix any $\rho \in \text{Irrep}(G)$ such that $\dim(\rho) \geq 2$. Assume towards contradiction that there exists $1 \leq p, q, r \leq \dim(\rho)$ such that

$$\mathbb{E}_{e(u,v) \in E} [|\Theta_{p,q,r}^e(\text{low})|] \geq \left| \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{low})] \right| > \frac{\delta}{2|G|^3}.$$

We show that in this case, $\mathcal{H}$ has a $> \frac{\delta^2}{10|G|^{10C}}$ -satisfying assignment which is a contradiction. Consider the term $\text{Term}^e(\alpha, \beta)$ when $\alpha = (\rho_1, \rho_2, \dots, \rho_R)$ and $\beta = (\tau_1, \tau_2, \dots, \tau_L)$.

$$\begin{aligned} \text{Term}^e(\alpha, \beta) &= \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\text{tr}(\hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi)) \cdot \text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1})) \right] \\ &= \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\sum_{i,k} (\hat{g}(\alpha) \hat{g}'(\alpha))_{ik} \alpha(\mathbf{b} \circ \pi)_{ki} \hat{h}(\beta)_{i'k'} \beta(\mathbf{b}^{-1})_{k'i'} \right] \\ &= \dim(\alpha) \dim(\beta) \sum_{i,k} (\hat{g}(\alpha) \hat{g}'(\alpha))_{ik} \hat{h}(\beta)_{i'k'} \mathbb{E}_{\mathbf{b}} [\alpha(\mathbf{b} \circ \pi)_{ki} \cdot \beta(\mathbf{b}^{-1})_{k'i'}], \end{aligned}$$

where (i, k) are the tuples $i = (i_1, i_2, \dots, i_R)$ and $k = (k_1, k_2, \dots, k_R)$ such that for all $\ell \in [R]$, $1 \leq i_\ell, j_\ell \leq \dim(\rho_\ell)$. Similarly, (i', k') are the tuples $i' = (i'_1, i'_2, \dots, i'_L)$ and $k' = (k'_1, k'_2, \dots, k'_L)$ such that for all $\ell' \in [L]$, $1 \leq i'_{\ell'}, j'_{\ell'} \leq \dim(\tau_{\ell'})$. Now,

$$\begin{aligned} \mathbb{E}_{\mathbf{b}} [\alpha(\mathbf{b} \circ \pi)_{ki} \cdot \beta(\mathbf{b}^{-1})_{k'i'}] &= \mathbb{E}_{\mathbf{b}} \left[\prod_{\ell'=1}^L \tau_{\ell'}(b_{\ell'}^{-1})_{k'_{\ell'} i'_{\ell'}} \prod_{\ell \in \pi^{-1}(\ell')} \rho_{\ell}(b_{\ell})_{k_{\ell} i_{\ell}} \right] \\ &= \prod_{\ell'=1}^L \mathbb{E}_{\mathbf{b}} \left[\tau_{\ell'}(b_{\ell'}^{-1})_{k'_{\ell'} i'_{\ell'}} \prod_{\ell \in \pi^{-1}(\ell')} \rho_{\ell}(b_{\ell})_{k_{\ell} i_{\ell}} \right]. \end{aligned}$$

Now suppose there exists an ℓ' such that for all $\ell \in \pi^{-1}(\ell')$, $\dim(\rho_{\ell}) = 1$. Since the product of dimension 1 representations is also a dimension 1 representation, for the expectation to be nonzero, $\dim(\tau_{\ell'})$ must be 1, by Proposition 2.9. Thus, if $\dim(\beta) \geq 2$, then there must exist an ℓ' such that $\dim(\tau_{\ell'}) \geq 2$ and an $\ell \in \pi^{-1}(\ell')$ such that $\dim(\rho_{\ell}) \geq 2$ (again, for the expectation to be non-zero). Therefore, we conclude that the terms that are nonzero in the expression Θ_{low} are all (α, β) such that for all $\ell' \in [L]$ whenever $\dim(\tau_{\ell'}) \geq 2$, there exists an $\ell \in \pi^{-1}(\ell')$ such that $\dim(\rho_{\ell}) \geq 2$. Let us define $\pi_{\geq 2}(\alpha) = \{\ell' \in [L] \mid \exists \ell \in \pi^{-1}(\ell'), \dim(\rho_{\ell}) \geq 2\}$ and $\beta_{\geq 2} = \{\ell' \mid \dim(\tau_{\ell'}) \geq 2\}$. We get

