Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 99 (2009)

Improved Inapproximability Results for Maximum k-Colorable Subgraph

VENKATESAN GURUSWAMI* ALI KEMAL SINOP*

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213.

Abstract

We study the maximization version of the fundamental graph coloring problem. Here the goal is to color the vertices of a k-colorable graph with k colors so that a maximum fraction of edges are properly colored (i.e. their endpoints receive different colors). A random k-coloring properly colors an expected fraction $1-\frac{1}{k}$ of edges. We prove that given a graph promised to be k-colorable, it is NP-hard to find a k-coloring that properly colors more than a fraction $\approx 1-\frac{1}{33k}$ of edges. Previously, only a hardness factor of $1-O\big(\frac{1}{k^2}\big)$ was known. Our result pins down the correct asymptotic dependence of the approximation factor on k. Along the way, we prove that approximating the Maximum 3-colorable subgraph problem within a factor greater than $\frac{32}{33}$ is NP-hard.

Using semidefinite programming, it is known that one can do better than a random coloring and properly color a fraction $1-\frac{1}{k}+\frac{2\ln k}{k^2}$ of edges in polynomial time. We show that, assuming the 2-to-1 conjecture, it is hard to properly color (using k colors) more than a fraction $1-\frac{1}{k}+O\left(\frac{\ln k}{k^2}\right)$ of edges of a k-colorable graph.

*Research supported in part by a Packard Fellowship. Email: guruswami@cmu.edu, asinop@cs.cmu.edu

ISSN 1433-8092

1 Introduction

1.1 Problem statement

A graph $G = (V, E)$ is said to be k-colorable for some positive integer k if there exists a k-coloring $\chi : V \to \{1, 2, \ldots, k\}$ such that for all edges $(u, v) \in E$, $\chi(u) \neq \chi(v)$. For $k \geqslant 3$, finding a k-coloring of a k-colorable graph is a classic NP-hard problem. The problem of coloring a graph with the fewest number of colors has been extensively studied. In this paper, our focus is on hardness results for the following maximization version of graph coloring: Given a k-colorable graph (for some fixed constant $k \geqslant 3$), find a k-coloring that maximizes the fraction of properly colored edge. (We say an edge is properly colored under a coloring if its endpoints receive distinct colors.) Note that for $k = 2$ the problem is trivial — one can find a proper 2-coloring in polynomial time when the graph is bipartite (2-colorable).

We will call this problem Max k-Colorable Subgraph. The problem is equivalent to partitioning the vertices into k parts so that a maximum number of edges are cut. This problem is more popularly referred to as Max k-Cut in the literature; however, in the Max k-Cut problem the input is an arbitrary graph that need not be k-colorable. To highlight this difference that our focus is on the case when the input graph is k-colorable, we use Max k-Colorable Subgraph to refer to this variant. We stress that we will use this convention throughout the paper: Max k-Colorable Subgraph **always** *refers to the "perfect completeness" case, when the input graph is k-colorable.*[1] Since our focus is on hardness results, we note that this restriction only makes our results stronger.

A factor $\alpha = \alpha_k$ approximation algorithm for Max k-Colorable Subgraph is an efficient algorithm that given as input a k-colorable graph outputs a k-coloring that properly colors at least a fraction α of the edges. We say that Max k-Colorable Subgraph is NP-hard to approximate within a factor β if no factor β approximation algorithm exists for the problem unless $\mathrm{P} = \mathrm{NP}$. The goal is to determine the approximation threshold of Max k-Colorable Subgraph: the largest α as a function of k for which a factor α approximation algorithm for Max k-Colorable Subgraph exists.

1.2 Previous results

The algorithm which simply picks a random k-coloring, without even looking at the graph, properly colors an expected fraction $1 - 1/k$ of edges. Frieze and Jerrum [3] used semidefinite programming to give a polynomial time factor $1 - 1/k + 2\ln k/k^2$ approximation algorithm for Max k-Cut, which in particular means the algorithm will color at least this fraction of edges in a k-colorable graph. This remains the best known approximation guarantee for Max k-Colorable Subgraph to date. Khot, Kindler, Mossel, and O'Donnell [7] showed that obtaining an approximation factor of $1-1/k+2\ln k/k^2+\Omega(\ln\ln k/k^2)$ for Max k-Cut is Unique Games-hard, thus showing that the Frieze-Jerrum algorithm is essentially the best possible. However, due to the "imperfect completeness" inherent to the Unique Games conjecture, this hardness result does *not* hold for Max k-Colorable Subgraph when the input is required to be k-colorable.

For Max k-Colorable Subgraph, the best hardness known prior to our work was a factor $1 - \Theta(1/k^2)$. This is obtained by combining an inapproximability result for Max 3-Colorable Subgraph due to Petrank [11] with a reduction from Papadimitriou and Yannakakis [10]. It is a natural question whether is an efficient algorithm that could properly color a fraction $1 - 1/k^{1+\varepsilon}$ of edges given a k-colorable graph for some absolute constant $\varepsilon > 0$. The existing hardness results do not rule out the possibility of such an algorithm.

[1] While a little non-standard, this makes our terminology more crisp, as we can avoid repeating the fact that the hardness holds for k-colorable graphs in our statements.

For Max k-Cut, a better hardness factor was shown by Kann, Khanna, Lagergren, and Panconesi [5] — for some absolute constants $\beta > \alpha > 0$, they showed that it is NP-hard to distinguish graphs that have a k-cut in which a fraction $(1-\alpha/k)$ of the edges cross the cut from graphs whose Max k-cut value is at most a fraction $(1-\beta/k)$ of edges. Since MaxCut is easy when the graph is 2-colorable, this reduction does not yield any hardness for Max k-Colorable Subgraph.

1.3 Our results

Petrank [11] showed the existence of a $\gamma_0 > 0$ such that it is NP-hard to find a 3-coloring that properly colors more than a fraction $(1-\gamma_0)$ of the edges of a 3-colorable graph. The value of γ_0 in [11] was left unspecified and would be very small if calculated. The reduction in [11] was rather complicated, involving expander graphs and starting from the weak hardness bounds for bounded occurrence satisfiability. We prove that the NP-hardness holds with $\gamma_0 = \frac{1}{33}$. In other words, it is NP-hard to obtain an approximation ratio bigger than $\frac{32}{33}$ for Max 3-Colorable Subgraph. The reduction is from the constraint satisfaction problem corresponding to the adaptive 3-query PCP with perfect completeness from [4].

By a reduction from Max 3-Colorable Subgraph, we prove that for every $k \geqslant 3$, the Max k-Colorable Subgraph is NP-hard to approximate within a factor greater than $\approx 1 - \frac{1}{33k}$ (Theorem 4). This identifies the correct asymptotic dependence on k of the best possible approximation factor for Max k-Colorable Subgraph. The reduction is similar to the one in [5], though some crucial changes have to be made in the construction and some new difficulties overcome in the soundness analysis when reducing from Max 3-Colorable Subgraph instead of MaxCut.

In the quest for pinning down the *exact* approximability of Max k-Colorable Subgraph, we prove the following *conditional* result. Assuming the so-called 2-to-1 conjecture, it is hard to approximate Max k-Colorable Subgraph within a factor $1 - \frac{1}{k} + O\left(\frac{\ln k}{k^2}\right)$. In other words, the Frieze-Jerrum algorithm is optimal up to lower order terms in the approximation ratio *even for instances of* Max k-Cut *where the graph is k-colorable.*

Unlike the Unique Games Conjecture (UGC), the 2-to-1 conjecture allows perfect completeness, i.e., the hardness holds even for instances where an assignment satisfying *all* constraints exists. The 2-to-1 conjecture was used by Dinur, Mossel, and Regev [2] to prove that for every constant c, it is NP-hard to color a 4-colorable graph with c colors. We analyze a similar reduction for the k-coloring case when the objective is to maximize the fraction of edges that are properly colored by a k-coloring. Our analysis uses some of the machinery developed in [2], which in turn extends the invariance principle of [8]. The hardness factor we obtain depends on the spectral gap of a certain $k^2 \times k^2$ stochastic matrix.

