THEORY OF COMPUTING, Volume 13 (931), 2017, pp. 1–19
www.theoryofcomputing.org

On The Hardness of Approximating Balanced Homogeneous 3-Lin

Johan Håstad* Rajsekar Manokaran†

Received November 17, 2014; Revised May 18, 2017; Published September 18, 2017

Abstract: We consider systems of homogeneous linear equations modulo 2 with three variables in each equation and study balanced assignments as solutions to such equations. We prove that it is hard to distinguish systems where there is a balanced assignment that satisfies a fraction $1-\varepsilon$ of the equations from systems where the best balanced assignment satisfies a fraction $\frac{1}{2}+\varepsilon$ of the equations assuming that NP is not contained in quasipolynomial time. This improves on a similar result by Holmerin and Khot who relied on the assumption that NP is not contained in subexponential time. The key for the improvement is to replace long codes used by Holmerin and Khot by the low-degree long code.

1 Introduction

Many natural problems can be formulated in terms of *constraint satisfaction problems* (CSPs) where we have a large number of constraints on a large number of variables, but each constraint involves only a constant number of variables. To determine if there is an assignment that satisfies all the constraints is, except in a few special cases, NP-complete, and we turn to the question of satisfying as many constraints as possible. Clearly, finding the maximum fraction of constraints satisfiable is NP-hard and hence one

*KTH - Royal Institute of Technology. Work supported by ERC Grant 226-203
†IIT Madras. Work supported by ERC Grant 226-203

ACM Classification: F.2.2, F.1.3

AMS Classification: 68Q17, 68Q25

Key words and phrases: hardness of approximation, inapproximability, probabilistically checkable proofs, constraint satisfaction problems, bisection, approximation resistance

DOI: 10.4086/toc.2017.v013a931

turns to approximation algorithms. An algorithm is an α-approximation algorithm if its output is always within a factor α of the optimal value.

A huge effort has been spent on finding the best value of α achievable by a polynomial time algorithm for different problems. A key question here is whether an efficient algorithm beats the value of α obtained by simply picking an assignment at random. Problems where this is not possible are called *approximation resistant*, and a surprising number of problems have this property (see for instance [12, 8]). These results establish NP-hardness of the discussed problems and if we settle for Unique Games hardness, Austrin and Mossel [4] gave very general conditions for approximation resistance that apply to a vast majority of all constraint predicates [2]. Going even further, also in the case of Unique Games hardness, Khot et al. [16], gave necessary and sufficient conditions for "strong approximation resistance," a slightly stronger notion and thus we have, in broad terms, a fairly good understanding of the complexity of approximating CSPs. There are two special cases where the picture is less clear and more information is needed. One is that of *perfect completeness* (i. e., when we are guaranteed that there is an assignment that satisfies all conditions) and the other is when the problem also comes with *global constraints* and this is the case of interest in the current paper.

The simplest and possibly the most natural CSPs with global constraints are Max-Bisection and Min-Bisection where an assignment *bisecting* the vertices of a graph while maximizing or minimizing the cut edges is required. It is not difficult to see that Max-Bisection is at least as hard as the Max-Cut problem wherein any partition of the vertices is allowed. Thus, it is Unique Games hard to approximate this problem within the Goemans-Williamson constant (approximately 0.8786) [10, 15]. The algorithm by Austrin et al. [3] comes very close to obtaining this approximation ratio (it obtains a factor 0.8776) and it is a tantalizing open problem to understand whether the problem is strictly harder than Max-Cut.

When it comes to Min-Bisection, the situation is even more open. The best algorithm, by Räcke [18], achieves an approximation ratio that is logarithmic in the instance size. On the other hand, the best inapproximability, due to Khot [14], rules out a $1+\varepsilon$-approximation unless $\mathsf{NP} \subseteq \mathsf{DTIME}(\exp(n^{\varepsilon'}))$ where $\varepsilon' \to 0$ as $\varepsilon \to 0$.

In this paper we study a problem given by a system of linear equations modulo 2 with three variables in each equation where a solution is required to be *balanced*, i.e. give values 0 and 1 to (almost) the same number of variables. The basic variant of problem, called Max-3-LIN, without the condition of balance between the number of 0s and 1s is known to be NP-hard [12] to approximate within a factor of $\frac{1}{2}+\varepsilon$ for any $\varepsilon > 0$. As the good solution in the completeness case in the reduction of [12] is balanced the same hardness applies to the problem requiring a balanced assignment. In this paper, however, we study the homogeneous case, denoted by Max-Hom-Bal-3-LIN, where all the right hand sides are 0. In this situation allowing a general assignment makes the problem trivial as the identically 0 assignment satisfies all equations. The reason for studying this problem is at least two-fold. The fact that the balance condition is essential highlights this global condition which in the general case just happens to be satisfied. Secondly, having hardness for this simpler problem can give a more powerful starting point to use in further reductions.

Max-Hom-Bal-3-LIN, has previously been studied by Holmerin and Khot [13] who show that that it is hard to distinguish systems where a fraction $1-\varepsilon$ of the constraints can be satisfied, from systems where the best balanced assignment satisfies a fraction $1/2+\varepsilon$ of the constraints under a complexity assumption: a polynomial time distinguishing algorithm implies that $\mathsf{NP} \subseteq \mathsf{DTIME}(\exp(n^{\delta}))$ for all $\delta > 0$. We note

that an inapproximability of homogeneous linear equations, albeit with a stronger guarantee on the structure of the system, forms the core of the inapproximability of Min-Bisection [14].

1.1 Our Results and Techniques Used

We improve the complexity assumption in the above, proving the following statement.

Theorem 1.1 (Main). *For every $\varepsilon > 0$, a language $\mathsf{L} \in \mathsf{NP}$ can be reduced to* Max-Hom-Bal-3-LIN *so that a $w \in \mathsf{L}$ maps to an instance, $\mathcal{I}$, with value at least $1-\varepsilon$, while if $w \notin \mathsf{L}$, the instance $\mathcal{I}$ output has value at most $1/2+\varepsilon$. Further, the reduction is deterministic, and the running-time (and the size of $\mathcal{I}$) is at most* $\exp\bigl(\log(|w|)^{O(\log 1/\varepsilon)}\bigr)$.

As an immediate corollary, we see that even quasipolynomial time algorithms fail to approximate Max-Hom-Bal-3-LIN better than a random assignment.

Corollary 1.2. *For any $\varepsilon > 0$, there is no algorithm that runs in time* $\exp((\log n)^{O(1)})$ *and distinguishes* Max-Hom-Bal-3-LIN *instances with value at least $1-\varepsilon$ from instances with value at most $1/2+\varepsilon$ unless* $\mathsf{NP} \subseteq \mathsf{DTIME}\bigl(\exp\bigl((\log n)^{O(1)}\bigr)\bigr)$.

Our proof of inapproximability follows rather closely that of Holmerin and Khot [13] and is obtained by reducing graph 3-colorability to Max-Hom-Bal-3-LIN. The reduction uses a PCP where the verifier reads three bits from a proof and verifies that their exclusive-or is 0. The verifier is a composition of an outer verifier with an inner verifier. The former starts with a 3-colorability question and expects to find an encoding of a coloring as a low-degree polynomial. As noted by Holmerin and Khot [13], low-degree polynomials satisfy certain crucial properties that allow the inner verifier of Håstad [12] to output homogeneous linear equations.

Our main new ingredient is the use of the low-degree long code, introduced under the name "short code" by Barak et al. [6]. To adopt this to our inapproximability proof, we follow the approach by Dinur and Guruswami [9] and Guruswami et al. [11] who used it to prove lower bounds for some maximum constraint satisfaction problems and hypergraph colorability.

The technical problems encountered in the adaptation are not substantial as soon as one realizes that the line-point test for low-degree testing is a linear test in the bits in the proof. We believe that our result reinforces the proposition that the low-degree long code is versatile and may prove useful in many situations to get improved asymptotic bounds.

As already noted in [13], lower bounds for Max-Hom-Bal-3-LIN imply, through a simple gadget reduction, results for Max-Bisection. The result is that for any $\varepsilon > 0$ it is hard to approximate the latter problem within $\frac{15}{16}+\varepsilon$. The assumption needed is the same as for Theorem 1.2. This compares favorably to the factor $\frac{16}{17}+\varepsilon$ for which the problem is known to be NP-hard, but, of course, falls short Goemans-Williamson constant for which Unique Games hardness is known.

The next section introduces our notation and some standard facts used in the rest of the paper. Section 3 describes the outer verifier and its analysis. Section 4 describes the inner verifier and the proof of the main theorem. For completeness we give most details of the reduction to Max-Bisection in Section 5.

2 Preliminaries

The focus of this article is the problem Max-Hom-Bal-3-LIN, defined below.