$$\begin{aligned} \Theta_{p,q,r}^e(\text{low}) &= \sum_{\substack{\alpha, \beta, \\ \dim(\alpha) \geq 2, \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) \leq C, \\ \beta_{\geq 2} \subseteq \pi_{\geq 2}(\alpha)}} \dim(\alpha) \dim(\beta) \sum_{i,k,i',k'} (\hat{g}(\alpha) \hat{g}'(\alpha))_{ik} \hat{h}(\beta)_{i'k'} \mathbb{E}_{\mathbf{b}} [\alpha(\mathbf{b} \circ \pi)_{ki} \cdot \beta(\mathbf{b}^{-1})_{k'i'}]. \end{aligned}$$

Now, let $F_{\alpha}^{ik} : G^L \rightarrow \mathbb{C}$ be the following function:

$$F_{\alpha}^{ik}(\mathbf{b}) := \alpha(\mathbf{b}^{-1} \circ \pi)_{ki}.$$

Note that,

$$\sum_k \|F_{\alpha}^{ik}\|_2^2 = \sum_k \mathbb{E}_{\mathbf{b}} [|\alpha(\mathbf{b}^{-1} \circ \pi)_{ki}|^2] = \mathbb{E}_{\mathbf{b}} \sum_k |\alpha(\mathbf{b}^{-1} \circ \pi)_{ki}|^2 = 1, \quad (7)$$

where the last equality uses the fact that the sum expression is exactly the norm of the column i of representation α , which is 1 ($\alpha(\cdot)$ is unitary). Using this one can show the following upper bound. For details we refer to the full version of this paper [6].

$$\begin{aligned} |\Theta_{p,q,r}^e(\text{low})|^2 &\leq |G|^C \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) \leq C, \\ \beta_{\geq 2} \subseteq \pi_{\geq 2}(\alpha)}} \dim(\alpha) \dim(\beta) \|\hat{g}'(\alpha)\|_{\text{HS}}^2 \|\hat{h}(\beta)\|_{\text{HS}}^2. \end{aligned}$$

We now show that if $\Theta_{p,q,r}^e(\text{low})$ is large for a typical e , then it can be used to get a good labeling to the Label Cover instance $\mathcal{H}$.

Randomized labeling. Consider the following randomized labeling. For each $v \in \mathcal{V}$, consider $g_{pq} : G^R \rightarrow \mathbb{C}$ which is defined as $g_{pq}(\mathbf{x}) := \rho(f_v(\mathbf{x}))_{pq}$. Select $\alpha = (\rho_1, \rho_2, \dots, \rho_R)$ with probability $\dim(\alpha) \|\hat{g}_{pq}(\alpha)\|_{\text{HS}}^2$. Select a uniformly random $\ell_v \in [R]$ such that $\dim(\rho_{\ell_v}) \geq 2$ and assign the label ℓ_v to v .

For each $u \in \mathcal{U}$, consider $h_{qr} : G^L \rightarrow \mathbb{C}$ be defined as $h_{qr}(\mathbf{y}) := \rho(f_u(\mathbf{y}))_{qr}$. Select $\beta = (\tau_1, \tau_2, \dots, \tau_L)$ with probability proportional to $\dim(\beta) \|\hat{h}_{qr}(\beta)\|_{\text{HS}}^2$. Select a uniformly random $\ell_u \in [L]$ such that $\dim(\tau_{\ell_u}) \geq 2$ and assign the label ℓ_u to u .

