

Improved Inapproximability Results for Maximum k -Colorable Subgraph*

Venkatesan Guruswami[†]

Ali Kemal Sinop[‡]

May 12, 2013

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 $1 - O(\frac{1}{k^2})$ 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 $1 - \frac{1}{k} + O(\frac{\ln k}{k^2})$ of edges of a k -colorable graph.

*An extended abstract of this paper appeared in the Proceedings of the 12th International Workshop on Approximation, Randomization, and Combinatorial Optimization, 2009 [7].

[†]Research supported in part by a Packard Fellowship

[‡]Research supported in part by a Packard Fellowship

ACM Classification: F.2.2, G.1.6, G.2.2

AMS Classification: 68W17, 68R10

Key words and phrases: Hardness of approximation, Graph coloring, Constraint Satisfaction, Probabilistically checkable proof,

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 $\chi : V \rightarrow \{1, 2, \dots, k\}$ such that for all edges $(u, v) \in E$, $\chi(u) \neq \chi(v)$. For $k \geq 3$, 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 $k \geq 3$), find a k -coloring that maximizes the fraction of properly colored edges. Throughout this paper, for a given coloring, we say an edge is miscolored by this coloring if both of its endpoints receive the same color and properly colored otherwise. If no such edge exists, we call the coloring proper. 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, 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 $P = NP$. 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 [4] 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 approximation guarantee for Max k -Colorable Subgraph known to date. Khot, Kindler, Mossel, and O’Donnell [11] 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 [15] with a reduction from Papadimitriou and Yannakakis [14]. 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.

For Max k -Cut, an NP-hardness factor was shown by Kann, Khanna, Lagergren, and Panconesi [9] — for some absolute constants $\beta > \alpha > 0$: They showed that it is NP-hard to distinguish graphs that have a

¹While 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.

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. This hardness result was proven by reduction from MaxCut. 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

Petrant [15] 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 [15] was left unspecified and would be very small if calculated. The reduction in [15] was rather complicated, involving expander graphs and starting from the weaker hardness bounds for bounded occurrence satisfiability (as opposed to, say, problems for which deciding between perfect completeness and $\frac{1}{2}$ soundness is hard). 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 [6]. Note that we are interested in making the constant γ_0 as large as possible, which is the reason why we went through a different reduction than simply re-using Petrank's hardness result [15].

By a reduction from Max 3-Colorable Subgraph, we prove that for every $k \geq 3$, the Max k -Colorable Subgraph is NP-hard to approximate within a factor greater than $\approx 1 - \frac{1}{33k}$ (Theorem 2.9). 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 [9], 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. Furthermore the constant in front of $\frac{1}{k}$ is the same with constant γ_0 of 3-coloring case. Any improvement for 3-coloring hardness will immediately apply to general k -coloring case as well.

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 (see Definition 3.4, Section 3.1), 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 [3] to prove that for every constant $c \geq 4$, 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 [3], which in turn extends the invariance principle of [12]. The hardness factor we obtain depends on the spectral gap of a certain $k^2 \times k^2$ stochastic matrix.

Remark 1.1. 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 [13] who showed a tight hardness for approximating satisfiable constraint satisfaction problems on three 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 [16], the approximability of *satisfiable* constraint satisfaction problems remains a mystery to understand in any generality.

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 [15] 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 and uses expander graphs. We prove a much better inapproximability factor below, via a simpler proof.

Theorem 2.1 (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 every constant $\varepsilon > 0$.*

Remark 2.2. In a recent work, Austrin *et al.* [1] improved the hardness of approximation factor in our Theorem 2.1 to $\frac{16}{17} + \varepsilon$.

For the proof of this theorem, we will reduce from a hard to approximate constraint satisfaction problem (CSP), stated below. This CSP is related to the adaptive 3-query PCP (with perfect completeness and soundness $1/2 + \varepsilon$ for any desired constant ε) given in [6], which in turn is based on Håstad's PCP for showing tight inapproximability of Max 3SAT on satisfiable instances.