Definition 2.1 (Max-Hom-Bal-3-LIN)**.** An instance $\mathcal{I}$ of Max-Hom-Bal-3-LIN is specified by a set of variables, $\mathcal{X}$, and a distribution $\mathcal{C}$ over 3-tuples of $\mathcal{X}$. An assignment, λ, to $\mathcal{I}$ is a map $\mathcal{X} \to \mathbb{F}_2$. Such an assignment is said to be *balanced* if the sets $\lambda^{-1}(0)$ and $\lambda^{-1}(1)$ differ in cardinality by at most 1. The value of an assignment is:

$$\mathsf{val}(\lambda;\mathcal{I}) \triangleq \mathop{\mathbf{Pr}}_{(x_1,x_2,x_3)\leftarrow\mathcal{C}}\{\lambda(x_1)+\lambda(x_2)+\lambda(x_3)=0\}.$$

Max-Hom-Bal-3-LIN asks us to identify a balanced assignment with maximum value.

We are interested in the size of an instance but as polynomial blow-up does not change the statement of the results it does not matter if we measure the size in the number of variables or the number of constraints. In view of this, for simplicity, we define the size to be the number of variables.

The Label Cover (LC) problem is a canonical starting point in many optimal inapproximability results. An intermediate step in our reduction produces LC instances with a special structure which is crucial in the subsequent parts of the proof. We call these instances of the *Certified LC* problem and the definition of the problem and its properties is given below. The key property that we need is that the constraints σ_u given below are low-degree polynomials and that the mappings π_{uv} are affine mappings. This is crucial for our ability to use the low-degree long code to code assignments.

Definition 2.2 (Certified LC)**.** An instance of Certified LC problem, henceforth CLC, is a tuple $\mathcal{L} = (U,V,E,\Pi,\Sigma)$ where:

- U and V are finite sets, E is a distribution over $U \times V$, and
- there exists $L,R \in \mathbb{Z}$ such that:

$$\Pi = \left\{\pi_{uv} : \mathbb{F}_2^L \to \mathbb{F}_2^R \mid (u,v) \in \mathsf{supp}(E)\right\}, \text{ and } \Sigma = \left\{\sigma_u : \mathbb{F}_2^L \to \mathbb{F}_2 \mid u \in U\right\}.$$

The integers L and R are called the left and right bit-length respectively while Π and Σ are called the projection and certification constraints. An assignment to the instance is a pair of functions (Λ_L,Λ_R) where $\Lambda_L : U \to \mathbb{F}_2^L$ satisfies $\sigma_u(\Lambda_L(u)) = 0$ for every $u \in U$ and $\Lambda_R : V \to \mathbb{F}_2^R$. The value of an assignment $\Lambda = (\Lambda_L,\Lambda_R)$ is defined as:

$$\mathsf{val}_{\mathcal{L}}(\Lambda) \triangleq \mathop{\mathbf{Pr}}_{(uv)\leftarrow E}\{\Lambda_R(v) = \pi_{uv}(\Lambda_L(u))\}.$$

We use the notation $\mathbf{Pr}_{v|_u}\{E\}$ to denote the probability of event E when choosing v as a random neighbor of u and we have a similar notation $\mathbf{E}_{v|_u}[X]$ for expectation.

We construct CLC instances that satisfy two crucial properties, mixing and smoothness, that are defined below.

Definition 2.3 (Smoothness)**.** A CLC instance is said to be ξ-smooth if for every $u \in U$ and every two distinct $a, b \in \mathbb{F}_2^L$,

$$\Pr_{v|_u}\{\pi_{uv}(a) = \pi_{uv}(b)\} \leq \xi.$$

Definition 2.4 (Mixing)**.** A CLC instance is said to be η-mixing if for every $\psi : V \to [0,1]$,

$$\mathop{\mathbf{E}}_{u}\left[\left|\mathop{\mathbf{E}}_{v|_u}[\psi(v)] - \mathop{\mathbf{E}}_{v}[\psi(v)]\right|\right] \leq \eta.$$

Our result invokes the analysis of two different low-degree testers and their set up and analysis are described next.

2.1 Testing Low-Degree Functions

For a field $\mathbb{F}$ and a positive integer m, $\mathbb{F}^m$ denotes the m-dimensional vector space over $\mathbb{F}$. We consider polynomials with coefficients in $\mathbb{F}$ and the degree of a polynomial $p : \mathbb{F}^m \to \mathbb{F}$ is the maximum over the degrees of the monomials where the degree of the monomial $\prod x_i^{d_i}$ is $\sum_i d_i$. The set of all polynomials of degree at most d is a vector space over $\mathbb{F}$ and is denoted by $\mathbb{P}_d(\mathbb{F}^m)$, or simply $\mathbb{P}_d$ if the domain of the polynomials is clear from the context. The *agreement* between two functions $f, g : \mathbb{F}^m \to \mathbb{F}$ is the fraction of inputs where they are equal and this quantity is denoted by $\alpha(f,g)$. The agreement of a function f with $\mathbb{P}_d$, $\alpha(f,\mathbb{P}_d)$, is the maximum of $\alpha(f,p)$ over $p \in \mathbb{P}_d$. The well known fact that low-degree polynomials have small pairwise agreement is essential for us and in coding theory terms this is a statement about the minimal distance of a Reed-Muller code. This is also known as the Schwarz-Zippel lemma and we use it in the following form.

Lemma 2.5. *Any two distinct $p, q \in \mathbb{P}_d(\mathbb{F}^m)$ satisfy $\alpha(p,q) \leq {}^{d}/_{|\mathbb{F}|}$.*

The first low-degree tester we describe is designed for large fields. It was developed in the context of the PCP theorem and uses geometric structures such as lines. We present it here as a CLC to suit our needs. A *line* passing through two distinct points $u, v \in \mathbb{F}^m$ has a parametric representation $\ell : \mathbb{F} \to \mathbb{F}^m$ given by $\ell(x) = u + x \cdot (v - u)$. Similarly, for a positive integer k, and $k+1$ points $u_0, \ldots, u_k$, we have the function $p : (x_1, \ldots, x_k) = u_0 + \sum_i x_i(u_i - u_0)$ whose image, $\{p(x) \mid x_i \in \mathbb{F}\}$ is a *flat*, denoted by $\mathbb{S}(u_0, \ldots, u_k)$. When the elements $(u_i - u_0)$ are all linearly independent, the set is a *k-flat*. Note that different choices of u_i might yield the same flat, and we fix an arbitrary canonical parametric representation of each flat. The space of all k-flats is denoted by $\mathbb{S}_k$, or $\mathbb{S}_k(\mathbb{F}^m)$ when the vector space needs to be emphasized.

Fix a positive integer t and let $\mathbb{F}$ be the field of order $q = 2^t$. The *line-point* test is the CLC with $V = \mathbb{F}^m$, $U = \mathbb{S}_1(\mathbb{F}^m)$, and $\sigma_u \equiv 0$ for every $u \in U$. The assignments to V are expected to be the evaluation of a polynomial $p \in \mathbb{P}_d$ and thus we set $R = t$. The restriction of p of degree d to a line ℓ is $p \circ \ell$ and is a univariate polynomial of degree at most d. The assignment to the line $\ell \in U$ is to be the representation of this polynomial. We set $L = t(d+1)$ and treat an assignment as a polynomial $p_\ell : \mathbb{F} \to \mathbb{F}$. For a $x \in \mathbb{F}$, $\ell(x) \in \mathbb{F}^m = V$ the constraint $\pi_{\ell,\ell(x)}$ maps p_ℓ to $p_\ell(x)$. Finally, the distribution E picks $\ell \leftarrow U$, and $x \leftarrow \mathbb{F}$ uniformly at random and produces $(\ell, \ell(x))$. We denote the CLC by $\mathcal{L}^{\mathrm{LP}}_{m,d}(\mathbb{F})$. The following analysis of this test is due to Arora and Sudan [1].

Theorem 2.6 (Theorem 16, [1]; reformulation[1]). *There exists constants $c > 0$ and $\gamma_0 > 0$ such that, if $q > d^c$, $\gamma < \gamma_0$ the following holds. For every assignment $\Lambda_R : V \to \mathbb{F}$ to $\mathcal{L}^{\mathrm{LP}}_{m,d}(\mathbb{F})$, if for some Λ_L: (Λ_L, Λ_R) has value at least $1-\gamma$, then there is a $p \in \mathbb{P}_d$ such that $\alpha(\Lambda_R, p) \geq 7/12$.*

Points on a line are pairwise independent and hence the line-point test yields a mixing CLC.