Now fix an edge $e(u, v)$. The probability p_e that this edge is satisfied by the randomized labeling, is lower bounded by:

$$\begin{aligned} p_e &\geq \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) \leq C, \\ \beta_{\geq 2} \subseteq \pi_{\geq 2}(\alpha)}} \dim(\alpha) \dim(\beta) \|\hat{g}'(\alpha)\|_{\text{HS}}^2 \|\hat{h}(\beta)\|_{\text{HS}}^2 \cdot \frac{1}{\dim_{\geq 2}(\alpha)} \\ &\geq \frac{1}{C} \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) \leq C, \\ \beta_{\geq 2} \subseteq \pi_{\geq 2}(\alpha)}} \dim(\alpha) \dim(\beta) \|\hat{g}'(\alpha)\|_{\text{HS}}^2 \|\hat{h}(\beta)\|_{\text{HS}}^2 \\ &\geq \frac{1}{C \cdot |G|^C} |\Theta_{p,q,r}^e(\text{low})|^2. \end{aligned}$$

Therefore the expected number of edges satisfied by the randomized labeling is lower bounded by

$$\begin{aligned} \mathbb{E}_{e \in E} [p_e] &\geq \mathbb{E}_e \left[\frac{1}{C \cdot |G|^C} |\Theta_{p,q,r}^e(\text{low})|^2 \right] \\ &\geq \frac{1}{C \cdot |G|^C} \mathbb{E}_e \left[|\Theta_{p,q,r}^e(\text{low})|^2 \right] \quad (\text{Using convexity}) \\ &\geq \frac{1}{C \cdot |G|^C} \cdot \frac{\delta^2}{4|G|^6} > \frac{\delta^2}{10|G|^{10C}}. \end{aligned}$$

Since the expected fraction of the edges that are satisfied is strictly greater than $\frac{\delta^2}{10|G|^{10C}}$, by conditional expectation, there exists a labeling to the Label Cover instance $\mathcal{H}$ that satisfies more than $\frac{\delta^2}{10|G|^{10C}}$ fraction of the edges, which is a contradiction. $\square$

4.1.2 Bounding the high terms. We now show the following claim:

CLAIM 8. (Restatement of Claim 6) Let C be a constant such that $C - d_0/2 \leq \frac{\delta^2}{12|G|^6}$, where d_0 is the constant from Theorem 2.2. For every $1 \leq p, q, r \leq \dim(\rho)$,

$$\left| \mathbb{E}_{e(u,v) \in E} [\Theta_{p,q,r}^e(\text{high})] \right| \leq \frac{\delta}{2|G|^3}.$$

PROOF. Recall,

$$\Theta_{p,q,r}^e(\text{high}) := \sum_{\substack{\alpha, \beta, \\ \dim(\alpha), \dim(\beta) \geq 2, \\ \dim_{\geq 2}(\alpha) > C}} \text{Term}^e(\alpha, \beta).$$

Let's analyze the expression $\Theta_{p,q,r}^e(\text{high})$ more carefully. For the notational convenience, we suppress the conditions on α, β and simply write the sum over pairs α, β . We will analyze the

complete sum with the extra conditions on α, β , once we simplify the expression.

Let $U(e, \alpha)$ be the transformation (i.e., change of basis) which takes a representation $\alpha(\cdot)$ and converts it into a direct sum of irreducible representations of $\pi_e(G^R) := \{\mathbf{x} \circ \pi_e \mid \mathbf{x} \in G^L\}$ which is a subgroup of G^R isomorphic to the group G^L . For simplicity, we denote this unitary matrix by U . Recall that the decomposition is unique.

We extend the definition of the *block diagonal matrices* to include any permutation of columns of a block diagonal matrix. For clarity, we call such general matrices *block matrices*. Note that with this extended definition, it still makes sense to talk about the 'blocks', except that the blocks are not contiguous and not necessarily along the diagonal. For a given (e, α) , we apply a column-permutation matrix $P(e, \alpha)$ to the block diagonal matrix $U\alpha U^*$. We will get back to the specific choice of $P(e, \alpha)$ later in the proof, but for now just write the permutation matrix as P for notational convenience.

$$\begin{aligned} \text{tr}(\hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi)) &= \text{tr}(U \hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi) U^*) \\ &\quad (\text{cyclic property of tr, and } UU^* = I) \\ &= \text{tr}(U \hat{g}(\alpha) \hat{g}'(\alpha) U^* U \alpha(\mathbf{b} \circ \pi) U^*) \\ &= \text{tr}(U \hat{g}(\alpha) \hat{g}'(\alpha) U^* P^{-1} P U \alpha(\mathbf{b} \circ \pi) U^*) \quad (8) \end{aligned}$$

In this last expression, $U\alpha(\mathbf{b} \circ \pi)U^*$ is a block diagonal matrix, whereas $P U \alpha(\mathbf{b} \circ \pi) U^*$ is a block matrix.