Remark 1. In general it is far from clear which Unique Games-hardness results can be extended to hold with perfect completeness by assuming, say, the 2-to-1 (or some related) conjecture. In this vein, we also mention the result of O'Donnell and Wu [9] who showed a tight hardness for approximating satisfiable constraint satisfaction problems on 3 Boolean variables assuming the d-to-1 conjecture for any fixed d. While the UGC assumption has led to a nearly complete understanding of the approximability of constraint satisfaction problems [12], the approximability of *satisfiable* constraint satisfaction problems remains a mystery to understand in any generality.

Remark 2. It has been shown by Crescenzi, Silvestri and Trevisan [1] that any hardness result for weighted instances of Max k-Cut carries over to unweighted instances assuming the total edge weight is polynomially bounded. In fact, their reduction preserves k-colorability, so an inapproximability result for the weighted Max k-Colorable Subgraph problem also holds for the unweighted version. Therefore all our hardness results hold for the unweighted Max k-Colorable Subgraph problem.

2 Unconditional Hardness Results for Max k-Colorable Subgraph

We will first prove a hardness result for Max 3-Colorable Subgraph, and then reduce this problem to Max k-Colorable Subgraph.

2.1 Inapproximability result for Max 3-Colorable Subgraph

Petrank [11] showed that Max 3-Colorable Subgraph is NP-hard to approximate within a factor of $(1-\gamma_0)$ for some constant $\gamma_0 > 0$. This constant γ_0 is presumably very small, since the reduction starts from bounded occurrence satisfiability (for which only weak inapproximability results are known) and uses expander graphs. We prove a much better inapproximability factor below, via a simpler proof.

Theorem 1 (Max 3-Colorable Subgraph Hardness)**.** *The* Max 3-Colorable Subgraph *problem is NP-hard to approximate within a factor of* $\frac{32}{33} + \varepsilon$ *for any constant* $\varepsilon > 0$.

Proof. For the proof of this theorem, we will use reduce from a hard to approximate constraint satisfaction problem (CSP) underlying the adaptive 3-query PCP given in [4]. This PCP has perfect completeness and soundness $1/2 + \varepsilon$ for any desired constant ε (which is the best possible for 3-query PCPs).

We first state the properties of the CSP. An instance of the CSP will have variables partitioned into three parts $\mathcal{X}, \mathcal{Y}$ and $\mathcal{Z}$. Each constraint will be of the form $(x_i \vee (Y_j = z_k)) \wedge (\overline{x_i} \vee (Y_j = z_l))$, where $x_i \in \mathcal{X}$, $z_k, z_l \in \mathcal{Z}$ are variables (unnegated) and Y_j is a literal ($Y_j \in \{y_j, \overline{y_j}\}$ for some variable $y_j \in \mathcal{Y}$). For YES instances of the CSP, there will be a Boolean assignment that satisfies **all** the constraints. For NO instances, every assignment to the variables will satisfy at most a fraction $(1/2 + \varepsilon)$ of the constraints.

Remark 3. We remark the condition that the instance is tripartite, and that the variables in $\mathcal{Z}$ never appear negated are not explicit in [4]. But these can be ensured by an easy modification to the PCP construction in [4]. The PCP in [4] has a bipartite structure: the proof is partitioned into two parts called the A-tables and B-tables, and each test consists of probing one bit $A(f)$ from an A table and 3 bits $B(g), B(g_1), B(g_2)$ from the B table, and checking $(A(f) \vee (B(g) = B(g_1)) \wedge (\overline{A(f)} \vee (B(g) = B(g_2))$. Further these tables are *folded* which is a technical condition that corresponds to the occurrence of negations in the CSP world. If the queries at locations g_1 and g_2 are made in a parallel C-table, and even if the C-table is not folded (though the A and B tables need to be folded), one can verify that the analysis of the PCP construction still goes through. This then translates to a CSP with the properties claimed above.

Let $\mathcal{I}$ be an instance of such a CSP with m constraints of the above form on variables $\mathcal{V} = \mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$. Let $\mathcal{X} = \{x_1, x_2, \ldots, x_{n_1}\}$, $\mathcal{Y} = \{y_1, y_2, \ldots, y_{n_2}\}$ and $\mathcal{Z} = \{z_1, z_2, \ldots, z_{n_3}\}$. From the instance $\mathcal{I}$ we create a graph G for the Max 3-Colorable Subgraph problem as follows. There is a node x_i for each variable $x_i \in \mathcal{X}$, a node z_l for each $z_l \in \mathcal{Z}$, and a pair of nodes $\{y_j, \overline{y_j}\}$ for the two literals corresponding to each $y_j \in \mathcal{Y}$. There are also three global nodes $\{R, T, F\}$ representing boolean values which are connected in a triangle with edge weights $m/2$ (see Fig. 1).

For each constraint of the CSP, we place the local gadget specific to that constraint shown in Figure 2. Note that there are 10 edges of unit weight in this gadget. The nodes $y_j, \overline{y_j}$ are connected to node R by a triangle whose edge weights equal $w_j = \frac{\Delta(y_j)+\Delta(\overline{y_j})}{2}$. Here $\Delta(X)$ denotes the total number of edges going from node X *into* all the local gadgets. The nodes x_i and z_l connected to R with an edge of weight $\Delta(x_i)/2$ and $\Delta(z_l)/2$ respectively. The proofs of the following lemmas appear in Appendix A.

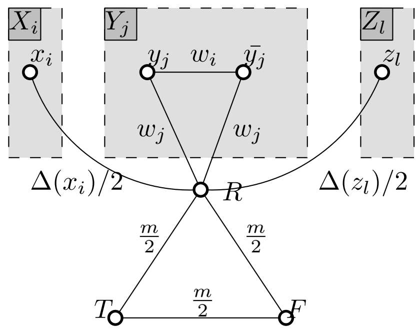

Figure 1: Global gadget for truth value assignments. Blocks X_i, Y_j and Z_l are replicated for all vertices in $\mathcal{X}$, $\mathcal{Y}$ and $\mathcal{Z}$. Edge weights are shown next to each edge.

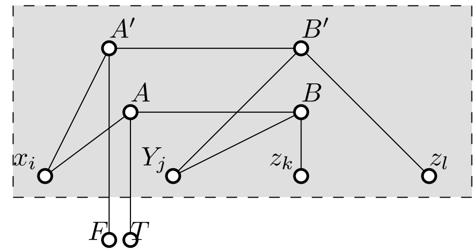

Figure 2: Local gadget for each constraint of the form $(x_i \vee Y_j = z_k) \wedge (\overline{x_i} \vee Y_j = z_l)$. All edges have unit weight. Labels A, A', B, B' refer to the local nodes in each gadget.

Lemma 2 (Completeness). *Given an assignment of variables $\sigma : \mathcal{V} \to \{0,1\}$ which satisfies at least c of the constraints, we can construct a 3-coloring of G with at most $m - c$ improperly colored edges (each of weight 1).*

Lemma 3 (Soundness). *Given a 3-coloring of G, χ, such that the total weight of edges that are not properly colored by χ is at most $\tau < m/2$, we can construct an assignment $\sigma' : \mathcal{V} \to \{0,1\}$ to the variables of the CSP instance that satisfies at least $m - \tau$ constraints.*

Returning to the proof of Theorem 1, the total weight of edges in G is

$$10m + \frac{3m}{2} + \underbrace{\sum_{i=1}^{n_1} \frac{\Delta(x_i)}{2}}_{m} + \sum_{j=1}^{n_2} 3w_j + \underbrace{\sum_{l=1}^{n_3} \frac{\Delta(z_l)}{2}}_{m} = \frac{27}{2}m + \frac{3}{2}\underbrace{\sum_{j=1}^{n_2}(\Delta(y_i) + \Delta(\overline{y_j}))}_{2m} = \frac{33}{2}m\,.$$

By the completeness lemma, YES instances of the CSP are mapped to graphs G that are 3-colorable. By the soundness lemma, NO instances of the CSP are mapped to graphs G such that every 3-coloring miscolors at least a fraction $\frac{(1/2-\varepsilon)}{33/2} = \frac{1-2\varepsilon}{33}$ of the total weight of edges. Since $\varepsilon > 0$ is an arbitrary constant, the proof of Theorem 1 is complete.[2] □

2.2 Max k-Colorable Subgraph Hardness

Theorem 4. *For every integer $k \geqslant 3$ and every $\varepsilon > 0$, it is NP-hard to approximate* Max k-Colorable Subgraph *within a factor of* $1 - \frac{1}{33(k+c_k)+c_k} + \varepsilon$ *where* $c_k = k \mod 3 \leqslant 2$.