Lemma 2.7. *The instance $\mathcal{L}^{\mathrm{LP}}_{m,d}(\mathbb{F})$ is $\sqrt{1/q}$-mixing.*

Proof. Fix a function $\psi : V \to [0,1]$. Now,

$$\mathop{\mathbf{E}}_{\ell \leftarrow \mathbb{S}_1}\left[\mathop{\mathbf{E}}_{t \leftarrow \mathbb{F}}[\psi(\ell(t))]\right] = \mathop{\mathbf{E}}_{u \leftarrow \mathbb{F}^m}[\psi(u)],$$

$$\text{and, } \mathop{\mathbf{E}}_{\ell \leftarrow \mathbb{S}_1}\left[\left(\mathop{\mathbf{E}}_{t \leftarrow \mathbb{F}} \psi(\ell(t))\right)^2\right] = \mathop{\mathbf{E}}_{\ell, t_1, t_2}[\psi(\ell(t_1))\psi(\ell(t_2))].$$

For fixed $t_1 \neq t_2$, conditioned on $x = \ell(t_1)$, $\ell(t_2)$ on a random line ℓ takes every value in V other than $\ell(t_1)$ with equal probability and thus in this case

$$\mathop{\mathbf{E}}_{\ell}[\psi(\ell(t_1))\psi(\ell(t_2))] = \psi(x)\frac{q^m \mathbf{E}[\psi] - \psi(x)}{q^m}.$$

Since $t_1 = t_2$ with probability $1/q$,

$$\mathop{\mathbf{E}}_{\ell \leftarrow \mathbb{S}_1}\left[\left(\mathop{\mathbf{E}}_{t \leftarrow \mathbb{F}} \psi(\ell(t))\right)^2\right] \leq \frac{1}{q} + \frac{q-1}{q}(\mathbf{E}[\psi])^2 \leq \frac{1}{q} + (\mathbf{E}[\psi])^2.$$

An application of Jensen's inequality yields the result. □

2.2 Fourier Analysis over $\mathbb{P}_d(\mathbb{F}_2^L)$

The second low-degree testing result is used in analyzing maps from $\mathbb{P}_d(\mathbb{F}_2^L)$ to $\mathbb{R}$. We review standard facts about the Fourier transform of such maps towards stating the testing result in this context. In what follows, $\mathbb{F} = \mathbb{F}_2$ and the vector space is $\mathbb{F}^L$ unless stated otherwise. Every function $f : \mathbb{F}^L \to \mathbb{F}$ can be written as a multilinear polynomial of degree at most L and thus the space of maps, $\mathbb{F}^L \to \mathbb{F}$, is equal to $\mathbb{P}_L$. For two maps f, g, we have the natural inner product: $\langle f, g\rangle \triangleq \sum_x f(x)g(x)$ where the operations are over $\mathbb{F}$. As is well known, the dual code of $\mathbb{P}_d$ under this inner product is $\mathbb{P}_{L-d-1}$. For a $\beta \in \mathbb{P}_L$, define the character map $\chi_\beta : \mathbb{P}_L \to \mathbb{R}$ by:

$$\chi_\beta(g) \triangleq (-1)^{\langle \beta, g\rangle}.$$

If $\beta_1 = \beta_2 + p$ where $p \in \mathbb{P}_{L-d-1}$ then, as $\mathbb{P}_{L-d-1}$ is the dual of $\mathbb{P}_d$, $\chi_{\beta_1}(g) = \chi_{\beta_2}(g)$ for any $g \in \mathbb{P}_d$. By standard terminology β_1 and β_2 are in this case said to belong to the same *coset* of $\mathbb{P}_{L-d-1}$.

The Hamming weight of β, denoted by $\mathsf{wt}(\beta)$, is the number of x such that $\beta(x) = 1$. From each coset of $\mathbb{P}_{L-d-1}$ in $\mathbb{P}_L$, we pick a β of least Hamming weight (breaking ties arbitrarily) and denote the

[1] The number $7/12$ can be replaced by any number between $1/2$ and $2/3$. The constant γ_0 is such that setting $\rho \geq (1-\gamma_0)$ in Theorem 16 satisfies $2\rho/3 \geq 7/12$.

set of these representatives by $\mathbb{X}_d$. A standard fact about the Fourier transform states that a function $\lambda : \mathbb{P}_d \to \mathbb{R}$ can be written as:

$$\lambda(g) = \sum_{\beta \in \mathbb{X}_d} \hat{\lambda}(\beta)\chi_\beta(g), \tag{2.1}$$

where the *Fourier coefficients* $\hat{\lambda}(\beta)$ are real numbers. The characters form an orthonormal basis as for every $\beta_1, \beta_2 \in \mathbb{X}_d$,

$$\mathop{\mathbf{E}}_{g \leftarrow \mathbb{P}_d}\left[\chi_{\beta_1}(g)\chi_{\beta_2}(g)\right] = \begin{cases} 1 & \text{if } \beta_1 = \beta_2 \\ 0 & \text{otherwise.} \end{cases} \tag{2.2}$$

An *affine form* in L variables is the polynomial $a_0 + \sum a_i x_i$ where each a_i is in $\mathbb{F}_2$. A list of affine forms $\pi^{(1)}, \ldots, \pi^{(R)} : \mathbb{F}^L \to \mathbb{F}$, can be viewed as an affine map $\pi : \mathbb{F}^L \to \mathbb{F}^R$. In this situation, for $f \in \mathbb{P}_d(\mathbb{F}^R)$, we have that also $f \circ \pi = f(\pi(x))$ belong to $\mathbb{P}_d(\mathbb{F}^L)$. Let $\pi^{-1}(y)$ be the set of all x such that $\pi(x) = y$, and for a $\beta : \mathbb{F}^L \to \mathbb{F}$ define the projected map $\pi_2(\beta) : \mathbb{F}^R \to \mathbb{F}$ as

$$\pi_2(\beta)(y) \triangleq \sum_{x \in \pi^{-1}(y)} \beta(x). \tag{2.3}$$

It is not difficult to see that the identity $\chi_\beta(f \circ \pi) = \chi_{\pi_2(\beta)}(f)$ holds.

Let $\mathbb{L}_d$ denote the space of all polynomials which are the product of exactly d linearly independent affine forms. The result of Bhattacharyya et al. [7] on testing functions in $\mathbb{P}_{L-d-1}(\mathbb{F}_2^L)$ shows the following.

Theorem 2.8 (Theorem 1, [7]; stated here as in Proposition 14, [9])**.** *There is a $\rho_0 < 1$ such that for every $\beta \in \mathbb{X}_d$,*

$$\mathop{\mathbf{E}}_{\zeta \leftarrow \mathbb{L}_d}\left[\chi_\beta(\zeta)\right] \le \max\left\{1 - \frac{\mathsf{wt}(\beta)}{2^d}, \rho_0\right\}.$$

3 Outer Verifier

In this section we show that the values of CLC instances satisfying all the necessary structural properties are still hard to approximate. The proof is via a reduction from the 3-coloring problem on graphs and is essentially due to Holmerin and Khot [13]. Our statement (see Theorem 3.2) allows for a more flexible setting of the parameters controlling the size of the reduction and proves a few additional properties that are relevant to its use in section 4.

A graph is said to be θ-far from being 3-colorable if every 3-coloring of the vertices has at least a θ-fraction of the edges between vertices of the same color. The work of Petrank [17] (see also [5]) shows the following inapproximability of the 3-coloring problem.

Theorem 3.1 (Theorem 3.3, [17])**.** *There is a number $\theta > 0$ such that 3-colorable graphs are* NP*-hard to distinguish from graphs that are θ-far from being 3-colorable.*

In what follows, we reserve the symbol n for the number of vertices in the graph we reduce from. A parameter, say m, that depends on the size of the graph is a function $\mathbb{Z}^+ \to \mathbb{Z}^+$ of n, but we omit the argument writing m instead of $m(n)$.

Theorem 3.2. *There exist constants c, d_0 such that, for parameters m, d, and t satisfying*

$$\binom{m}{d} \geq n;\ d \geq d_0;\ \text{and } q = 2^t \geq d^c, \tag{3.1}$$

and for any $\delta > 0$, there is a reduction from a graph G on n vertices to a CLC $\mathcal{L} = (U, V, E, \Pi, \Sigma)$ such that the statements below are true.