We reiterate that the identity in Equation (8) holds for any unitary matrix U and column-permutation matrix P . For a fixed (e, α) , we will be using an arbitrary fixed $U(e, \alpha)$ (any unitary transformation which converts the representation α into a block diagonal matrix). The choice of $P(e, \alpha)$ will be delicate and in Claim 9, we will show an existence of a permutation matrix $P(e, \alpha)$ using which we can bound $\Theta_{p,q,r}^e(\text{high})$ effectively.

From this point onward, the choice of the unitary matrix does not matter as long as it converts $\alpha(\cdot)$ into a block diagonal matrix (also the arrangement of blocks along the diagonal does not matter). We are going to suppress the use of U and write:

$$U \hat{g}(\alpha) = A(\alpha), \hat{g}'(\alpha) U^* P^{-1} = A'(\alpha) \text{ and } U \alpha(\mathbf{b} \circ \pi) U^* = B(\alpha)(\mathbf{b}).$$

Note that by Claim 2, the $\|\cdot\|_{\text{HS}}$ of the matrices are preserved, i.e., $\|A(\alpha)\|_{\text{HS}} = \|\hat{g}(\alpha)\|_{\text{HS}}$ and $\|A'(\alpha)\|_{\text{HS}} = \|\hat{g}'(\alpha)\|_{\text{HS}}$. Coming back to the task of simplifying the expression:

$$\begin{aligned} \Theta_{p,q,r}^e(\text{high}) &= \sum_{\alpha, \beta} \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\text{tr}(\hat{g}(\alpha) \hat{g}'(\alpha) \alpha(\mathbf{b} \circ \pi)) \cdot \text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1})) \right] \\ &= \sum_{\alpha, \beta} \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\frac{\text{tr}(U \hat{g}(\alpha) \hat{g}'(\alpha) U^* P^{-1} P U \alpha(\mathbf{b} \circ \pi) U^*)}{\text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1}))} \right] \\ &= \sum_{\alpha, \beta} \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\text{tr}(A(\alpha) A'(\alpha) P B(\alpha)(\mathbf{b})) \cdot \text{tr}(\hat{h}(\beta) \beta(\mathbf{b}^{-1})) \right] \\ &= \sum_{\alpha, \beta} \dim(\alpha) \dim(\beta) \mathbb{E}_{\mathbf{b}} \left[\frac{\sum_{i,k} (A(\alpha) A'(\alpha))_{ik} \cdot (P B(\alpha)(\mathbf{b}))_{ki}}{\sum_{i',k'} \hat{h}(\beta)_{i'k'} \beta(\mathbf{b}^{-1})_{k'i'}} \right] \\ &= \sum_{\alpha, \beta} \frac{\dim(\alpha)}{\dim(\beta)} \sum_{i,k} (A(\alpha) A'(\alpha))_{ik} \sum_{i',k'} \hat{h}(\beta)_{i'k'} \mathbb{E}_{\mathbf{b}} \left[(P B(\alpha)(\mathbf{b}))_{ki} \cdot \beta(\mathbf{b}^{-1})_{k'i'} \right]. \end{aligned}$$

Consider the expectation:

$$\begin{aligned} & \mathbb{E}_{\mathbf{b}} \left[(PB(\alpha)(\mathbf{b}))_{ki} \cdot \beta(\mathbf{b}^{-1})_{k'i'} \right] \\ &= \mathbb{E}_{\mathbf{b}} \left[(P(\oplus_{m=1}^t n_m \beta_m(\mathbf{b})))_{ki} \cdot \beta(\mathbf{b}^{-1})_{k'i'} \right], \end{aligned}$$

where β_m s are the block along the diagonal of the block diagonal matrix $B(\alpha)(\mathbf{b})$ with multiplicity n_m .