Proof. We will reduce Max 3-Colorable Subgraph to Max k-Colorable Subgraph and then apply Theorem 1. Throughout the proof, we will assume k is divisible by 3. At the end, we will cover the remaining cases also. The reduction is inspired by the reduction from MaxCut to Max k-Cut given by Kann *et al.* [5] (see Remark 4). Some modifications to the reduction are needed when we reduce from Max 3-Colorable

[2] Our reduction produced a graph with edge weights, but by Remark 2, the same inapproximability factor holds for unweighted graphs as well.

Subgraph, and the analysis has to handle some new difficulties. The details of the reduction and its analysis follow.

Let $G = (V, E)$ be an instance of Max 3-Colorable Subgraph. By Theorem 1, it is NP-hard to tell if G is 3-colorable or every 3-colors miscolors a fraction $\frac{1}{33} - \varepsilon$ of edges. We will construct a graph H such that H is k-colorable when G is 3-colorable, and a k-coloring which miscolors at most a fraction μ of the total weight of edges of H implies a 3-coloring of G with at most a fraction μk of miscolored edges. Combined with Theorem 1, this gives us the claimed hardness of Max k-Colorable Subgraph.

Let $K'_{k/3}$ denote the complete graph with loops on $k/3$ vertices. Let G' be the tensor product graph between $K_{k/3}$ and G, $G' = K'_{k/3} \otimes G$ as defined by Weichsel [14]. Identify each node in G' with $(u, i), u \in V(G), i \in \{1, 2, \dots, k/3\}$. The edges of G' are $((u, i), (v, i'))$ for $(u, v) \in E$ and any $i, i' \in \{1, \dots, k/3\}$. Next we make 3 copies of G', and identify the nodes with $(u, i, j), (u, i) \in V(G'), j \in \{1, 2, 3\}$, then put edges between all nodes of the form (u, i, j) and (u, i', j') if either $i \neq i'$ or $j \neq j'$ with weight $\frac{2}{3} d_u$, where d_u is degree of node u. The total weight of edges in this new construction H equals

$$\sum_{u \in V} \left(\binom{k}{2} \frac{2}{3} d_u + \frac{3}{2} \binom{k}{3}^2 d_u \right) \leqslant k^2 m \; .$$

Lemma 5. *If G is 3-colorable, then H is k-colorable.*

Proof. Let $\chi_G : V(G) \to \{1, 2, 3\}$ be a 3-coloring of G. Consider the following coloring function for H, $\chi_H : V(H) \to \{1, 2, \dots, k\}$. For node (u, i, j), let $\chi_H((u, i, j)) = \pi^j(\chi_G(u)) + 3(i - 1)$. Here π is the permutation $\begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}$, and $\pi^j(x) = \underbrace{\pi(\dots(\pi(}_{\text{j times}} x)))$. Equivalently $\pi(x) = x \mod 3 + 1$.

Consider edges of the form $\{(u, i, j), (v, i', j)\}$. If $i \neq i'$, then colors of the endpoints are different. Else we have $\chi((u, i, j)) - \chi((v, i, j)) \equiv \chi(u) - \chi(v) \not\equiv 0 \mod 3$. For edges of the form $\{(u, i, j), (u, i', j')\}$, if $i \neq i'$, clearly edge is satisfied. When $i = i', j \neq j'$, $\chi((u, i, j)) - \chi((u, i, j')) \equiv \pi^j(u) - \pi^{j'}(u) \equiv j - j' \not\equiv 0 \mod 3$. ☐

Lemma 6. *If H has a k-coloring that properly colors a set of edges with at least a fraction $(1 - \mu)$ of the total weight, then G has a 3-coloring which colors at least a fraction $(1 - \mu k)$ of its edges properly.*

Proof. Let χ_H be the coloring of H, $\mathsf{Sugg}_u^j = \{\chi_H((u, i, j)) \mid 1 \leqslant i \leqslant k/3\}$ and $\mathsf{Sugg}_u = \bigcup_j \mathsf{Sugg}_u^j$. Denote the total weight of uncut edges in this solution as

$$C^{total} = \sum_{u \in V(G)} \frac{2}{3} d_u C_u^{within} + C^{between}, \tag{1}$$

where C_u^{within} and $C^{between}$ denotes the number of improperly colored edges within the copies of node u and between copies of different vertices $u, v \in V(G)$ respectively. We have the following relations:

$$\begin{aligned} C^{between} &= \textstyle\sum_{j=1}^{3} \sum_{uv \in E(G)} \sum_{1 \leqslant i \leqslant i' \leqslant k/3} 1_{\chi_H((u,i,j)) = \chi_H((v,i',j))} \\ &\geqslant \textstyle\sum_{j=1}^{3} \sum_{uv \in E(G)} |\mathsf{Sugg}_u^j \cap \mathsf{Sugg}_v^j| \end{aligned} \tag{2}$$

$$\begin{array}{rclr} C_u^{within} & = & \sum_{c \in \mathsf{Sugg}_u} \binom{|\chi_H^{-1}(c) \cap B_u|}{2} & (B_u = \{(u,i,j) | \forall i,j\}) \\ & = & \sum_{c \in \mathsf{Sugg}_u} \frac{|B_{u,c}|^2}{2} - \frac{k}{2} & (B_{u,c} = B_u \cap \chi_H^{-1}(c)) \\ & \geqslant & \frac{1}{2|\mathsf{Sugg}_u|} \left(\sum_{c \in \mathsf{Sugg}_u} |B_{u,c}| \right)^2 - \frac{k}{2} & \text{(Cauchy-Schwarz)} \\ & \geqslant & \frac{k}{2} \left(\frac{k}{|\mathsf{Sugg}_u|} - 1 \right) \geqslant \frac{k}{2} \frac{|\overline{\mathsf{Sugg}_u}|}{|\mathsf{Sugg}_u|} \geqslant \frac{|\overline{\mathsf{Sugg}_u}|}{2} & \end{array} \quad (3)$$

Now we will find a (random) 3-coloring χ_G for G. Pick c from $\{1, 2, \ldots, k\}$ uniformly at random. If $c \notin \mathsf{Sugg}_u$, select $\chi_G(u)$ uniformly at random from $\{1, 2, 3\}$. If $c \in \mathsf{Sugg}_u$, set $\chi_G(u) = j$ if j is the smallest index for which $c \in \mathsf{Sugg}^j(u)$. With this coloring $\chi_G(u)$, the probability that an edge $(u, v) \in E(G)$ will be improperly colored is:

$$\begin{aligned} \mathsf{Pr}\left[\chi_G(u) = \chi_G(v)\right] \;\; \leqslant \;\; & \sum_{j=1}^{3} \mathsf{Pr}_c \left[c \in \mathsf{Sugg}_u^j \cap \mathsf{Sugg}_v^j\right] + \frac{1}{3}\mathsf{Pr}_c\left[c \in \overline{\mathsf{Sugg}_u}, c \in \mathsf{Sugg}_v\right] \\ & + \frac{1}{3}\mathsf{Pr}_c\left[c \in \mathsf{Sugg}_u, c \in \overline{\mathsf{Sugg}_v}\right] + \frac{1}{3}\mathsf{Pr}_c\left[c \in \overline{\mathsf{Sugg}_u}, c \in \overline{\mathsf{Sugg}_v}\right] \\ \leqslant \;\; & \sum_{j=1}^{3} \frac{|\mathsf{Sugg}_u^j \cap \mathsf{Sugg}_v^j|}{k} + \frac{|\overline{\mathsf{Sugg}_u}|}{3k} + \frac{|\overline{\mathsf{Sugg}_v}|}{3k} \end{aligned}$$