Definition 2.3 (The GLST CSP). An instance of the GLST constraint satisfaction problem will have variables partitioned into three parts \mathcal{X}, \mathcal{Y} and \mathcal{Z} . Each constraint will be of the form $(x_i \vee (Y_j = z_k)) \wedge (\bar{x}_i \vee (Y_j = z_l))$, where $x_i \in \mathcal{X}$, $z_k, z_l \in \mathcal{Z}$ are variables (unnegated) and Y_j is a literal ($Y_j \in \{y_j, \bar{y}_j\}$ for some variable $y_j \in \mathcal{Y}$). The goal is to find a Boolean assignment to the variables in $\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$ that satisfies as many constraints as possible.

The following hardness result will be our starting point in the reduction establishing Theorem 2.1.

Proposition 2.4. *For all constants $\varepsilon > 0$, given an instance of the GLST CSP, it is NP-hard to distinguish between the following two cases:*

- **YES instances:** *There is a Boolean assignment to the variables that satisfies **all** the constraints.*
- **NO instances:** *Every Boolean assignment to the variables satisfies at most a fraction $(1/2 + \varepsilon)$ of the constraints.*

Remark 2.5. Without the restriction that the instance is tripartite (i.e., the x , Y , and z variables come from disjoint parts) and that the z -variables never appear negated, the above hardness result follows immediately from [6]. However, these extra structural restrictions can be ensured by an easy modification to the PCP construction in [6]. The PCP in [6] 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 $(A(f) \vee (B(g) = B(g_1)) \wedge (\bar{A}(f) \vee (B(g) = B(g_2)))$. Further these tables are *folded*, meaning that $A(f) = -A(-f)$ and $B(g) = -B(-g)$ are enforced by construction, and this leads to the occurrence of negations when the PCP checks are viewed as constraints in the CSP world. In fact, as argued in [5], one does not need the B -table to be folded if one makes a slight change to the predicate and checks the condition $(A(f) \vee (\bar{B}(g) = B(-g_1)) \wedge (\bar{A}(f) \vee (\bar{B}(g) = B(-g_2)))$. Not folding the B -table gives further control to how negations appear in the constraints.

The proof of Proposition 2.4 is based on the observation that the analysis of the PCP construction in [6] goes through with minor changes, even if (i) the queries at locations g_1 and g_2 are made in a parallel C -table instead of also being made in the B table, and (ii) the C -table is not required to be folded (though the A and B tables are still folded).² We skip a full formal proof of this as it will take us too far afield into the inner Fourier-analytic workings of the analysis of [6]. However, in Appendix A, we indicate the main changes needed in the analysis when a C -table is used. This should suffice to convince those generally familiar with Håstad's work [8] about the claim of Proposition 2.4.




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 $(x_i \vee Y_j = z_k) \wedge (\bar{x}_i \vee Y_j = z_l)$. All edges have unit weight. Labels A, A', B, B' refer to the local nodes in each gadget.

Proof. Let \mathcal{I} be an instance of the GLST CSP with m constraints of the above form on variables $\mathcal{V} = \mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$. Let $\mathcal{X} = \{x_1, x_2, \dots, x_{n_1}\}$, $\mathcal{Y} = \{y_1, y_2, \dots, y_{n_2}\}$ and $\mathcal{Z} = \{z_1, z_2, \dots, z_{n_3}\}$. 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 $\{y_j, \bar{y}_j\}$ 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, \bar{y}_j are connected to node R by a triangle whose edge weights equal $w_j = \frac{\Delta(y_j) + \Delta(\bar{y}_j)}{2}$. Here $\Delta(X)$ denotes the total number of edges going from node X into all the local gadgets. Note that $\Delta(x_i)$ equals twice the total number of occurrences of x_i in all constraints, and similarly for $\Delta(y_j)$ and $\Delta(\bar{y}_j)$. For the z -variables, $\Delta(z_l)$ equals the number of occurrences of z_l in constraints of the CSP. 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.