1. *If G is 3-colorable, then $\mathsf{val}(\mathcal{L}) = 1$.*
2. *If G is θ-far from being 3-colorable, then $\mathsf{val}(\mathcal{L}) \leq \delta$ (here θ is as in Theorem 3.1).*
3. *The instance $\mathcal{L}$ is $\mathrm{O}\big(\log(1/\delta)/\sqrt{q}\big)$-mixing, d/q-smooth and further, the marginal distribution of $v \in V$ obtained from E is uniform.*
4. *The left and right bit-lengths (L and R) are at most $\mathrm{O}\big(\log(1/\delta)td^3\big)$.*
5. *Each projection constraint $\pi_{uv} : \mathbb{F}_2^L \to \mathbb{F}_2^R$ in Π is a list of R affine maps $\mathbb{F}_2^L \to \mathbb{F}_2$.*
6. *Each certification constraint σ_u in Σ is decomposable into a list of polynomials $p_1, \ldots, p_k$ each belonging to $\mathbb{P}_2(\mathbb{F}_2^L)$ such that $\sigma_u(x) = 0$ if and only if $p_i(x) = 0$ for every i.*
7. *The reduction is deterministic and requires at most $(q^m n)^{\mathrm{O}(\log(1/\delta))}$ steps.*

The rest of this section proves this theorem. The reduction first embeds the coloring problem in the line-point CLC. We set up some notation before presenting the reduction. For parameters m, d, and t as in Theorem 3.2, $\mathbb{F}$ denotes the field of order 2^t. We need a representation of each element of $\mathbb{F}$ by a vector of t bits and use a natural representation such that field-addition becomes exclusive-or of vectors and multiplication by a fixed field element is a linear transformation. The zero of the field is represented by 0^t and we assume that the indentity is represented by the unit vector e_t with a one in the last component. We denote these two elemenst by $\mathbf{0}$ and $\mathbf{1}$, respectively. Finally we let $\mathbf{2}$ denote the field element represented by e_{t-1}. These three elements are used as colors and a 3-coloring of G is viewed as $\phi : G \to \{\mathbf{0}, \mathbf{1}, \mathbf{2}\}$.

A set $s \subseteq [m]$ is also viewed as a point $\mathbf{s} \in \mathbb{F}^m$ where $\mathbf{s}_k = \mathbf{1}$ if $k \in s$ and $\mathbf{s}_k = \mathbf{0}$ otherwise where we typeset the symbol in bold face to emphasize the latter view. Let $(\mathbb{S}_1(\mathbb{F}^m), \mathbb{F}^m, E_0, \Pi_0, \Sigma_0)$ denote the line-point CLC, $\mathcal{L}^{\mathrm{LP}}_{m,d}(\mathbb{F})$. Each element in U of the CLC we construct extends a line $\ell = \mathbb{S}(u_0, u_1) \in \mathbb{S}_1$ into a flat $\mathbb{S}(u_0, u_1, \mathbf{s_i}, \mathbf{s_j})$ for some choice of $\mathbf{s_i}, \mathbf{s_j} \in \mathbb{F}^m$. An assignment to this flat is a polynomial in $\mathbb{P}_d(\mathbb{F}^3)$ that describes the restriction of a polynomial in $\mathbb{P}_d(\mathbb{F}^m)$ along the parametric representation of the flat. Note that different choices of u_0 and u_1 may lead to the same flat and thus we fix some canonical parametric representation of the flat $\mathbb{S}(u_0, u_1, \mathbf{s_i}, \mathbf{s_j})$. In particular, given an assignment $p \in \mathbb{P}_d(\mathbb{F}^3)$, $p(\mathbf{0}, \mathbf{1}, \mathbf{0})$ and $p(\mathbf{0}, \mathbf{0}, \mathbf{1})$ are the values at $\mathbf{s_i}$ and $\mathbf{s_j}$ respectively.

Definition 3.3 (Reduction to CLC)**.**

Input: A graph $G = ([n], E_G)$; and parameters m, d and t.

Output: A CLC instance $\mathcal{L} = (U, V, E, \Pi, \Sigma)$.

1. Pick a collection $\mathcal{F} = \{s_1, \ldots, s_n\}$ of n distinct elements of $\binom{[m]}{d}$.

2. Set $V = \mathbb{F}^m$, and

$$U = \bigcup_{(i,j)\in E_G} \left\{ \mathbb{S}(u_0, u_1, \mathbf{s_i}, \mathbf{s_j}) \mid \mathbb{S}(u_0, u_1) \in \mathbb{S}_1(\mathbb{F}^m) \text{ and } \mathbb{S}(u_0, u_1, \mathbf{s_i}, \mathbf{s_j}) \text{ a 3-flat} \right\}.$$

   An assignment to $v \in V$ is an element of $\mathbb{F}$; thus $R = t$. Assignments to U are from $\mathbb{P}_d(\mathbb{F}^3)$ and thus $L = \binom{d+3}{3}t$.
3. Set E to be the distribution that
   (a) samples an edge $(i, j) \in E_G$ uniformly at random;
   (b) samples $(\ell = \mathbb{S}(u_0, u_1), v)$ from E_0;
   (c) picks a random 3-flat u containing $\tilde{u} = \mathbb{S}(u_0, u_1, \mathbf{s_i}, \mathbf{s_j}), v)$; and
   (d) produces (u, v).
4. Let (x_1, x_2, x_3) map to the point v in the canonical representation of the flat u picked above. Set $\pi_{uv} = p \to p(x_1, x_2, x_3)$, for $p \in \mathbb{P}_d(\mathbb{F}^3)$.
5. Finally, σ_u maps a polynomial p to zero if and only if $p(\mathbf{0}, \mathbf{1}, \mathbf{0})$ and $p(\mathbf{0}, \mathbf{0}, \mathbf{1})$ are valid and distinct colors.

Note that the collection $\mathcal{F}$ exists as long as eq. (3.1) is satisfied. We fix a CLC $\mathcal{L}$ obtained from a graph G and proceed with the analysis.

Lemma 3.4. *If G is 3-colorable, then* $\mathsf{val}(\mathcal{L}) = 1$.

Proof. For a set $s \subseteq [m]$, we write x_s for the monomial $\prod_{j\in s} x_j$. A 3-coloring of the graph, $\phi : [n] \to \{\mathbf{0}, \mathbf{1}, \mathbf{2}\}$ corresponds to the polynomial $p = \sum_{i\in[n]} x_{s_i}\phi(i)$ where s_i are from the collection $\mathcal{F}$. The evaluation of this polynomial is the intended assignment to the above instance. If G is 3-colorable, then setting ϕ to be a valid coloring assures that the restriction of p to the flat u, $p \circ u$, satisfies the constraint σ_u. Further, since each s_i has exactly d elements, p is degree-d. Clearly the assignment $v \to p(v)$ for $v \in V$, and $u \to p \circ u$ for the flats $u \in U$ satisfies the projection constraints, π_{uv}. □

Thinking of the set s_i as its characteristic vector we note that $p(s_i) = \phi(i)$. This is mostly a curiosity as values at specific points is not a "robust" property and it is much more interesting what values p takes at random points.

As a complement to Lemma 3.4 giving completeness, we have the following soundness statement which is the analogue of item 2 in Theorem 3.2.

Lemma 3.5. *$\exists \gamma_0 > 0, d_0$ so that $\forall \gamma < \gamma_0, d > d_0$, and G that is γ-far from being 3-colorable,* $\mathsf{val}(\mathcal{L}) \leq 1 - \gamma^2$.

Proof. Suppose $\Lambda = (\Lambda_L, \Lambda_R)$ is an assignment whose value is larger than $1 - \gamma^2$. Then, a fraction $1 - \gamma$ of edges $(i, j) \in E_G$ are *valuable* which we define to be that the probability, conditioned on picking one of these edges, of satisfying the projection is at least $1 - \gamma$.

The distribution E conditioned on an edge $(i, j) \in E_G$, is exactly E_0 of $\mathcal{L}^{\mathrm{LP}}_{m,d}(\mathbb{F})$. Conditioning on a valuable edge (i, j) and applying Theorem 2.6, we see that Λ_R is $7/12$ close to a $p \in \mathbb{P}_d$. Further, using Theorem 2.5, any other $p' \in \mathbb{P}_d$ agrees with p on at most a d/q-fraction and thus agrees with Λ_R on at

most a fraction $5/12+d/q$. Thus, for a sufficiently large d_0 (and hence d and $q \geq d^c$), the choice of p with agreement $7/12$ with Λ_R is unique. We claim, and prove below, that whenever (i,j) is valuable, $p(\mathbf{s_i})$ and $p(\mathbf{s_j})$ are valid and distinct colors. This implies that $\{p(\mathbf{s_i})\}_i$ contradicts G being γ-far from being 3-colorable. To see that $p(\mathbf{s_i})$ and $p(\mathbf{s_j})$ are valid and distinct suppose for contradiction that they are not.

As $\Lambda_L(u)(\mathbf{s_i})$ and $\Lambda_L(u)(\mathbf{s_j})$ are valid and distinct colors, p and $\Lambda_L(u)$ on u and are two distinct degree-d polynomials and hence they agree on a fraction at most d/q of the points.