We can thus bound the expected number of miscolored edges in the coloring χ_G as follows.

$$\begin{aligned} \mathbb{E}\left[\sum_{(u,v) \in E(G)} 1_{\chi_G(u) = \chi_G(v)}\right] \;\; \leqslant \;\; & \sum_{uv \in E} \left[\left(\sum_{j=1}^{3} \frac{|\mathsf{Sugg}_u^j \cap \mathsf{Sugg}_v^j|}{k}\right) + \frac{|\overline{\mathsf{Sugg}_u}|}{3k} + \frac{|\overline{\mathsf{Sugg}_v}|}{3k}\right] \\ \leqslant \;\; & \frac{1}{k}\Big(C^{between} + \sum_{u \in V(G)} \frac{d_u}{3} |\overline{\mathsf{Sugg}_u}|\Big) \quad \text{(using (2))} \\ \leqslant \;\; & \frac{1}{k}\Big(C^{between} + \sum_{u \in V(G)} \frac{2d_u}{3} C_u^{within}\Big) = \frac{C^{total}}{k} \end{aligned}$$

This implies that there exists a 3-coloring of G for which the number of improperly colored edges in G is at most $\frac{C^{total}}{k}$. Therefore if H has a k-coloring which improperly colors at most a total weight $\mu k^2 m$ of edges, then there is a 3-coloring of G which colors improperly at most a fraction $\frac{\mu k^2 m}{km} = \mu k$ of its edges. □

This completes the proof of Theorem 4 when k is divisible by 3. The other cases are easily handled by adding $k \mod 3$ extra nodes connected to all vertices by edges of suitable weight. See Appendix D for details. □

Remark 4 (Comparison to [5])**.** The reduction of Kann *et al* [5] converts an instance G of MaxCut to the instance $G' = K'_{k/2} \otimes G$ of Max k-Cut. Edge weights are picked so that the optimal k-cut of G' will give a set S_u of $k/2$ different colors to all vertices in each $k/2$ clique (u, i), $1 \leqslant i \leqslant k/2$. This enables converting a k-cut of G' into a cut of G based on whether a random color falls in S_u or not. In the 3-coloring case, we make 3 copies of G' in an attempt to enforce three "translates" of S_u, and use those to define a 3-coloring from a k-coloring. But we cannot ensure that each k-clique is properly colored, so these translates might overlap and a more careful soundness analysis is needed.

3 Conditional Hardness Results for Max k-Colorable Subgraph

We will first review the (exact) 2-to-1 Conjecture, and then construct a noise operator, which allows us to preserve k-colorability. Then we will bound the stability of coloring functions with respect to this noise operator. In the last section, we will give a PCP verifier which concludes the hardness result.

3.1 Preliminaries

We begin by reviewing some definitions and d-to-1 conjecture.

Definition 1. *An instance of a bipartite Label Cover problem represented as $\mathcal{L} = (U, V, E, W, R_U, R_V, \Pi)$ consists of a weighted bipartite graph over node sets U and V with edges $e = (u, v) \in E$ of non-negative real weight $w_e \in W$. R_U and R_V are integers with $1 \leqslant R_U \leqslant R_V$. Π is a collection of projection functions for each edge: $\Pi = \{\pi_{vu} : \{1, \ldots, R_V\} \to \{1, \ldots, R_U\} \big| u \in U, v \in V\}$. A labeling ℓ is a mapping $\ell : U \to \{1, \ldots, R_U\}$, $\ell : V \to \{1, \ldots, R_V\}$. An edge $e = (u, v)$ is satisfied by labeling ℓ if $\pi_e(\ell(v)) = \ell(u)$. We define the value of a labeling as sum of weights of edges satisfied by this labeling normalized by the total weight. $\mathsf{Opt}(\mathcal{L})$ is the maximum value over any labeling.*

Definition 2. *A projection $\pi : \{1, \ldots, R_V\} \to \{1, \ldots, R_U\}$ is called d-to-1 if for each $i \in \{1, \ldots, R_U\}$, $|\pi^{-1}(i)| \leqslant d$. It is called* exactly d-to-1 *if $|\pi^{-1}(i)| = d$ for each $i \in \{1, 2, \ldots, R_U\}$.*

Definition 3. *A bipartite Label-Cover instance $\mathcal{L}$ is called d-to-1 Label-Cover if all projection functions, $\pi \in \Pi$ are d-to-1.*

Conjecture 1 (d-to-1 Conjecture [6])**.** *For any $\gamma > 0$, there exists a d-to-1 Label-Cover instance $\mathcal{L}$ with $R_V = R(\gamma)$ and $R_U \leqslant dR_V$ many labels such that it is NP-hard to decide between two cases, $\mathsf{Opt}(\mathcal{L}) = 1$ or $\mathsf{Opt}(\mathcal{L}) \geqslant \gamma$. Note that although the original conjecture involves d-to-1 projection functions, we will assume that it also holds for* exactly d-to-1 *functions (so $R_U = dR_V$), which is the case in [2].*

Using the reductions from [2], it is possible to show that the above conjecture still holds given that the graph $(U \cup V, E)$ is left-regular and unweighted, i.e., $w_e = 1$ for all $e \in E$.

3.2 Noise Operators

For a positive integer M, we will denote by $[M]$ the set $\{0, 1, \ldots, M-1\}$. We will identify elements of $[M^2]$ with $[M] \times [M]$ in the obvious way, with the pair $(a, b) \in [M]^2$ corresponding $a + Mb \in [M^2]$.

Definition 4. *A Markov operator T is a linear operator which maps probability measures to other probability measures. In a finite discrete setting, it is defined by a stochastic matrix whose (x, y)'th entry $T(x \to y)$ is the probability of transitioning from x to y. Such an operator is called symmetric if $T(x \to y) = T(y \to x) = T(x \leftrightarrow y)$.*

Definition 5. *Given $\rho \in [-1, 1]$, the Beckner noise operator, T_ρ on $[q]$ is defined by as $T_\rho(x \to x) = \frac{1}{q} + \left(1 - \frac{1}{q}\right)\rho$ and $T_\rho(x \to y) = \frac{1}{q}(1 - \rho)$ for any $x \neq y$.*

Observation 1. *All eigenvalues of the operator T_ρ are given by $1 = \lambda_0(T_\rho) \geqslant \lambda_1(T_\rho) = \ldots = \lambda_{q-1}(T_\rho) = \rho$. Any orthonormal basis $\alpha_0, \alpha_1, \ldots, \alpha_{q-1}$ with α_0 being constant vector, is also a basis for T_ρ.*

Lemma 7. *For an integer $q \geqslant 6$, there exists a symmetric Markov operator T on $[q]^2$ whose diagonal entries are all 0 and with eigenvalues $1 = \lambda_0 \geqslant \lambda_1 \geqslant \ldots \geqslant \lambda_{q^2-1}$ such that the* spectral radius $\rho(T) = \max\{|\lambda_1|, |\lambda_{q^2-1}|\}$ *is at most $\frac{4}{q-1}$.*

Proof. Consider the symmetric Markov operator T on $[q]^2$ such that, for $x = (x_1, x_2), y = (y_1, y_2) \in [q]^2$,

$$T(x \leftrightarrow y) = \begin{cases} \alpha & \text{if } \{x_1, x_2\} \cap \{y_1, y_2\} = \emptyset \text{ and } x_1 \neq x_2, y_1 \neq y_2, \\ \beta & \text{if } x_1 \notin \{y_1, y_2\} \text{ and } x_1 = x_2, y_1 \neq y_2, \\ \beta & \text{if } y_1 \notin \{x_1, x_2\} \text{ and } x_1 \neq x_2, y_1 = y_2, \\ 0 & \text{else,} \end{cases}$$

where $\alpha = \frac{1}{(q-1)(q-3)}$ and $\beta = \frac{1}{(q-1)(q-2)}$. It is clear that T is symmetric and doubly stochastic.