Remark 2.6. It has been shown by Crescenzi, Silvestri and Trevisan [2] 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.

²Note that these restrictions in the PCP world exactly correspond to the structural restrictions of the CSP instance in Definition 2.3.

Let us now prove the completeness and soundness properties of the reduction.

Lemma 2.7 (Completeness). *Given an assignment of variables $\sigma : \mathcal{V} \rightarrow \{0, 1\}$ 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).*

Proof. We define the coloring $\chi : V(G) \rightarrow [3]$ 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

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

Now, for the constraints satisfied by this assignment, $(x_i \vee (Y_j = z_k)) \wedge (\bar{x}_i \vee (Y_j = z_l))$, consider the corresponding gadget. Let $\text{Sugg}(A) = [3] \setminus \{\chi(x_i), \chi(T)\}$ and $\text{Sugg}(B) = [3] \setminus \{\chi(Y_j), \chi(z_k)\}$ 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 $\text{Sugg}(A)$ and $\text{Sugg}(B)$ has two elements in it. Hence there exists a coloring of A and B from sets $\text{Sugg}(A)$ and $\text{Sugg}(B)$ 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 $\text{Sugg}(A)$ or $\text{Sugg}(A')$ has one element. Augmenting that set with the color $\chi(x_i)$ will cause only one edge to be violated. \square

Lemma 2.8 (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 $\sigma' : \mathcal{V} \rightarrow \{0, 1\}$ to the variables of the CSP instance that satisfies at least $m - \tau$ constraints.*

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 \bar{y}_j for a variable in Y . If $\chi(y_j) = \chi(R)$, $\chi(\bar{y}_j) = \chi(R)$ or $\chi(y_j) = \chi(\bar{y}_j)$, then randomly choosing $(\chi(y_j), \chi(\bar{y}_j))$ 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 \bar{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 $\chi(y_j) \neq \chi(\bar{y}_j)$. 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 $(x_i \vee (Y_j = z_k)) \wedge (\bar{x}_i \vee (Y_j = z_l))$. 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 . \square

Returning to the proof of Theorem 2.1, the total weight of edges in G is

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

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 2.1 is complete.³ \square

2.2 Max k -Colorable Subgraph Hardness

Theorem 2.9. *For every fixed integer $k \geq 3$ 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 $c_k = k \bmod 3 \leq 2$.*

Remark 2.10. The hardness factor in Theorem 2.9 can be improved to $1 - \frac{1}{17(k+c_k)+c_k} + \varepsilon$ using the earlier mentioned improved hardness of approximation result for Max 3-Colorable Subgraph due to Austrin *et al.* [1].

Proof. We will reduce Max 3-Colorable Subgraph to Max k -Colorable Subgraph and then apply Theorem 2.1. 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.* [9] (see Remark 2.13). Some modifications to the reduction are needed when we reduce from Max 3-Colorable 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 2.1, it is NP-hard to tell if G is 3-colorable or every 3-coloring 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 2.1, this gives us the claimed hardness of Max k -Colorable Subgraph.

Let $K'_{k/3}$ denote the complete graph including 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 the following: 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 $i, i' \in \{1, \dots, k/3\}$.

We construct another graph H by making three copies of G' , and identifying the nodes with $(u, i, j), (u, i) \in V(G'), j \in \{1, 2, 3\}$. The edge set of this new graph is defined in the following way: There are edges between all nodes of the form (u, i, j) and (u, i', j') with weight $\frac{2}{3}d_u$ if either $i \neq i'$ or $j \neq j'$. Here d_u is degree of node u .

The total weight of edges in this new graph, H , is equal to

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

Lemma 2.11. *If G is 3-colorable, then H is k -colorable.*

³Our reduction produced a graph with edge weights, but by Remark 2.6, the same inapproximability factor holds for unweighted graphs as well.