Let ℓ be the line through $\mathbf{s_i}$ and $\mathbf{s_j}$ and consider the distribution of u and v. It is easy to see that the probability that v is on ℓ is $O(q^{1-m})$ and otherwise it is uniformly distributed on u outside ℓ. As each point in the space outside ℓ is equally likely to belong to u we have the the expected agreement of p and Λ_R on u outside ℓ is at least $7/12 - O(q^{1-m})$. As ℓ is a fraction $1/q^2$ of u it follows that the probability that $\Lambda_L(u)$ agrees with Λ_R at v is at most $5/12 + O(q^{1-m}) + d/q + 1/q^2$. This contradicts the assumption that (i,j) is valuable and thus we can conclude that $p(\mathbf{s_i})$ and $p(\mathbf{s_j})$ are valid and distinct colors and the proof is complete. □

3.1 Gap Amplification

To amplify the guarantee from Theorem 3.5, we apply the parallel repetition theorem. We let $U^{\oplus k}$ be the k-fold direct product of U and use the notations $V^{\oplus k}$ and $E^{\oplus k}$ similarly. Given a CLC $\mathcal{L}$, and a positive integer k, parallel repetition produces an instance $\mathcal{L}^{\oplus k} = (U_k, V_k, E_k, \Sigma_k, \Pi_k)$ obtained by setting:

- $U_k = U^{\oplus k}$; $V_k = V^{\oplus k}$; E as the k-wise product distribution $E^{\oplus k}$;
- $\pi_{(u_1,\ldots,u_k)(v_1,\ldots,v_k)} = (\pi_{u_1v_1},\ldots,\pi_{u_kv_k})$; and $\sigma_{u_1,\ldots,u_k}(p_1,\ldots,p_k) = 0$ if $\forall i\ \sigma_{u_i}(p_i) = 0$.

The CLC problem, akin to Label Cover, falls in the category of 2-prover 1-round projection games and thus we have the following guarantee on $\mathsf{val}(\mathcal{L}^{\oplus k})$ from the work of Rao [19].

Theorem 3.6 (Theorem 4, [19]; restated for CLCs)**.** *There is an $\alpha > 0$ such that if a CLC instance $\mathcal{L}$ has value at most $1-\gamma$, then $\forall k \in \mathbb{Z}^+$, we have $\mathsf{val}(\mathcal{L}^{\oplus k}) \leq (1-\gamma/2)^{\alpha\gamma k}$. In particular, if $k \geq \frac{2\log(1/\delta)}{\alpha\gamma^2}$, then $\mathsf{val}(\mathcal{L}^{\oplus k}) \leq \delta$.*

Further, as shown by Holmerin and Khot [13], the smoothness is preserved while mixing deteriorates linearly with k.

Theorem 3.7 (Theorem B.1, [13])**.** *If $\mathcal{L}$ is ξ-smooth, then so is $\mathcal{L}^{\oplus k}$ for any positive integer k.*

Theorem 3.8 (Theorem B.5, [13])**.** *If $\mathcal{L}$ is η-mixing, then $\mathcal{L}^{\oplus k}$ is $k\eta$-mixing for any $k \in \mathbb{Z}^+$.*

We now finish the proof of the main theorem of this section.

Proof of Theorem 3.2. Set d_0 as in Theorem 3.5 and fix parameters m, d, and t satisfying the premises and fix a $\delta > 0$. Set $\gamma = \max\{\gamma_0, \theta\}$ (γ_0 as in Theorem 3.5, and θ as in Theorem 3.1) and set $k = {}^{2\log(1/\delta)}/_{\alpha\gamma^4}$ with α as in Theorem 3.6. The reduction transforms G to $\mathcal{L}$ by running Theorem 3.3 and produces $\mathcal{L}^{\oplus k}$. We show that $\mathcal{L}^{\oplus k}$ satisfies all the properties stated.

Indeed, if G is 3-colorable, then Theorem 3.4 gives an assignment (Λ_L, Λ_R) that satisfies the constraints in $\mathcal{L}$. The assignment $(\Lambda_L^{(k)}, \Lambda_R^{(k)})$ where $\Lambda_L^{(k)}(u_1, \ldots, u_k) = (\Lambda_L(u_1), \ldots, \Lambda_L(u_k))$ and similarly, $\Lambda_R^{(k)}(v_1, \ldots, v_k) = (\Lambda_R(v_1), \ldots, \Lambda_R(v_k))$ satisfies all the constraints in $\mathcal{L}^{\oplus k}$. Thus we have item 1.

To establish item 2 we note that, by Theorem 3.5, if G is θ-far from 3-colorable then $\mathsf{val}(\mathcal{L}) \leq 1 - \gamma^2$ and hence by Theorem 3.6 and the choice of k we get the desired conclusion, and we proceed to consider item 3.

The projection constraints of $\mathcal{L}$ are evaluations of the polynomials in $\mathbb{P}_d(\mathbb{F}^3)$ at specific points. Thus, smoothness of $\mathcal{L}$ follows immediately from Theorem 2.5, and applying Theorem 3.7 proves it for $\mathcal{L}^{\oplus k}$. For each choice of an edge (i, j) in the distribution E, the tests are a copy of $\mathcal{L}^{\mathrm{LP}}_{m,d}(\mathbb{F})$ over the same set $V = V_0$, Thus, mixing follows from Theorem 2.7 and Theorem 3.8. The claim on the marginal distribution follows from uniformity of points on a random line. This proves item 3.

An assignment to an element in $V^{\oplus k}$ is a list of k field elements and hence $R = kt$. Similarly, an element in $\mathbb{P}_d(\mathbb{F}^3)$ is identified by specifying $\binom{d+3}{3}$ elements and hence $L = \binom{d+3}{3}tk$ bits encode an assignment to $U^{\oplus k}$. Thus, we have item 4.

Recall that the elements of $\mathbb{F}$ are represented by vectors in $\mathbb{F}_2^t$ where field additions correspond to vector additions while multiplication by an element is multiplication by an invertible matrix. Thus, if the polynomial is written as a vector in $\mathbb{F}_2^L$, the evaluation at a fixed point is a multiplication by a matrix in $\mathbb{F}_2^{R \times L}$. This proves item 5.

Given a $p \in \mathbb{P}_d(\mathbb{F}^3)$, verifying the constraint σ_u for $u \in U$ requires verifying that $p(\mathbf{0}, \mathbf{1}, \mathbf{0})$ and $p(\mathbf{0}, \mathbf{0}, \mathbf{1})$ are valid and distinct colors. A color $y \in \mathbb{F}_2^t$ is valid if all but the last 2 bits are 0 and if the last two bits are not both 1. The former constraints are linear and the latter degree-2 constraints and thus may be encoded as a list of $t - 1$ polynomial equations in the coefficients of p, all of degree at most two. We have that, $x, y \in \mathbb{F}_2^t$, both coding correct colors, are distinct iff $(1 + x_t + y_t) \cdot (1 + x_{t-1} + y_{t-1}) = 0$. Each of these bits is an affine form in the bits representing the polynomial p and hence we have a list of $2(t-1)+1$ polynomials ($2(t-2)$ affine polynomials and 3 quadratic polynomials) that encode the coloring constraints. Concatenating such lists obtained for different u_i yields the necessary list for a $u \in U^{\oplus k}$ hence proving item 6.

Finally, the size of the set U is $\mathrm{O}\left(n^2 \cdot q^{2m}\right)$ and hence the size of $\mathcal{L}^{\oplus k}$ is at most $(q^{2m} n^2)^{\mathrm{O}(k)}$. This proves item 7, concluding the proof. □

4 Inner Verifier

In this section, we prove the main theorem by reducing from CLC to Max-Hom-Bal-3-LIN. The reduction replaces each element of U and V with the low-degree long code (a.k.a the short code) introduced in the work of Barak et al. [6] and uses folding ideas from the work of Dinur and Guruswami [9]. We now describe these concepts, and follow it up with our reduction and the analysis.

We have a parameter s, to be determined later, governing our construction. The low-degree long code encoding of an element $a \in \mathbb{F}_2^L$ is the map $\mathsf{SC}_s^{(a)} : \mathbb{P}_s(\mathbb{F}_2^L) \to \mathbb{F}_2$ defined as:

$$\mathsf{SC}_s^{(a)}(g) \triangleq g(a) \tag{4.1}$$

For each $u \in U$, the constraint σ_u can be decomposed to a list of polynomials $p_1, \ldots, p_k \in \mathbb{P}_2(\mathbb{F}_2^L)$ as

in item 6 of Theorem 3.2. Define the subspaces[2]

$$\mathbb{I}_u = \left\{ p \in \mathbb{P}_s(\mathbb{F}_2^L) \mid p = \sum p_i q_i; \quad q_i \in \mathbb{P}_{s-2}(\mathbb{F}_2^L) \right\} \text{ for each } u \in U.$$

Let $P^{(u)}$ denote the cosets of $\mathbb{I}_u$ in $\mathbb{P}_s(\mathbb{F}_2^L)$. Note that if $a \in \mathbb{F}_2^L$ satisfies $\sigma_u(a) = 0$, then $\mathsf{SC}_s^{(a)}(f) = \mathsf{SC}_s^{(a)}(f')$ if f and f' belong to the same coset. The reduction uses this to replace each vertex u by a table indexed by $P^{(u)}$ instead of a complete low-degree code. This is viewed as the complete code folded so that elements belonging to a fixed coset are implicitly mapped to a single element in $P^{(u)}$.