To bound the spectral radius of T, we will bound the second largest eigenvalue $\lambda_1(T^2)$ of T^2. Notice that T^2 is also a symmetric Markov operator. Moreover $\lambda_i(T^2) = \lambda_i^2(T)$, therefore $\lambda_1(T^2) \geqslant \max(\lambda_1^2(T), \lambda_{q^2-1}^2(T)) \geqslant \rho(T)^2$.

Notice that $T^2(x \leftrightarrow y) > 0$ for all pairs $x, y \in [q]^2$. Consider the variational characterization of $1 - \lambda_1(T^2)$ [13]:

$$\min_{\psi} \frac{\sum_{x,y} (\psi(x) - \psi(y))^2 \pi(x) T^2(x \leftrightarrow y)}{\sum_{x,y} (\psi(x) - \psi(y))^2 \pi(x) \pi(y)} \geqslant \min_{\psi} \min_{x,y} \frac{\pi(x) (\psi(x) - \psi(y))^2 T^2(x \leftrightarrow y)}{(\psi(x) - \psi(y))^2 \pi(x) \pi(y)} = \min_{x,y} q^2 T^2(x \leftrightarrow y)$$

For any two pairs $(x_1, x_2), (y_1, y_2) \in [q]^2$, let $l = |[q] \setminus \{x_1, x_2, y_1, y_2\}|$. Then we have

$$\begin{aligned} T^2((x_1, x_2) \leftrightarrow (y_1, y_2)) &= \begin{cases} l(l-1)\beta^2 \geqslant (q-2)(q-3)\beta^2 & \text{if } x_1 = x_2 \text{ and } y_1 = y_2, \\ l(l-1)\alpha\beta \geqslant (q-3)(q-4)\alpha\beta & \text{if } x_1 \neq x_2 \text{ and } y_1 = y_2, \\ l(l-1)\alpha\beta \geqslant (q-3)(q-4)\alpha\beta & \text{if } x_1 = x_2 \text{ and } y_1 \neq y_2, \\ l(l-1)\alpha^2 + l\beta^2 \geqslant (q-4)(q-5)\alpha^2 + (q-4)\beta^2 & \text{if } x_1 \neq x_2 \text{ and } y_1 \neq y_2. \end{cases} \\ &\geqslant \frac{(q-5)(q-4)}{(q-3)^2(q-2)(q-1)} \end{aligned}$$

So $\rho(T) \leqslant \sqrt{\lambda_1(T^2)} \leqslant \sqrt{1 - \frac{(q-5)(q-4)q^2}{(q-3)^2(q-2)(q-1)}} \leqslant \frac{3}{q} + \frac{8}{q^2} \leqslant \frac{4}{q-1}$ for $q \geqslant 6$. □

3.3 q-ary Functions, Influences, Noise stability

We define inner product on space of functions from $[q]^N$ to $\mathbb{R}$ as $\langle f, g \rangle = \mathbb{E}_{x \sim [q]^N} [f(x) g(x)]$. Here $x \sim \mathcal{D}$ denotes sampling from distribution $\mathcal{D}$ and $\mathcal{D} = [q]^N$ denotes the uniform distribution on $[q]^N$.

Given a symmetric Markov operator T and $x = (x_1, \ldots, x_N) \in [q]^N$, let $T^{\otimes N} x$ denote the product distribution on $[q]^N$ whose i^{th} entry y_i is distributed according to $T(x_i \leftrightarrow y_i)$. Therefore $T^{\otimes N} f(x) = \mathbb{E}_{y \sim T^{\otimes N} x} [f(y)]$.

Definition 6. *Let $\alpha_0, \alpha_1, \ldots, \alpha_{q-1}$ be an orthonormal basis of $\mathbb{R}^q$ such that α_0 is all constant vector. For $x \in [q]^N$, we define $\alpha_x \in \mathbb{R}^{q^N}$ as*

$$\alpha_x = \alpha_{x_1} \otimes \ldots \otimes \alpha_{x_N}.$$

Definition 7 (Fourier coefficients)**.** *For a function $f : [q]^N \to \mathbb{R}$, define $\hat{f}(\alpha_x) = \langle f, \alpha_x \rangle$.*

Definition 8. *Let $f : [q]^N \to \mathbb{R}$ be a function. The* influence *of i^{th} variable on f, $\mathsf{Inf}_i(f)$ is defined by*

$$\mathsf{Inf}_i(f) = \mathbb{E}\left[\mathsf{Var}\left[f(x)|x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_N\right]\right]$$

where $x_1, \ldots, x_N$ are uniformly distributed. Equivalently, $\mathsf{Inf}_i(f) = \sum_{x:x_i \neq 0} \hat{f}^2(\alpha_x)$.

Definition 9. *Let $f : [q]^N \to \mathbb{R}$ be a function. The* low-level influence *of i^{th} variable of f is defined by*

$$\mathsf{Inf}_i^{\leqslant t}(f) = \sum_{x:x_i \neq 0,\ |x| \leqslant t} \hat{f}^2(\alpha_x).$$

Observation 2. *For any function f, $\sum_i \mathsf{Inf}_i^{\leqslant t}(f) = \sum_{x:|x|\leqslant t} \hat{f}^2(\alpha_x)|x| \leqslant t \sum_x \hat{f}^2(\alpha_x) = t\|f\|_2^2$. If $f : [q]^N \to [0,1]$, then $\|f\|_2^2 \leqslant 1$, so $\sum_i \mathsf{Inf}_i^{\leqslant t}(f) \leqslant t$.*

Definition 10 (Noise stability)**.** *Let f be a function from $[q]^N$ to $\mathbb{R}$, and let $-1 \leqslant \rho \leqslant 1$. Define the* noise stability *of f at ρ as*

$$\mathbb{S}_\rho(f) = \langle f, T_\rho^{\otimes n} f \rangle = \sum_x \rho^{|x|} \hat{f}_i^2(\alpha_x)$$

where T_ρ is the Beckner operator as in Definition 5.

A natural way to think about a q-coloring function is as a collection of q-indicator variables summing to 1 at every point. To make this formal:

Definition 11. *Define the unit q-simplex as $\Delta_q = \{(x_1, \ldots, x_q) \in \mathbb{R}^q \mid \sum x_i = 1, x_i \geqslant 0\}$.*

Observation 3. *For positive integers Q, q and any function $f = (f_1, \ldots, f_q) : [Q]^N \to \Delta_q$, $\sum_i \mathsf{Inf}_i^{\leqslant t}(f) = \sum_i \sum_j \mathsf{Inf}_i^{\leqslant t}(f_j) \leqslant t \sum_j \|f_j\|^2 \leqslant t$.*

We want to prove a lower bound on the stability of q-ary functions with noise operators T. The following proposition is generalization of Proposition 11.4 in [7] to general symmetric Markov operators T with small spectral radii. Its proof appears in Appendix B.

Proposition 1. *For integers $Q, q \geqslant 3$, and a symmetric Markov operator T on $[Q]$ with spectral radius $\rho(T) \leqslant \frac{c}{q-1}$, for some $c > 0$, there is a small enough $\delta = \delta(q) > 0$ and $t = t(q) > 0$ such that for any function $f = (f_1, \ldots, f_q) : [Q]^N \to \Delta_q$ with $\mathsf{Inf}_i^{\leqslant t}(f) \leqslant \delta$, for all i, satisfies*

$$\sum_{j=1}^{q} \langle f_j, T^{\otimes N} f_j \rangle \geqslant 1/q - 2c \ln q / q^2 - C \ln \ln q / q^2$$

for some universal constant $C < \infty$.