For a block matrix $PB(\alpha)(\cdot)$, let $\mathcal{B}(PB(\alpha))$ be the indices (i, j) that belong to the blocks in the matrix. Let $\beta_{P,U,\alpha,k}^{\text{row}}$ ($\beta_{P,U,\alpha,i}^{\text{col}}$) denotes the irreducible representation of G^L present in the i^{th} column (k^{th} row) of the block matrix $PB(\alpha)(\cdot)$.⁷ For a fixed (α, k) following are the only scenarios when the expectation is nonzero:

- i must be such that (k, i) belongs to some block β_m in the block matrix $PB(\alpha)(\cdot)$, i.e., $(k, i) \in \mathcal{B}(PB(\alpha))$ (as otherwise $(P(\oplus_{m=1}^t n_m \beta_m(\mathbf{b})))_{ki} = 0$).
- β must be equal to $\beta_{P,U,\alpha,i}^{\text{col}}$ (Proposition 2.9). Furthermore, the entry of the matrix β_m given by (k, i) must be the “transpose” of (k', i') (Proposition 2.9). This also means that if we vary (i, k) inside a block β_m , then we get distinct (k', i') (i.e., transpose of (k, i) in that block) for which the expectation is non-zero (we will use this fact later). Thus, (α, i, k) uniquely determines (i', k') for which the expectation is non-zero. We denote this map by $(i', k') \leftarrow (\alpha, i, k)$.
- If both the above conditions are true, then the expectation is $\frac{1}{\dim(\beta_m)}$ (again, using Proposition 2.9).

Therefore, we have

$$\begin{aligned} & \Theta_{p,q,r}^e(\text{high}) \\ &= \sum_{\alpha} \dim(\alpha) \sum_{\substack{k, \\ i|(k,i) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,i,k)}} (A(\alpha)A'(\alpha))_{ik} \cdot \hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'} \\ &= \sum_{\alpha} \dim(\alpha) \sum_{\substack{k, \\ i|(k,i) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,i,k)}} \sum_j A(\alpha)_{ij} \cdot A'(\alpha)_{jk} \cdot \hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'} \\ &= \sum_{\alpha} \sum_{j,k} \dim(\alpha) A'(\alpha)_{jk} \cdot \sum_{i|(k,i) \in \mathcal{B}(PB(\alpha))} A(\alpha)_{ij} \cdot \hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'}. \end{aligned}$$

(rearranging)

We now apply the Cauchy–Schwartz inequality twice to simplify the expression.

$$\begin{aligned} |\Theta_{p,q,r}^e(\text{high})|^2 &\leq \left(\sum_{\alpha} \sum_{j,k} \dim(\alpha) |A'(\alpha)_{jk}|^2 \right) \cdot \\ &\quad \sum_{\alpha} \sum_{j,k} \dim(\alpha) \left| \sum_{\substack{i|(k,i) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,i,k)}} A(\alpha)_{ij} \cdot \hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'} \right|^2 \end{aligned}$$

Consider the first summation,

$$\begin{aligned} \sum_{\alpha} \sum_{j,k} \dim(\alpha) |A'(\alpha)_{jk}|^2 &= \sum_{\alpha} \dim(\alpha) \sum_{j,k} |A'(\alpha)_{jk}|^2 \\ &= \sum_{\alpha} \dim(\alpha) \|A'\|_{\text{HS}}^2 = \sum_{\alpha} \dim(\alpha) \|\hat{g}'(\alpha)'\|_{\text{HS}}^2, \end{aligned}$$

⁷Since we allow permutation of columns of a block diagonal matrices, $\beta_{P,U,\alpha,i}^{\text{row}}$ and $\beta_{P,U,\alpha,i}^{\text{col}}$ may be different, and hence the superscript.

where in the last step we use the fact that $\|A'(\alpha)\|_{\text{HS}} = \|\hat{g}'(\alpha)\|_{\text{HS}}$. Using Proposition 2.20, this is upper bounded by $\|\hat{g}'\|_2^2$ which is at most 1. Therefore,

$$|\Theta_{p,q,r}^e(\text{high})|^2 \leq \sum_{\alpha} \sum_{j,k} \dim(\alpha) \left| \sum_{\substack{i|(k,i) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,i,k)}} A(\alpha)_{ij} \cdot \hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'} \right|^2$$