Proof. Let $\chi_G : V(G) \rightarrow \{1, 2, 3\}$ be a 3-coloring of G . Consider the following coloring function for H , $\chi_H : V(H) \rightarrow \{1, 2, \dots, k\}$. 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 $\pi^j(x) = \underbrace{\pi(\dots(\pi(x)))}_{j \text{ times}}$. Equivalently $\pi(x) = x \bmod 3 + 1$.

First let us consider edges of the form $\{(u, i, j), (v, i', j')\}$. If $i \neq i'$, then colors of the endpoints belong to different windows of size 3 and are therefore different. If $i = i'$, we have $\chi((u, i, j)) - \chi((v, i, j)) \equiv \chi(u) - \chi(v) \not\equiv 0 \bmod 3$ since χ is a proper 3-coloring of G .

Let us now consider edges of the form $\{(u, i, j), (u, i', j')\}$. Once again, if $i \neq i'$, the two endpoints receive different colors. When $i = i'$, $j \neq j'$, $\chi((u, i, j)) - \chi((u, i, j')) \equiv \pi^j(u) - \pi^{j'}(u) \equiv j - j' \not\equiv 0 \bmod 3$. \square

Lemma 2.12. *If H has a k -coloring that properly colors a set of edges with at least a fraction $(1 - \mu)$ of the total weight for some $\mu \in [0, 1]$, then G has a 3-coloring which colors at least a fraction $(1 - \mu k)$ of its edges properly.*

Proof. Here we only handle the case of k being divisible by 3. We defer the proof for the case of k being not divisible by 3 to Section 2.3.

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

$$C^{\text{total}} = \sum_{u \in V(G)} \frac{2}{3} d_u C_u^{\text{within}} + C^{\text{between}}, \quad (2.1)$$

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

$$\begin{aligned} C^{\text{between}} &= \sum_{j=1}^3 \sum_{\{u,v\} \in E(G)} \sum_{1 \leq i, i' \leq k/3} 1_{\chi_H((u, i, j)) = \chi_H((v, i', j))} \\ &\geq \sum_{j=1}^3 \sum_{\{u,v\} \in E(G)} |\text{Sugg}_u^j \cap \text{Sugg}_v^j|. \end{aligned} \quad (2.2)$$

$$\begin{aligned} C_u^{\text{within}} &= \sum_{c \in \text{Sugg}_u} \binom{|\chi_H^{-1}(c) \cap B_u|}{2} & (B_u = \{(u, i, j) \mid \forall i, j\}) \\ &= \frac{1}{2} \left(\sum_{c \in \text{Sugg}_u} |B_{u,c}|^2 \right) - \frac{k}{2} & (B_{u,c} = B_u \cap \chi_H^{-1}(c)) \\ &\geq \frac{1}{2|\text{Sugg}_u|} \left(\sum_{c \in \text{Sugg}_u} |B_{u,c}| \right)^2 - \frac{k}{2} & (\text{Cauchy-Schwarz}) \\ &= \frac{k}{2} \left(\frac{k}{|\text{Sugg}_u|} - 1 \right) = \frac{k}{2} \frac{|\text{Sugg}_u|}{|\text{Sugg}_u|} \geq \frac{|\text{Sugg}_u|}{2}. \end{aligned} \quad (2.3)$$

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

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

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

$$\begin{aligned}
\mathbb{E} \left[\sum_{\{u,v\} \in E(G)} 1_{\chi_G(u) \neq \chi_G(v)} \right] &\leq \sum_{\{u,v\} \in E} \left[\left(\sum_{j=1}^3 \frac{|\text{Sugg}_u^j \cap \text{Sugg}_v^j|}{k} \right) + \frac{|\overline{\text{Sugg}}_u|}{3k} + \frac{|\overline{\text{Sugg}}_v|}{3k} \right] \\
&\leq \frac{1}{k} \left(C^{\text{between}} + \sum_{u \in V(G)} \frac{d_u}{3} |\overline{\text{Sugg}}_u| \right) \quad (\text{using (2.2)}) \\
&\leq \frac{1}{k} \left(C^{\text{between}} + \sum_{u \in V(G)} \frac{2d_u}{3} C_u^{\text{within}} \right) = \frac{C^{\text{total}}}{k}.
\end{aligned}$$