Definition 12 (Moving between domains)**.** *For any $x = (x_1, \ldots, x_{2N}) \in [q]^{2N}$, denote $\overline{x} \in [q^2]^N$ as*

$$\overline{x} = ((x_1, x_2), \ldots, (x_{2N-1}, x_{2N}))\ .$$

Similarly for $y = (y_1, \ldots, y_N) \in [q^2]^N$, denote $\underline{y} \in [q]^{2N}$ as

$$\underline{y} = (y_{1,1}, y_{1,2}, \ldots, y_{N,1}, y_{N,2}),$$

where $y_i = y_{i,1} + y_{i,2} q$ such that $y_{i,1}, y_{i,2} \in [q]$. For a function f on $[q]^{2N}$, define $\overline{f}$ on $[q^2]^N$ as $\overline{f}(y) = f(\underline{y})$.

The relationship between influences of variables for functions f and $\overline{f}$ are given by the following claim (Claim 2.7 in [2]).

Claim 1. *For any function* $f : [q]^{2N} \to \mathbb{R}$*,* $i \in \{1, \ldots, N\}$ *and any* $t \geqslant 1$*,* $\mathsf{Inf}_i^{\leqslant t}(\overline{f}) \leqslant \mathsf{Inf}_{2i-1}^{\leqslant 2t}(f) + \mathsf{Inf}_{2i}^{\leqslant 2t}(f)$*.*

3.4 PCP Verifier for Max k-Colorable Subgraph

This verifier uses ideas similar to the Max k-Cut verifier given in [7] and the 4-coloring hardness reduction in [2]. Let $\mathcal{L} = (U, V, E, R, 2R, \Pi)$ be a 2-to-1 bipartite, unweighted and left regular Label-Cover instance as in Conjecture 1. Assume the proof is given as the Long Code over $[k]^{2R}$ of the label of every vertex $v \in V$. Below for a permutation σ on $\{1, \ldots, n\}$ and a vector $x \in \mathbb{R}^n$, $x \circ \sigma$ denotes $(x_{\sigma(1)}, x_{\sigma(2)}, \cdots, x_{\sigma(n)})$. For a function f on $\mathbb{R}^n$, $f \circ \sigma$ is defined as $f \circ \sigma(x) = f(x \circ \sigma)$.

- Pick u uniformly at random from U, $u \sim U$.
- Pick v, v' uniformly at random from u's neighbors. Let π, π' be the associated projection functions, $\chi_v, \chi_{v'}$ be the (supposed) Long Codes for the labels of v, v' respectively.
- Let T be the Markov operator on $[k]^2$ given in Lemma 7. Pick $x \sim [k^2]^R$ and $y \sim T^{\otimes R}x$. Let $\sigma_v, \sigma_{v'}$ be two permutations of $\{1, \ldots, 2R\}$ such that $\pi(\sigma_v^{-1}(2i-1)) = \pi(\sigma_v^{-1}(2i)) = \pi'(\sigma_{v'}^{-1}(2i-1)) = \pi'(\sigma_{v'}^{-1}(2i))$ (both π and π' are exactly 2-to-1, so such permutations exist).
- Accept iff $\chi_v \circ \sigma_v(\underline{x})$ and $\chi_{v'} \circ \sigma_{v'}(\underline{y})$ are different.

The proofs of the following two lemmas appear in Appendix C.

Lemma 8 (Completeness)**.** *If the original* 2-to-1 *Label-Cover instance* $\mathcal{L}$ *has a labeling which satisfies all constraints, then there is a proof which makes the above verifier always accept.*

Lemma 9 (Soundness)**.** *There is a constant* C *such that, if the above verifier passes with probability exceeding* $1 - 1/k + O(\ln k/k^2)$*, then there is a labeling of* $\mathcal{L}$ *which satisfies* $\gamma' = \gamma'(k)$ *fraction of the constraints independent of label set size* R*.*

Note that our PCP verifier makes "k-coloring" tests. By the standard conversion from PCP verifiers to CSP hardness, and Remark 2 about conversion to unweighted graphs with the same inapproximability factor, we conclude the main result of this section by combining Lemmas 8 and 9.

Theorem 10. *For any constant* $k \geqslant 3$*, assuming* 2-to-1 *Conjecture, it is NP-hard to approximate* Max k-Colorable Subgraph *within a factor of* $1 - 1/k + O(\ln k/k^2)$*.*

References

[1] P. Crescenzi, R. Silvestri, and L. Trevisan. On weighted vs unweighted versions of combinatorial optimization problems. *Information and Computation*, 167(1):10–26, 2001.

[2] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 344–353, 2006.

[3] A. M. Frieze and M. Jerrum. Improved approximation algorithms for max k-cut and max bisection. *Algorithmica*, 18(1):67–81, 1997.

[4] V. Guruswami, D. Lewin, M. Sudan, and L. Trevisan. A tight characterization of NP with 3 query PCPs. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages 8–17, 1998.

[5] V. Kann, S. Khanna, J. Lagergren, and A. Panconesi. On the hardness of approximating max k-cut and its dual. *Chicago J. Theor. Comput. Sci.*, 1997, 1997.

[6] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 767–775, 2002.

[7] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007.

[8] E. Mossel, R. O'Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 21–30, 2005.

[9] R. O'Donnell and Y. Wu. Conditional hardness for satisfiable CSPs. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 493–502, 2009.

[10] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991.

[11] E. Petrank. The hardness of approximation: Gap location. *Computational Complexity*, 4:133–157, 1994.

[12] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th ACM Symposium on Theory of Computing*, pages 245–254, 2008.

[13] A. Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1:351–370, 1992.

[14] P. M. Weichsel. The kronecker product of graphs. *Proceedings of the American Mathematical Society*, 13(1):47–52, 1962.

A Proofs from Section 2.1

A.1 Proof of Lemma 2

Proof. We define the coloring $\chi : V(G) \to [3]$ in the obvious way, with nodes T, R and F fixed to different colors. Then define

$$\chi(x_i) = \begin{cases} \chi(T) & \text{if } \sigma(x_i) = 1, \\ \chi(F) & \text{else.} \end{cases}$$

and similarly for the nodes y_j, z_l. Define

$$\chi(\overline{y_i}) = \begin{cases} \chi(F) & \text{if } \sigma(y_j) = 1, \\ \chi(T) & \text{else.} \end{cases}$$

Now, for the constraints satisfied by this assignment, $(x_i \vee (Y_j = z_k)) \wedge (\overline{x_i} \vee (Y_j = z_l))$, consider the corresponding gadget. Let $\mathsf{Sugg}(A) = [3] \setminus \{\chi(x_i), \chi(T)\}$ and $\mathsf{Sugg}(B) = [3] \setminus \{\chi(Y_j), \chi(z_k)\}$ be the available colors to A and B which can properly color all edges incident to variables. Notice that none of these sets are empty and since $x_i \vee (Y_j = z_k)$ is true, at least one of these sets $\mathsf{Sugg}(A)$ and $\mathsf{Sugg}(B)$ has two elements in it. Hence there exists a coloring of A and B from sets $\mathsf{Sugg}(A)$ and $\mathsf{Sugg}(B)$ such that $\chi(A) \neq \chi(B)$. The same argument also holds for A' and B', therefore all edges in this gadget are properly colored.

For the violated constraints, either $\mathsf{Sugg}(A)$ or $\mathsf{Sugg}(A')$ has one element. Augmenting that set with the color $\chi(x_i)$ will cause only one edge to be violated. □

A.2 Proof of Lemma 3

Proof. Since $\tau < m/2$, the coloring χ must give three different colors to the nodes T, F, and R. If $\chi(x_i) = \chi(R)$, then randomly choosing $\chi(x_i)$ from $\{\chi(T), \chi(F)\}$ will, in expectation, make at most half of the local gadget edges going out of x_i improperly colored, which is exactly the value $\Delta(x_i)/2$ gained. So we can assume that $\chi(x_i) \in \{\chi(T), \chi(F)\}$ for each x_i. A similar argument holds for the nodes z_l. Now consider the nodes y_j and $\overline{y_j}$ for a variable in Y. If $\chi(y_j) = \chi(R)$, $\chi(\overline{y_j}) = \chi(R)$ or $\chi(x_j) = \chi(\overline{y_j})$, then randomly choosing $(\chi(y_j), \chi(\overline{y_j}))$ from $\{(\chi(T), \chi(F)), (\chi(F), \chi(T))\}$ will, in expectation, make at most half of the local gadget edges going out of nodes y_j and $\overline{y_j}$ improperly colored, which is exactly the value w_j gained.