By applying Cauchy–Schwartz inequality to the innermost summation, we get

$$\begin{aligned} & |\Theta_{p,q,r}^e(\text{high})|^2 \\ &\leq \sum_{\alpha} \sum_{j,k} \dim(\alpha) \left(\sum_{\substack{i|(k,i) \in \mathcal{B}(PB(\alpha))}} |A(\alpha)_{ij}|^2 \right) \left(\sum_{\substack{\tilde{i} |(k,\tilde{i}) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,\tilde{i},k)}} |\hat{h}(\beta_{P,U,\alpha,\tilde{i}}^{\text{col}})_{i'k'}|^2 \right). \end{aligned}$$

Now, let's look carefully at the summation. Fix the term $|A(\alpha)_{ij}|^2$. Note that this term appears for every k such that $(k, i) \in \mathcal{B}(PB(\alpha))$. On rearranging the summation,

$$\begin{aligned} & |\Theta_{p,q,r}^e(\text{high})|^2 \\ &\leq \sum_{\alpha} \sum_{i,j} \dim(\alpha) \sum_{k|(k,i) \in \mathcal{B}(PB(\alpha))} |A(\alpha)_{ij}|^2 \sum_{\substack{\tilde{i} |(k,\tilde{i}) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,\tilde{i},k)}} |\hat{h}(\beta_{P,U,\alpha,\tilde{i}}^{\text{col}})_{i'k'}|^2 \\ &\leq \sum_{\alpha} \sum_{i,j} \dim(\alpha) \cdot |A(\alpha)_{ij}|^2 \sum_{k|(k,i) \in \mathcal{B}(PB(\alpha))} \sum_{\substack{\tilde{i} |(k,\tilde{i}) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,\tilde{i},k)}} |\hat{h}(\beta_{P,U,\alpha,\tilde{i}}^{\text{col}})_{i'k'}|^2. \end{aligned}$$

Note that in the above expression, $\beta_{P,U,\alpha,\tilde{i}}^{\text{col}} = \beta_{P,U,\alpha,i}^{\text{col}}$ (because of the block matrix nature of $PB(\alpha)(\cdot)$). Therefore, we have

$$\begin{aligned} & |\Theta_{p,q,r}^e(\text{high})|^2 \\ &\leq \sum_{\alpha} \sum_{i,j} \dim(\alpha) \cdot |A(\alpha)_{ij}|^2 \sum_{\substack{k|(k,i) \in \mathcal{B}(PB(\alpha)), \\ \tilde{i} |(k,\tilde{i}) \in \mathcal{B}(PB(\alpha)), \\ (i',k') \leftarrow (\alpha,\tilde{i},k)}} |\hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'}|^2. \end{aligned}$$

As mentioned earlier, if we vary $(k, \tilde{i})$ inside a block β_m of

$$(P(\oplus_{m=1}^t n_m \beta_m(\cdot))),$$

then we get distinct (k', i') under the map $(i', k') \leftarrow (\alpha, \tilde{i}, k)$. The last sum is precisely varying inside one of the blocks (for a fixed (α, i, j))). Therefore,

$$\begin{aligned} & |\Theta_{p,q,r}^e(\text{high})|^2 \\ &\leq \sum_{\alpha} \sum_{i,j} \dim(\alpha) |A(\alpha)_{ij}|^2 \sum_{1 \leq i', k' \leq \dim(\beta_{P,U,\alpha,i}^{\text{col}})} |\hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})_{i'k'}|^2 \\ &= \sum_{\alpha} \dim(\alpha) \sum_{i,j} |U\hat{g}(\alpha)_{ij}|^2 \cdot \|\hat{h}(\beta_{P,U,\alpha,i}^{\text{col}})\|_{\text{HS}}^2. \end{aligned}$$

By taking a closer look at the expression above, it is not hard to see that there can be multiple scenarios when the expression is large. For instance, it might happen that some $\beta_{P,U,\alpha,i}^{\text{col}}$ s have small dimension and in this case we will not be able to get the advantage that we saw in the dictatorship test.

We avoid the above mentioned scenario by noting that when this happens then it must be the case that many distinct $\beta_{P,U,\alpha,i}^{\text{col}}$ s occur in the expression as we vary i . Thus, on average we can efficiently upper bound the expression, by using a careful choice of P given by the following claim, as long as $|\{\pi_{uv}\}_{\geq 2}(\alpha)|$ is large.