Definition 4.1 (Reduction to Max-Hom-Bal-3-LIN)**.**

Input: A CLC $\mathcal{L} = (U, V, E, \Pi, \Sigma)$; and a parameter s, where the functions in Π are affine and the constraints in Σ are polynomials of degree at most 2.

Output: A Max-Hom-Bal-3-LIN instance $\mathcal{I} = (\mathcal{X}, \mathcal{C})$ where,

$$\mathcal{X} = \left\{ (u, g) \mid u \in U, g \in P^{(u)} \right\} \cup \left(V \times \mathbb{P}_s(\mathbb{F}_2^R) \right);$$

and $\mathcal{C}$ is the distribution that:

1. samples a constraint $(u, v) \leftarrow E$;
2. samples $g \leftarrow \mathbb{P}_s(\mathbb{F}_2^L)$ and $h \leftarrow \mathbb{P}_s(\mathbb{F}_2^R)$ at random;
3. samples $w = 2^{3s/4}$ independent $\zeta_1, \ldots, \zeta_w \leftarrow \mathbb{L}_s(\mathbb{F}_2^L)$; sets $\zeta = \zeta_1 + \cdots + \zeta_w$; and
4. produces the tuple $((u, g), (v, h), (u, g + \zeta + h \circ \pi_{uv}))$.

The reduction is well-defined for inputs $\mathcal{L}$ guaranteed by Theorem 3.2 as item 5 ensures that $h \circ \pi$ is degree-s if h is degree-s. We analyze the value of the instance in relation to the value of $\mathcal{L}$.

Lemma 4.2 (Completeness)**.** *If $\mathcal{L}$ has an assignment (Λ_L, Λ_R) whose value is 1, then there is a balanced assignment to $\mathcal{I}$ with value at least $1 - 2^{-s/4}$.*

Proof. Consider the assignment: $(u, g) \to \mathsf{SC}_s^{(\Lambda_L(u))}(g)$; $(v, h) \to \mathsf{SC}_s^{(\Lambda_R(v))}(h)$. Since Λ satisfies the certification constraints, and hence $p(\Lambda_L(u)) = 0$ for any $p \in \mathbb{I}_u$, the assignment is well-defined. Further, (u, g) and $(u, 1 + g)$ as well as (v, h) and $(v, 1 + h)$ take distinct values and hence the assignment is balanced. Under this assignment, $(u, g + \zeta + h \circ \pi_{uv})$ gets the value

$$(g + \zeta + h \circ \pi_{uv})(\Lambda_L(u)) = g(\Lambda_L(u)) + \zeta(\Lambda_L(u)) + h \circ \pi_{uv}(\Lambda_L(u)).$$

Now, $h \circ \pi_{uv}(\Lambda_L(u)) = h(\Lambda_R(v))$ whenever the constraint π_{uv} is satisfied. Further, each $\zeta_i \in \mathbb{L}_s(\mathbb{F}_2^L)$ is non-zero on a fraction 2^{-s} of the inputs, and thus $\zeta(\Lambda_L(u))$ is non-zero on at most a fraction $w2^{-s}$ of inputs. Hence, the constraints are satisfied with probability at least $(1 - w2^{-s})$. □

[2]For the values i such that p_i is linear we could allow q_i to be of degree $s - 1$ but as this does not greatly improve the result and just causes notational problems, we ignore this point.

Next, we analyze a generic assignment λ to $\mathcal{I}$ taking values 0 and 1, assuming $\mathsf{val}(\mathcal{L})$ is small. For each $u \in U$, set $\lambda_u : \mathbb{P}_s(\mathbb{F}_2^L) \to \{+1,-1\}$ as $\lambda_u(f) = (-1)^{\lambda((u,f))}$ and similarly $\lambda_v : \mathbb{P}_s(\mathbb{F}_2^R) \to \{+1,-1\}$ for each $v \in V$.

We show that $\mathsf{val}(\lambda;\mathcal{I})$ is close to the quantity $\mathbf{Pr}_{v,h}\{\lambda_v(h) = 1\}$. In the proof of the main theorem, we show that $\mathcal{L}$ can be amended (by making many copies of each v) so that any balanced assignment λ satisfies $\mathbf{Pr}_{v,h}\{\lambda_v(h) = 1\} \simeq 1/2$. This proves our main inapproximability result in light of Theorems 3.1 and 3.2.

Theorem 4.3 (Soundness)**.** *For every $\eta, \xi > 0$, and every integer s, if $\mathcal{L}$ is η-mixing, ξ-smooth, then every assignment λ to the instance $\mathcal{I}$ output by Theorem 4.1 satisfies*

$$\mathsf{val}(\lambda;\mathcal{I}) \leq \max\left(\frac{1}{2}, \mathop{\mathbf{Pr}}_{\substack{(\cdot,v)\leftarrow E,\\ h\leftarrow\mathbb{P}_s(\mathbb{F}^R)}}\{\lambda_v(h) = 1\}\right) + \exp(-2^{s/4}) + \eta + 2^s\xi + 2^{s/2}\sqrt{\mathsf{val}(\mathcal{L})}.$$

Proof. Fix an assignment $\lambda = \{\lambda_u\}_{\{u\in U\}} \cup \{\lambda_v\}_{\{v\in V\}}$. Writing out the Fourier expansion, we have:

$$\begin{aligned}
\mathsf{val}(\lambda;\mathcal{I}) &= \frac{1}{2} + \frac{1}{2}\mathop{\mathbf{E}}_{u,v,g,h,\zeta}[\lambda_u(g)\,\lambda_v(h)\,\lambda_u(g+\zeta+h\circ\pi_e)] \\
&= \frac{1}{2} + \frac{1}{2}\mathop{\mathbf{E}}_{u,v}\left[\sum_{\alpha,\beta,\gamma}\hat{\lambda}_u(\alpha)\,\hat{\lambda}_v(\beta)\,\hat{\lambda}_u(\gamma)\mathop{\mathbf{E}}_{g,h,\zeta}\left[\chi_\alpha(g)\chi_\beta(h)\chi_\gamma(g+h\circ\pi_{uv}+\zeta)\right]\right] \\
&= \frac{1}{2} + \frac{1}{2}\mathop{\mathbf{E}}_{u,v}\left[\sum_{\alpha}(\hat{\lambda}_u(\alpha))^2\,\hat{\lambda}_v(\pi_2(\alpha))\mathop{\mathbf{E}}_{\zeta}[\chi_\alpha(\zeta)]\right]. \qquad (4.2)
\end{aligned}$$

The last equality follows as if $\alpha \neq \gamma$ then taking expectation over g yields 0 and similarly looking at expectation over h only terms with $\pi_2(\alpha) = \beta$ are nonzero.

We recall that $\alpha \in \mathbb{X}_s(\mathbb{F}_2^L)$ and bound the expectation in eq. (4.2) by a case distinction on $\mathsf{wt}(\alpha)$.

Case 1: $\mathsf{wt}(\alpha) \geq 2^{s/2}$. The contribution of these terms is bounded by applying Theorem 2.8. Indeed, for such an α, we have $|\mathbf{E}_{\zeta_i\leftarrow\mathbb{L}_s}[\chi_\alpha(\zeta_i)]| \leq 1 - 2^{-s/2}$. Since ζ is a sum of $2^{3s/4}$ such independent terms, we have:

$$\begin{aligned}
\sum_{\alpha|\mathsf{wt}(\alpha)\geq 2^{s/2}}(\hat{\lambda}_u(\alpha))^2\,\hat{\lambda}_v(\pi_2(\alpha))\mathop{\mathbf{E}}_{\zeta}[\chi_\alpha(\zeta)] &\leq \sum_{\alpha|\mathsf{wt}(\alpha)\geq 2^{s/2}}(\hat{\lambda}_u(\alpha))^2(1-2^{-s/2})^{2^{3s/4}} \qquad (4.3)\\
&\leq (1-2^{-s/2})^{2^{3s/4}} \leq \exp(-2^{s/4}). \qquad (4.4)
\end{aligned}$$

Case 2: $\mathsf{wt}(\alpha) = 0$. Here we use the mixing property. Since $\mathcal{L}$ is η-mixing,

$$\begin{aligned}
\mathop{\mathbf{E}}_{u,v}\left[(\hat{\lambda}_u(\emptyset))^2\,\hat{\lambda}_v(\emptyset)\right] &\leq \mathop{\mathbf{E}}_{u}\left[\max(0, \mathop{\mathbf{E}}_{v|_u}\left[\hat{\lambda}_v(\emptyset)\right])\right] \\
&\leq \eta + \max(0, \mathop{\mathbf{E}}_{v}\left[\hat{\lambda}_v(\emptyset)\right]) = \eta + \max(0, \mathop{\mathbf{E}}_{v,h}[\lambda_v(h)]) \qquad (4.5)
\end{aligned}$$

Case 3: $0 < \mathsf{wt}(\alpha) < 2^{s/2}$; $\mathsf{wt}(\pi_2(\alpha)) = 0$. Here we use smoothness. A necessary condition for $\pi_2(\alpha) = \emptyset$ is that at least two distinct elements of α map to the same element. Since $\mathcal{L}$ is ξ-smooth, for a fixed u, and α, the probability over the choice of v that any two elements of α map to the same element is bounded $\mathsf{wt}^2(\alpha)\xi \leq 2^s\xi$.