To summarize, we can assume that nodes T,F and R are colored differently, $\chi(x_i), \chi(Y_j), \chi(z_l) \in \{\chi(T), \chi(F)\}$ and $\chi(y_j) \neq \chi(\overline{y_j})$. Thus all edges other than the edges inside the local gadgets are properly colored by χ, and by assumption at most τ edges are miscolored by χ.

Now define the natural assignment σ' that assigns a variable of $\mathcal{V}$ the value 1 if the associated variable received the color $\chi(T)$, and the value 0 if its color is $\chi(F)$.

Consider a local gadget, with all edges properly colored, corresponding to the constraint $(x_i \vee (Y_j = z_k)) \wedge (\overline{x_i} \vee (Y_j = z_l))$. Assume $\sigma'(x_i) = 0$, which implies $\chi(A) = \chi(R)$. Then both neighbors of B besides A must have the same color, therefore $\sigma(Y_j) = \sigma(z_k)$. The other case when $\sigma'(x_i) = 1$ is similar. Hence the assignment σ' will satisfy this constraint.

Since the local gadgets corresponding to different constraints have disjoint sets of edges, it follows that the number of constraints violated by the assignment σ' is at most τ. □

B Proof of Proposition 1

Proof. Let $t = 4$, $f_i : [Q]^N \to [0, 1]$ denote the i^{th} coordinate function of f, and let $\mu_i = \mathbb{E}[f_i]$. Let $\alpha_0, \ldots, \alpha_{Q-1}$ be an orthonormal set of eigenvectors for T with corresponding eigenvalues $\lambda_0 \geqslant \ldots \geqslant \lambda_{Q-1}$, with $\rho = \rho(T) \leqslant \frac{c}{q-1}$ being the spectral radius of T. Notice that T is symmetric so $\lambda_0 = 1$ and α_0 is a constant vector. Therefore $\mathbb{E}[f_i] = \hat{f}_i(\alpha_0) = \mu_i$. Then (using the notation from [2]):

$$T^{\otimes N} \alpha_x = (\prod_{a \neq 0} \lambda_a^{|x|_a}) \alpha_x$$

and hence

$$T^{\otimes N} f_i = \sum_x (\prod_{a \neq 0} \lambda_a^{|x|_a}) \hat{f}_i(\alpha_x) \alpha_x.$$

At this point, consider the Beckner operator, T_ρ on $[Q]$. Since α_0 is the uniform distribution, it is a constant vector, thus $\alpha_0, \alpha_1, \ldots, \alpha_{Q-1}$ is also an orthonormal basis for T_ρ. Consequently,

$$\langle f_i, T_\rho^{\otimes N} f_i \rangle = \sum_x (\prod_{a \neq 0} \rho^{|x|_a}) \hat{f}_i^2(\alpha_x) = \sum_x \rho^{|x|} \hat{f}_i^2(\alpha_x) = \mathbb{S}_\rho(f_i)$$

Thus

$$\begin{aligned} \langle f_i, T^{\otimes N} f_i \rangle &= \hat{f}_i^2(\alpha_0) - \hat{f}_i^2(\alpha_0) + \sum_x \underbrace{(\prod_{a \neq 0} \lambda_a^{|x|_a})}_{\begin{cases} \geqslant -\rho^{|x|} & \text{if } |x| \neq 0, \\ = 1 & \text{else.} \end{cases}} \hat{f}_i^2(\alpha_x) \\ &\geqslant 2\mu_i^2 - \sum_x \rho^{|x|} \hat{f}_i^2(\alpha_x) = 2\mu_i^2 - \sum_{x:|x| \leqslant 4} \rho^{|x|} \hat{f}_i^2(\alpha_x) - \sum_{x:|x|>4} \rho^{|x|} \hat{f}_i^2(\alpha_x) \\ &\geqslant 2\mu_i^2 - \sum_{x:|x| \leqslant 4} \rho^{|x|} \hat{f}_i^2(\alpha_x) - \rho^4 \\ &\geqslant 2\mu_i^2 - \sum_{x:|x| \leqslant 4} \rho^{|x|} \hat{f}_i^2(\alpha_x) - q^{-3} \end{aligned}$$

At this point, let $\tilde{f}_i(x) = \sum_{x:|x| \leqslant 4} (\prod_{a \neq 0} \lambda_a^{|x|_a}) \hat{f}_i(\alpha_x) \alpha_x$ be the function having the same low-level coefficients with $f_i(x)$ and 0 for the higher-levels. It is easy to verify that $\mathbb{E}\left[\tilde{f}_i\right] = \mu_i$, $\mathsf{Inf}_i(f_j) \geqslant \mathsf{Inf}_i(\tilde{f}_j) = \mathsf{Inf}_i^{\leqslant 4}(f_j)$ and $\mathbb{S}_\rho(\tilde{f}_j) = \sum_{x:|x| \leqslant 4} \rho^{|x|} \hat{f}_i^2(\alpha_x)$. In particular, our assumption $\sum_j \mathsf{Inf}_i^{\leqslant t}(f_j) = \sum_j \mathsf{Inf}_i^{\leqslant 4}(f_j) \leqslant \delta$ implies $\sum_j \mathsf{Inf}_i(\tilde{f}_j) \leqslant \delta$.

Let δ be a small enough constant such that $\mathbb{S}_{\frac{c}{q-1}}(\tilde{f}_i) \leqslant \Gamma_{\frac{c}{q-1}}(\mu_i) + \varepsilon$ for some small $\varepsilon \leqslant \frac{1}{q^3}$, from the Majority is Stablest Theorem [8]. In [7], $\Lambda_\eta(\mu)$ is used for $\Gamma_\eta(\mu)$ and we will follow that convention instead. Below, for a real x, $[x]^+$ denotes $\max\{x, 0\}$. Then

$$\begin{aligned} \sum_i \langle f_i, T^{\otimes N} f_i \rangle &\geqslant \sum_i \left[2\mu_i^2 - \mathbb{S}_\rho(\tilde{f}_i)\right] - q^{-2} \\ &\geqslant \sum_i \left[2\mu_i^2 - \mathbb{S}_{\frac{c}{q-1}}(\tilde{f}_i)\right] - q^{-2} \\ &\geqslant \sum_i \left[2\mu_i^2 - \Lambda_{\frac{c}{q-1}}(\mu_i)\right]^+ - 2q^{-2} \\ &\geqslant \frac{1}{q} - \frac{2c \ln q}{q^2} - O\left(\frac{\ln \ln q}{q^2}\right) \end{aligned}$$