This implies that there exists a 3-coloring of G for which the number of improperly colored edges in G is at most $\frac{C^{\text{total}}}{k}$. 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. \square

This completes the proof of Theorem 2.9 when k is divisible by 3. The other cases are easily handled by adding $k \bmod 3$ extra nodes connected to all vertices by edges of suitable weight, which we describe in the following section. \square

Remark 2.13 (Comparison to [9]). The reduction of Kann *et al* [9] 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) , $1 \leq i \leq k/2$. 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.

2.3 Handling k not divisible by 3

Proof for the case of k not divisible by 3 in Theorem 2.9. We now argue how to handle the case when $k \bmod 3 \neq 0$ in the statement of Theorem 2.9. Assume k is of the form $K + L$, where $K \equiv 0 \pmod{3}$ and $L = k \bmod 3 \in \{1, 2\}$. 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, \dots, 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 $M\left(\frac{1}{33K} - \varepsilon\right)$. 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 $\{u_1, \dots, u_L\}$ 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 $M\left(\frac{1}{33K} - \varepsilon\right)$ edges are miscolored by χ , as desired.

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

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

of the total weight of edges in H . Since $L = k \bmod 3$, the bound stated in Theorem 2.9 holds. \square

3 Conditional Hardness Results for Max k -Colorable Subgraph

We will first review the 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 the d -to-1 conjecture.

Definition 3.1. 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 $1 \leq R_U \leq R_V$. Π is a collection of projection functions for each edge: $\Pi = \{\pi_{vu} : \{1, \dots, R_V\} \rightarrow \{1, \dots, R_U\} \mid u \in U, v \in V\}$. A labeling ℓ is a mapping $\ell : U \rightarrow \{1, \dots, R_U\}$, $\ell : V \rightarrow \{1, \dots, R_V\}$. An edge $e = (u, v)$ is satisfied by the labeling ℓ if $\pi_e(\ell(v)) = \ell(u)$. We define the value of a labeling as the sum of weights of edges satisfied by this labeling normalized by the total weight. $\text{Opt}(\mathcal{L})$ is the maximum value over any labeling.

Definition 3.2. A projection $\pi : \{1, \dots, R_V\} \rightarrow \{1, \dots, R_U\}$ is called d -to-1 if for each $i \in \{1, \dots, R_U\}$, $|\pi^{-1}(i)| \leq d$. It is called *exactly* d -to-1 if $|\pi^{-1}(i)| = d$ for each $i \in \{1, 2, \dots, R_U\}$.

Definition 3.3. 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.

The following conjecture was first introduced in a seminal paper by Khot [10]. We use a slight refinement of it due to Dinur et al. [3].

Conjecture 3.4 (d -to-1 Conjecture, as stated in [3]). *For every constant $\gamma > 0$, the following decision problem is NP-hard. Given an exactly d -to-1 Label-Cover instance \mathcal{L} with $R_V = R(\gamma)$ and $R_U = dR_V$ many labels:*

- **Yes:** $\text{Opt}(\mathcal{L}) = 1$,
- **No:** $\text{Opt}(\mathcal{L}) \leq \gamma$.

Using the reductions from [3], 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 $\{0, 1, \dots, M-1\}$. 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 3.5. 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 $T(x \rightarrow y)$ is the probability of transitioning from x to y . Such an operator is called symmetric if $T(x \rightarrow y) = T(y \rightarrow x) = T(x \leftrightarrow y)$.

Definition 3.6. Given $\rho \in [-1, 1]$, the Beckner noise operator, T_ρ on $[q]$ is defined as $T_\rho(x \rightarrow x) = \frac{1}{q} + \left(1 - \frac{1}{q}\right)\rho$ and $T_\rho(x \rightarrow y) = \frac{1}{q}(1 - \rho)$ for any $x \neq y$.