Case 4: $0 < \mathsf{wt}(\alpha) < 2^{s/2}$; $\mathsf{wt}(\pi_2(\alpha)) \neq 0$. Here we use a standard list-decoding argument as done by Dinur and Guruswami [9] where they prove the following lemma.

Lemma 4.4 (see Lemma 13, [9])**.**

$$\mathop{\mathbf{E}}_{u,v}\left[\sum_{\substack{\mathsf{wt}(\alpha)<2^{s/2},\\ \pi_2(\alpha)\neq 0}} (\hat{\lambda}_u(\alpha))^2(\hat{\lambda}_v(\pi_2(\alpha)))^2\right] \leq 2^{s/2}\,\mathsf{val}(\mathcal{L}).$$

In view of the above lemma, and using Cauchy-Schwartz, we see that:

$$\begin{aligned}\mathop{\mathbf{E}}_{u,v}\left[\sum_{\substack{\mathsf{wt}(\alpha)<2^{s/2},\\ \pi_2(\alpha)\neq 0}} (\hat{\lambda}_u(\alpha))^2\,\hat{\lambda}_v(\pi_2(\alpha))\mathop{\mathbf{E}}_{\zeta}[\chi_\alpha(\zeta)]\right] &\leq \sqrt{\mathop{\mathbf{E}}_{u,v}\left[\sum_{\substack{\mathsf{wt}(\alpha)<2^{s/2},\\ \pi_2(\alpha)\neq 0}} (\hat{\lambda}_u(\alpha))^2(\hat{\lambda}_v(\pi_2(\alpha)))^2\right]}\\ &\leq \sqrt{2^{s/2}\,\mathsf{val}(\mathcal{L})} \qquad (4.6)\end{aligned}$$

These exhaust all the cases and substituting these bounds in eq. (4.2), we have:

$$\begin{aligned}\mathsf{val}(\lambda;\mathcal{J}) &= \frac{1}{2} + \frac{1}{2}\mathop{\mathbf{E}}_{u,v}\left[\sum_{\alpha}(\hat{\lambda}_u(\alpha))^2\,\hat{\lambda}_v(\pi_2(\alpha))\mathop{\mathbf{E}}_{\zeta}[\chi_\alpha(\zeta)]\right]\\ &\leq \frac{1}{2} + \frac{1}{2}\left(\exp(-2^{s/4}) + \eta + \max(0, \mathop{\mathbf{E}}_{v,h}[\lambda_v(h)]) + 2^s\xi + \sqrt{2^s\,\mathsf{val}(\mathcal{L})}\right)\\ &\leq \max\left(\frac{1}{2}, \mathop{\mathbf{Pr}}_{\substack{(\cdot,v)\leftarrow E,\\ h\leftarrow\mathbb{P}_s(\mathbb{F}^R)}}\{\lambda_v(h)=1\}\right) + \exp(-2^{s/4}) + \eta + 2^s\xi + \sqrt{2^s\,\mathsf{val}(\mathcal{L})} \qquad \square\end{aligned}$$

Now, we prove the main theorem.

Proof of Theorem 1.1. Fix a language $\mathsf{L} \in \mathsf{NP}$. We write $|w|$ to denote the size of an instance $w \in \mathsf{L}$. Applying Theorem 3.1, we have a graph G which is 3-colorable if $w \in \mathsf{L}$ and θ-far from being 3-colorable if $w \notin \mathsf{L}$. Further, $n = |G| = |w|^{\mathrm{O}(1)}$.

For an $\varepsilon > 0$, set $s = 4\log(1/\varepsilon)$ so that Theorem 4.2 applied on a $\mathcal{L}$ with value 1 shows that $\mathsf{val}(\mathcal{J}) \geq 1-\varepsilon$. Set $d = \Omega(\log(n)/s)$, $m = d^{3s}$, and $q = 2^t = d^c$ so that the pre-conditions (eq. (3.1)) of Theorem 3.2 are satisfied and apply this theorem with $\delta = \varepsilon^7$. We note that while ε and hence s are constants, d and hence q are $\omega(1)$ and hence both mixing, η and smoothness ξ are $o(1)$.

Running the reduction promised in Theorem 3.2 yields a CLC $\mathcal{L}$ which, in turn, yields a Max-Hom-Bal-3-LIN instance $\mathcal{J}$ through Theorem 4.1. Set $\Gamma = |\mathcal{J}|/(\varepsilon^2|V|)$ and construct $\mathcal{L}'$ by making Γ identical

copies of each $v \in V$ in $\mathcal{L}$. Let $\mathcal{I}'$ be the output of Theorem 4.1 run on $\mathcal{L}'$. If $w \in \mathsf{L}$, then $\mathcal{L}$ and hence $\mathcal{L}'$ have value 1. Thus, from Theorem 4.2, we know that $\mathsf{val}(\mathcal{I}') \geq 1-\varepsilon$.

We claim that if λ is a balanced assignment to $\mathcal{I}'$, then

$$\mathop{\mathbf{Pr}}_{\substack{(\cdot,v)\leftarrow E,\\ h\leftarrow \mathbb{P}_s(\mathbb{F}^R)}} \{\lambda_v(h)=1\} \leq \frac{1}{2}+\varepsilon^2$$

as the variables $(u,\cdot)$ account only for an ε^2 fraction of all variables. Now, applying Theorem 4.3 yields that $\mathsf{val}(\mathcal{I}') \leq 1/2+\varepsilon$ asymptotically. Indeed both the terms η and $2^s\xi$ are $o(1)$ (as a function of n) while $\exp(-2^{s/4})$ and $2^{s/2}\sqrt{\mathsf{val}(\mathcal{L})}$ are $O(1)$ as a function of n but $o(\varepsilon)$ as function of ε.

Finally, we verify the claim on the size of the reduction. From item 7 of Theorem 3.2, we know that $|\mathcal{L}| \leq (q^m n)^{\mathrm{O}(\log(1/\delta))} \leq \exp(\mathrm{O}(\log n \log(1/\varepsilon)))$. The length of the low-degree long code is at most $2^{L^{\mathrm{O}(s)}}$ and thus $|\mathcal{I}| \leq \exp((\log n)^{\mathrm{O}(\log(1/\varepsilon))})$. Therefore,

$$\left|\mathcal{I}'\right| \leq |\mathcal{L}| \cdot |\mathcal{I}|/\varepsilon^2 \cdot \exp(L^{\mathrm{O}(s)}) \leq \exp\Big((\log n)^{\mathrm{O}(\log(1/\varepsilon))}\Big) \qquad \square$$

5 Hardness for Max-Bisection

Let us just observe that, as also pointed out in [13], we can use the gadget of [20] to get hardness for Max-Bisection.

Theorem 5.1. *For any $\varepsilon > 0$, there is no algorithm that runs in time $\exp((\log n)^{\mathrm{O}(1)})$ and approximates Max-Bisection within $\frac{15}{16}+\varepsilon$ unless $\mathsf{NP} \subseteq \mathsf{DTIME}\big(\exp\big((\log n)^{\mathrm{O}(1)}\big)\big)$.*

Proof. Let us describe the properties needed of the gadget-reduction of [20]. It introduces a global variable z and for each equation of the form $x_1+x_2+x_3=0$ it introduces 8 auxiliary variables $(y_i)_{i=0}^7$ which are local to this equation and outputs 18 equations of the form $a+b=1$. The constructed equations have the property that if $z=0$ and the equation $x_1+x_2+x_3=0$ is satisfied, then there is a setting of the auxiliary variables satisfying 16 equations. On the other if $z=0$ and $x_1+x_2+x_3=1$ then the best setting of the variables satisfies 14 equations. Note that if we complement all variables then this preserves any sum of two variables and thus it is possible to focus on the case when $z=0$.

Now we start with an instance of Max-Hom-Bal-3-LIN constructed to prove Theorem 1.2. Suppose that it has N variables and contains M equations. If we straightforwardly apply the gadget construction to each equation we get a new system of equations each being on the form $a+b=1$ such that this system contains $1+N+8M$ variables and $18M$ equations.