The last inequality is proved in the same way as Proposition 11.4 in [7]. The only difference is that we have

$$F(\mu_i) = \mu_i^2 + \frac{c}{q-1} 2\mu_i^2 \ln(1/\mu_i) \cdot \left(1 + C \frac{\ln \ln q}{\ln q}\right)$$

and

$$\sum_{i=1}^q \left[2\mu_i^2 - \Lambda_{\frac{c}{q-1}}(\mu_i)\right]^+ \geqslant \sum_{i=1}^q (2\mu_i^2 - F(\mu_i))$$

which is convex because $\mu_i \leqslant (1/q)^{1/10}$ and minimized at $\mu_i = 1/q$. In this case, we have

$$\sum_{i=1}^{q} (2\mu_i^2 - F(\mu_i)) \geqslant q\left(q^{-2} - 2cq^{-3}\ln q(1 + C\ln\ln q/\ln q\right)$$

from which the above claim follows. ◻

C Analysis of PCP verifier for Max k-Colorable Subgraph

C.1 Proof of Lemma 8

Proof. Let $\ell : V \to \{1, \ldots, 2R\}$ be a labeling for $\mathcal{L}$ satisfying all constraints in Π. Pick χ_v as the Long Code encoding of $\ell(v)$. Given any pair of vertices $v, v' \in V$ which share a common neighbor $u \in U$, and $x, y \in [k]^{2R}$ pairs such that

$$\Pr\left[\overline{y} \sim T^{\otimes R}(\overline{x})\right] = \prod_i T((x_{2i-1}, x_{2i}) \leftrightarrow (y_{2i-1}, y_{2i})) > 0 ,$$

let π, π' be the projection functions and $\sigma_v, \sigma_{v'}$ be the permutations as defined in the description of the verifier. We have $\chi_v(x \circ \sigma_v) = x_{\sigma(\ell(v))}$ and $\chi_{v'}(y \circ \sigma_{v'}) = y_{\sigma'(\ell(v'))}$. Since $\pi(\ell(v)) = \pi'(\ell(v'))$, this implies $\sigma_v(\ell(v)), \sigma_{v'}(\ell(v')) \in \{2i-1, 2i\}$ for some $i \leqslant R$. But

$$T((x_{2i-1}, x_{2i}) \leftrightarrow (y_{2i-1}, y_{2i})) > 0 \implies \{x_{2i-1}, x_{2i}\} \cap \{y_{2i-1}, y_{2i}\} = \emptyset ,$$

therefore $\chi_v \circ \sigma_v(x) = x_{\sigma_v(\ell(v))} \neq y_{\sigma_{v'}(\ell(v'))} = \chi_{v'} \circ \sigma_{v'}(y)$. So the verifier always accepts. ◻

C.2 Proof of Lemma 8

Proof. For each node $v \in V$, let $f^v : [k]^{2R} \to \Delta_k$ be the function $f^v(x) = e_{\chi_v(x)}$ where e_i is the indicator vector of the i^{th} coordinate. Let $\Gamma(u)$ denote the set of vertices adjacent to u in the Label Cover graph.

After arithmetizing, we can write the verifier's acceptance probability as

$$\begin{aligned}
\Pr[\text{acc}] &= \mathbb{E}_{u,v,v'}\left[1 - \textstyle\sum_j \langle \overline{f_j^v \circ \sigma_v}, T^{\otimes R}\overline{(f_j^{v'} \circ \sigma_{v'})}\rangle\right] \\
&= 1 - \mathbb{E}_u\left[\textstyle\sum_j \mathbb{E}_{v,v'}\left[\langle \overline{f_j^v \circ \sigma_v}, T^{\otimes R}\overline{(f_j^{v'} \circ \sigma_{v'})}\rangle\right]\right] \\
&= 1 - \mathbb{E}_u\left[\textstyle\sum_j \langle \mathbb{E}_v\left[\overline{f_j^v \circ \sigma_v}\right], T^{\otimes R}\mathbb{E}_{v'}\left[\overline{f_j^{v'} \circ \sigma_{v'}}\right]\rangle\right] \\
&= 1 - \mathbb{E}_u\left[\textstyle\sum_j \langle g_j^u, T^{\otimes R} g_j^u\rangle\right] \qquad \left(g_j^u = \mathbb{E}_{v\sim\Gamma(u)}\left[\overline{f_j^v \circ \sigma_v}\right]\right) \\
&\geqslant 1 - 1/k + C\ln k/k^2
\end{aligned}$$

where $g^u : [k^2]^R \to \Delta_k$ and some constant C. By averaging, for at least a fraction $\delta = (\varepsilon/2)\ln k/k^2$ of vertices in U, we have

$$\sum_j \langle g_j^u, T^{\otimes R} g_j^u\rangle \leqslant 1/k - C\ln k/k^2$$

Let these be "good" vertices. For a good vertex, by Proposition 1, there exist constants $\delta = \delta(k)$, $t = t(k)$ and i such that $\mathsf{Inf}_i^{\leqslant t}(g^u) \geqslant \delta$. Let $\mathsf{Sugg}_u = \{i | i \in \{1, \ldots, R\} \wedge \mathsf{Inf}_i^{\leqslant t}(g^u) \geqslant \delta\}$, so $|\mathsf{Sugg}_u| \geqslant 1$. By Observation 3, $|\mathsf{Sugg}_u| \leqslant t/\delta$. For a good vertex u, and $j \in \mathsf{Sugg}_u$:

$$\delta \leqslant \mathsf{Inf}_j^{\leqslant t}(g^u) = \mathbb{E}_{v\sim\Gamma(u)}\left[\mathsf{Inf}_j^{\leqslant t}(\overline{f^v \circ \sigma_v})\right]$$

Therefore, for at least a fraction $\delta/2$ of neighbors v of u, $\mathsf{Inf}_j^{\leqslant t}(\overline{f^v \circ \sigma_v}) \geqslant \delta/2$. For such v and j, by Claim 1, $\mathsf{Inf}_{2j-1}^{\leqslant 2t}(f^v \circ \sigma_v) + \mathsf{Inf}_{2j}^{\leqslant 2t}(f^v \circ \sigma_v) \geqslant \delta/2$. Therefore for some $j \in [2R]$, $\mathsf{Inf}_j^{\leqslant 2t}(f^v) \geqslant \delta/4$. Let $\mathsf{Sugg}_v = \{j | j \in \{1, \ldots, 2R\} \wedge \mathsf{Inf}_j^{\leqslant 2t}(f^v) \geqslant \delta/4\}$. Again, Sugg_v is not empty and $|\mathsf{Sugg}_v| \leqslant 8t/\delta$.

Following the decoding procedure in [7], we deduce that it is possible to satisfy a fraction $\gamma' = \gamma'(\delta, t) = \gamma'(k)$ of the constraints. □

D Handling k not divisible by 3 in Theorem 4

We now argue how to handle the case when $k \mod 3 \neq 0$ in the statement of Theorem 4. Assume k is of the form $K + L$, where $K \equiv 0 \pmod 3$ and $L = k \mod 3 \in \{1, 2\}$. We will give a reduction from Max K-Colorable Subgraph, which we already showed to be NP-hard to approximate within a factor $1 - \frac{1}{33K} + \varepsilon$, to Max k-Colorable Subgraph.

Let G_K be an (unweighted) instance of Max K-Colorable Subgraph with M edges. Construct a graph H by adding L new vertices $u_1, \ldots, u_L$ to G_K. Each u_i is connected by an edge of weight $\frac{d_v}{K}$ to each vertex $v \in V(G_K)$, where d_v is the degree of v in G_K. If $L > 1$, (u_1, u_2) is an edge in H with weight $\frac{M}{33K}$. The total weight of edges in H equals

$$M' = M + \frac{2LM}{K} + \frac{M(L-1)}{33K} .$$

Clearly if G_K is K-colorable, then H is k-colorable. For the soundness part, suppose every K-coloring of G_K miscolors at least $\left(\frac{1}{33K} - \varepsilon\right)M$ edges. Let χ be an optimal k-coloring of H. We will prove that χ miscolors edges with total weight at least $M(\frac{1}{33K} - \varepsilon)$. This will certainly be the case if $L > 1$ and $\chi(u_1) = \chi(u_2)$. So we can assume χ uses L colors for the newly added vertices u_i. If $\chi(v) = \chi(u_i)$ for some $v \in V(G_K)$, we can change $\chi(v)$ to one of the K colors not used to color $\{u_1, \ldots, u_L\}$ so that the weight of miscolored edges does not increase. Therefore, we can assume that χ uses only K colors to color the G_K portion of H. But this implies at least $M(\frac{1}{33K} - \varepsilon)$ edges are miscolored by χ, as desired.

Thus every k-coloring of H miscolors at least a fraction

$$\frac{M(1/(33K) - \varepsilon)}{M'} = \frac{(1/(33K) - \varepsilon)}{1 + 2L/K + (L-1)/(33K)} \geqslant \frac{1}{33(k+L) + (L-1)} - \varepsilon$$

of the total weight of edges in H. Since $L = k \mod 3$, the bound stated in Theorem 4 holds.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de