The following claim uses Lemma 2.26, for the proof see the full version of this paper [6].

CLAIM 9. Let $\varepsilon_0 \in (0, \frac{1}{2}]$. Suppose α and $e(u, v)$ are such that $|\pi_{uv}(\alpha)| \geq c$, where $c \geq 10|G|\log(\frac{1}{\varepsilon_0})$, then there exists a column-permutation matrix $\tilde{P}$ such that

$$\begin{aligned} & \sum_{i,j} |U\hat{g}(\alpha)_{ij}|^2 \cdot \|\hat{h}(\beta_{\tilde{P},U,\alpha,i}^{\text{col}})\|_{\text{HS}}^2 \\ & \leq \|\hat{g}(\alpha)\|_{\text{HS}}^2 \cdot \left(\varepsilon_0 + \sqrt{\max_{\beta \mid \dim(\beta) \geq c} \|\hat{h}(\beta)\|_{\text{HS}}^2} \right). \end{aligned}$$

Finishing the proof. We now proceed to upper bound the high terms. Let $\eta := C^{-d_0}$ where d_0 is the constant given in Theorem 2.2, $c := \frac{1}{\eta}$ and $\varepsilon_0 := \sqrt{\eta}$. Note that the condition on C and d_0 implies that $c \geq 10|G|\log(\frac{1}{\varepsilon_0})$. Next, we use the smoothness property of our Label Cover instance in order to apply Claim 9. With these settings of the parameters, the property says that for every $v \in \mathcal{V}$ and α such that $\dim_{\geq 2}(\alpha) > C$, for at least $(1 - \eta)$ fraction of the neighbors $u \sim v$ of v , $|\pi_{uv}(\alpha)| \geq c$. In what follows, we use the column-permutation matrix $\tilde{P} = P(e, \alpha)$, given by the Claim 9, for this setting of c and ε_0 .

$$\begin{aligned} & \mathbb{E}_{(u,v) \in E} \left[|\Theta_{p,q,r}^e(\text{high})|^2 \right] \\ & \leq \mathbb{E}_{(u,v) \in E} \left[\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) > C}} \dim(\alpha) \sum_{i,j} |U\hat{g}(\alpha)_{ij}|^2 \cdot \|\hat{h}(\beta_{\tilde{P},U,\alpha,i}^{\text{col}})\|_{\text{HS}}^2 \right] \\ & \leq \mathbb{E}_{(u,v) \in E} \left[\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) > C}} \dim(\alpha) \cdot \|\hat{g}(\alpha)\|_{\text{HS}}^2 \cdot \left(\varepsilon_0 + \sqrt{\max_{\beta \mid \dim(\beta) \geq c} \|\hat{h}(\beta)\|_{\text{HS}}^2} \right) \right] \\ & \quad + \eta. \\ & \leq \mathbb{E}_{(u,v) \in E} \left[\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) > C}} \dim(\alpha) \cdot \|\hat{g}(\alpha)\|_{\text{HS}}^2 \cdot \left(\sqrt{\max_{\beta \mid \dim(\beta) \geq c} \|\hat{h}(\beta)\|_{\text{HS}}^2} \right) \right] \\ & \quad + \eta + \varepsilon_0 \|g\|_2^2. \end{aligned}$$

Consider the summation,

$$\begin{aligned} & \sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) > C}} \dim(\alpha) \cdot \|\hat{g}(\alpha)\|_{\text{HS}}^2 \cdot \left(\sqrt{\max_{\beta \mid \dim(\beta) \geq c} \|\hat{h}(\beta)\|_{\text{HS}}^2} \right) \\ & \leq \sqrt{\frac{1}{c}} \left(\sum_{\substack{\alpha, \\ \dim_{\geq 2}(\alpha) > C}} \dim(\alpha) \cdot \|\hat{g}(\alpha)\|_{\text{HS}}^2 \right) \cdot \left(\sqrt{\sum_{\beta \mid \dim(\beta) \geq c} \dim(\beta) \cdot \|\hat{h}(\beta)\|_{\text{HS}}^2} \right) \\ & \leq \sqrt{\frac{1}{c}} \cdot \|g\|_2^2 \cdot \|h\|_2 \end{aligned}$$