Clearly the condition $a+b=1$ is a cut condition and we need to make sure to address the bisection property. In the constructed instance the new auxiliary variables are in majority but we change this by simply making copies of the original variables.

Let T be a parameter, to be determined, and let us make T copies $(x_i^j)_{j=0}^{T-1}$ of each original variable x_i. Now for each original equation $x_{i_1}+x_{i_2}+x_{i_3}=0$ we create T^3 gadgets, all using the same auxiliary variables while using all different combinations of copies of the original variables. This way we get an instance with $TN+8M+1$ variables and $18T^3M$ equations. Let us analyze this reduction.

If there is a balanced solution $x_i = a_i$ to the original set of equations that satisfies $(1-\varepsilon)M$ equations then we make a solution to the new system by setting $z = 0$, $x_i^j = a_i$ for all j and giving the auxiliary variables their optimal values. This satisfies at least $16(1-\varepsilon)T^3M$ equations. The solution is not, however, completely balanced has we have no control over the balance of the auxiliary variables. However, by changing the value of (all copies) of at most $8M/T$ original variables we can make the solution biased. Taking the least frequently occurring variables this may yield another $16T^3M \cdot \frac{16M}{NT}$ falsified equations, and we conclude that the system has a balanced assignment satisfying $16(1-\varepsilon-\frac{16M}{NT})T^3M$ equations. Provided that $T \geq \frac{16M}{N\varepsilon}$ this number is at least $16(1-2\varepsilon)T^3M$.

Now suppose we have as assignment that satisfies at least $(15+\delta)T^3M$ of the produced cut-equations. First we claim that we can modify the assignment such that all but one value of i, all copies of x_i take the same value maintaining at least as many satisfied equations.

Indeed suppose that both x_i and $x_{i'}$ have copies taking both values. Suppose each x_i^j has a neighbors taking the value 0 and b neighbors taking the value 1. Note that, by construction, these numbers do not depend on j. Let the similar numbers for $x_{i'}$ be a' and b'. Now suppose, without loss of generality that $a+b'-(a'+b) \geq 0$ and consider what happens if we change the value of a copy of x_i from 0 to 1 and and a copy of $x_{i'}$ from 1 to 0. For notational simplicity let us assume that these copies are x_i^0 and $x_{i'}^0$.

Suppose there are t equations containing both x_i^0 and $x_{i'}^0$. Then the number of satisfied equations containing either variable before the switch is $b+a'-t$ and after the switch $a+b'+t$. As $t \geq 0$ we see that we satisfy at least as many equations after the switch. By repeatedly doing similar switches we arrive at a solution with the same number of ones, that satisfies at least as many equations and such that at all but at most one of the original variables has all its copies take the same value.

Let us consider the assignment to the original system given by this unanimous value (taking either value for the "mixed-up" variable).

First we claim that this assignment is balanced within relative error $O(M/NT)$. This follows as the assignment to all the variables was biased and the fraction of auxiliary variables is $O(M/NT)$. Furthermore it still satisfies $(15+\delta)T^3M$ variables and hence the just defined assignment to the original variables must, by the property of the gadget satisfy at least $(\frac{1}{2}+\delta)M$ original equations. Setting $T = \omega(M/(\delta N))$, and adjusting the values of ε and δ, Theorem 4.3 is sufficient to establish the quality of the reduction.

The blow-up in size is only polynomial and hence we have established Theorem 5.1. □

6 Conclusion

Our paper proves tight inapproximability of Max-Hom-Bal-3-LIN assuming that NP-hard languages do not have quasipolynomial time algorithms. A standard NP-hardness of approximation still eludes us.

A situation where our understanding is even more limited is Min-Bisection where further progress in the form of a good algorithm or a stronger hardness result is badly needed.

Acknowledgment. We are grateful to three referees for a very careful reading of a preliminary version of the paper and in particular for pointing out a difference in two probability distributions claimed to be the same. We also thank one referee for reminding us of the reduction to Max-Bisection given in [13].

References

[1] SANJEEV ARORA AND MADHU SUDAN: Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003. 5, 6

[2] P. AUSTRIN AND J. HÅSTAD: Randomly supported independence and resistance. *SIAM Journal on Computing*, 40:1–27, 2011. 2

[3] PER AUSTRIN, SIAVOSH BENABBAS, AND KONSTANTINOS GEORGIOU: Better balance by being biased: A 0.8776-approximation for max bisection. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 277–294. SIAM, 2013. 2

[4] PER AUSTRIN AND ELCHANAN MOSSEL: Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18:249–271, 2009. 2

[5] PER AUSTRIN, RYAN O'DONNELL, LI-YANG TAN, AND JOHN WRIGHT: New np-hardness results for 3-coloring and 2-to-1 label cover. *ACM Trans. Comput. Theory*, 6(1):2:1–2:20, 2014. 7

[6] BOAZ BARAK, PARIKSHIT GOPALAN, JOHAN HÅSTAD, RAGHU MEKA, PRASAD RAGHAVENDRA, AND DAVID STEURER: Making the long code shorter. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pp. 370–379. IEEE Computer Society, 2012. 3, 11

[7] ARNAB BHATTACHARYYA, SWASTIK KOPPARTY, GRANT SCHOENEBECK, MADHU SUDAN, AND DAVID ZUCKERMAN: Optimal testing of reed-muller codes. In *51st Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pp. 488–497. IEEE Computer Society, 2010. 7

[8] SIU ON CHAN: Approximation resistance from pairwise independent subgroups. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, Chicago*, pp. 447–456. ACM, 2013. 2

[9] IRIT DINUR AND VENKATESAN GURUSWAMI: Pcps via low-degree long code and hardness for constrained hypergraph coloring. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pp. 340–349. IEEE Computer Society, 2013. 3, 7, 11, 14

[10] M. GOEMANS AND D. WILLIAMSON: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42:1115–1145, 1995. 2

[11] VENKATESAN GURUSWAMI, PRAHLADH HARSHA, JOHAN HÅSTAD, SRIKANTH SRINIVASAN, AND GIRISH VARMA: Super-polylogarithmic hypergraph coloring hardness via low-degree long codes. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pp. 614–623, New York, NY, USA, 2014. ACM. 3

[12] JOHAN HÅSTAD: Some optimal inapproximability results. *Journal of the ACM*, 48:798–859, 2001. 2, 3

[13] JONAS HOLMERIN AND SUBHASH KHOT: A new PCP outer verifier with applications to homogeneous linear equations and max-bisection. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pp. 11–20. ACM, 2004. 2, 3, 7, 10, 15, 16

[14] SUBHASH KHOT: Ruling out PTAS for graph min-bisection, dense k-subgraph, and bipartite clique. *SIAM J. Comput.*, 36(4):1025–1071, 2006. 2, 3

[15] SUBHASH KHOT, GUY KINDLER, ELCHANAN MOSSEL, AND RYAN O'DONNELL: Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. 2

[16] SUBHASH KHOT, MADHUR TULSIANI, AND PRATIK WORAH: A characterization of strong approximation resistance. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pp. 634–643, New York, NY, USA, 2014. ACM. 2

[17] EREZ PETRANK: The hardness of approximation: Gap location. *Computational Complexity*, 4:133–157, 1994. 7

[18] HARALD RÄCKE: Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pp. 255–264, New York, NY, USA, 2008. ACM. 2

[19] ANUP RAO: Parallel repetition in projection games and a concentration bound. *SIAM J. Comput.*, 40(6):1871–1891, 2011. 10

[20] LUCA TREVISAN, GREGORY B. SORKIN, MADHU SUDAN, AND DAVID P. WILLIAMSON: Gadgets, approximation, and linear programming. *SIAM J. Comput.*, 29(6):2074–2097, 2000. 15

AUTHORS

Johan Håstad
Professor
KTH Royal Institute of Technology, Sweden.
johanh@kth.se
http://www.csc.kth.se/~johanh

Rajsekar Manokaran
Assistant professor
Indian Institute of Technology, Madras, India.
rajsekar@gmail.com
http://www.cse.iitm.ac.in/~rajsekar

ABOUT THE AUTHORS

JOHAN HÅSTAD received his Bachelor of Science from Stockholm University in 1981, his Master of Science from Uppsala University in 1984, and his Ph. D. from MIT in 1986 under the supervision of Shafi Goldwasser. Johan was appointed Associate Professor at the Royal Institute of Technology in Stockholm, Sweden in 1988 and advanced to the level of Professor in 1992. He was elected a member of the Swedish Royal Academy of Sciences in 2001. He has research interests within several subareas of Theory of Algorithms and Complexity theory but has recently mainly focused on the inapproximability of NP-hard optimization problems.

RAJSEKAR MANOKARAN graduated from Princeton in 2012; his advisor was Sanjeev Arora. The subject of his thesis was the inapproximability of constraint satisfaction problems using convex relaxations. He was appointed as an assistant professor at the Indian Institute of Technology in Madras, India in 2013.