Therefore using the fact that $\|g\|_2, \|h\|_2 \leq 1$, we get

$$\mathbb{E}_{(u,v) \in E} \left[|\Theta_{p,q,r}^e(\text{high})|^2 \right] \leq \sqrt{\frac{1}{c}} + \eta + \varepsilon_0 \leq 3\sqrt{\eta} \leq 3C^{-d_0/2} \leq \left(\frac{\delta}{2|G|^3} \right)^2,$$

where the last inequality follows from the choice of C . This implies,

$$\mathbb{E}_{(u,v) \in E} \left[|\Theta_{p,q,r}^e(\text{high})| \right] \leq \frac{\delta}{2|G|^3},$$

as required. $\square$

REFERENCES

- [1] Sanjeev Arora and Boaz Barak. 2009. *Computational complexity: a modern approach*. Cambridge University Press, USA.
- [2] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. 1998. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)* 45, 3 (1998), 501–555. <https://doi.org/10.1145/278298.278306>
- [3] Sanjeev Arora and Shmuel Safra. 1998. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM (JACM)* 45, 1 (1998), 70–122. <https://doi.org/10.1145/273865.273901>
- [4] Per Austrin and Elchanan Mossel. 2009. Approximation resistant predicates from pairwise independence. *Computational Complexity* 18, 2 (2009), 249–271. <https://doi.org/10.1007/s00037-009-0272-6>
- [5] László Babai, Nikolay Nikolov, and László Pyber. 2008. Product growth and mixing in finite groups. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, USA, 248–257. <https://dl.acm.org/doi/10.5555/1347082.1347110>
- [6] Amey Bhangale and Subhash Khot. 2020. Optimal Inapproximability of Satisfiable k -LIN over Non-Abelian Groups. arXiv:2009.02815 [cs.CC]
- [7] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. 1993. Self-testing/correcting with applications to numerical problems. *Journal of computer and system sciences* 47, 3 (1993), 549–595. [https://doi.org/10.1016/0022-0000\(93\)90044-W](https://doi.org/10.1016/0022-0000(93)90044-W)
- [8] Lars Engebretsen and Jonas Holmerin. 2005. Three-query PCPs with perfect completeness over non-Boolean domains. *Random Structures & Algorithms* 27, 1 (2005), 46–75. <https://doi.org/10.1002/rsa.20050>
- [9] Lars Engebretsen, Jonas Holmerin, and Alexander Russell. 2004. Inapproximability results for equations over finite groups. *Theoretical Computer Science* 312, 1 (2004), 17–45. [https://doi.org/10.1016/S0304-3975\(03\)00401-8](https://doi.org/10.1016/S0304-3975(03)00401-8)
- [10] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. 1996. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM (JACM)* 43, 2 (1996), 268–292. <https://doi.org/10.1145/226643.226652>
- [11] Mikael Goldmann and Alexander Russell. 2002. The complexity of solving equations over finite groups. *Information and Computation* 178, 1 (2002), 253–262. <https://doi.org/10.1006/inco.2002.3173>
- [12] W.T. Gowers. 2008. Quasirandom Groups. *Combinatorics, Probability and Computing* 17, 3 (2008), 363–387. <https://doi.org/10.1017/S0963548307008826>
- [13] Johan Håstad. 2001. Some optimal inapproximability results. *Journal of the ACM (JACM)* 48, 4 (2001), 798–859. <https://doi.org/10.1145/502090.502098>
- [14] Elchanan Mossel. 2010. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis* 19, 6 (2010), 1713–1756. <https://doi.org/10.1007/s00039-010-0047-x>
- [15] Ryan O'Donnell. 2014. *Analysis of boolean functions*. Cambridge University Press, USA.
- [16] Ran Raz. 1998. A parallel repetition theorem. *SIAM J. Comput.* 27, 3 (1998), 763–803. <https://doi.org/10.1137/S0097539795280895>
- [17] Linqing Tang. 2009. Conditional hardness of approximating satisfiable Max 3CSP-q. In *International Symposium on Algorithms and Computation*. Springer, Berlin, Heidelberg, 923–932. https://doi.org/10.1007/978-3-642-10631-6_93
- [18] Audrey Terras. 1999. *Fourier analysis on finite groups and applications*. Cambridge University Press, USA.