Observation 3.7. All eigenvalues of the operator T_ρ are given by $1 = \lambda_0(T_\rho) \geq \lambda_1(T_\rho) = \dots = \lambda_{q-1}(T_\rho) = \rho$. Any orthonormal basis $\alpha_0, \alpha_1, \dots, \alpha_{q-1}$ with α_0 being constant vector, is also a basis for T_ρ .

Lemma 3.8. For an integer $q \geq 6$, there exists a symmetric Markov operator T on $[q]^2$ whose diagonal entries are all 0 and with eigenvalues $1 = \lambda_0 \geq \lambda_1 \geq \dots \geq \lambda_{q^2-1}$ 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 the set of eigenvalues of T^2 , $\{\lambda_i(T^2)\}_i$ is equal to $\{\lambda_i(T)^2\}_i$. Therefore $\lambda_1(T^2) \geq \max(\lambda_1^2(T), \lambda_{q^2-1}^2(T)) \geq \rho(T)^2$.

Notice that $T^2(x \leftrightarrow y) > 0$ for all pairs $x, y \in [q]^2$. Consider the variational characterization of $1 - \lambda_1(T^2)$ [17]:

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

where π corresponds to the stationary distribution for Markov operator T^2 .

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

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

So $\rho(T) \leq \sqrt{\lambda_1(T^2)} \leq \sqrt{1 - \frac{(q-5)(q-4)q^2}{(q-3)^2(q-2)(q-1)}} \leq \frac{4}{q-1}$ for $q \geq 6$. \square

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 $x = (x_1, \dots, x_N) \in [q]^N$, 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 3.9. Given $x \in [q]^N$ and $a \in [q]$, let $|x|_a$ be the number of coordinates equal to a , $a = |\{i \mid x_i = a\}|$. We will use $|x|$ to denote the number of non-zero coordinates of x , $|x| = \sum_{a \neq 0} |x|_a$.

Definition 3.10. Let $\alpha_0, \alpha_1, \dots, \alpha_{q-1}$ 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

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

Definition 3.11 (Fourier coefficients). For a function $f : [q]^N \rightarrow \mathbb{R}$, define $\hat{f}(\alpha_x) = \langle f, \alpha_x \rangle$.

Definition 3.12. Let $f : [q]^N \rightarrow \mathbb{R}$ be a function. The *influence* of i^{th} variable on f , $\text{Inf}_i(f)$ is defined by

$$\text{Inf}_i(f) = \mathbb{E} [\text{Var} [f(x) \mid x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N]]$$

where x_1, \dots, x_N are uniformly distributed.

Observation 3.13. $\text{Inf}_i(f) = \sum_{x: x_i \neq 0} \hat{f}^2(\alpha_x)$.

Definition 3.14. Let $f : [q]^N \rightarrow \mathbb{R}$ be a function. The *low-level influence* of i^{th} variable of f is defined by

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

Observation 3.15. For every function $f : [q]^N \rightarrow \mathbb{R}$,

$$\sum_i \text{Inf}_i^{\leq t}(f) = \sum_{x: |x| \leq t} \hat{f}^2(\alpha_x) |x| \leq t \sum_x \hat{f}^2(\alpha_x) = t \|f\|_2^2.$$

If $f : [q]^N \rightarrow [0, 1]$, then $\|f\|_2^2 \leq 1$, so $\sum_i \text{Inf}_i^{\leq t}(f) \leq t$.

Definition 3.16 (Noise stability). Let f be a function from $[q]^N$ to \mathbb{R} , and let $-1 \leq \rho \leq 1$. Define the *noise stability* of f at ρ as

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

where T_ρ is the Beckner operator as in Definition 3.6.

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 3.17. Define the unit q -simplex as $\Delta_q = \{(x_1, \dots, x_q) \in \mathbb{R}^q \mid \sum x_i = 1, x_i \geq 0\}$.

Observation 3.18. For positive integers Q, q and any function $f = (f_1, \dots, f_q) : [Q]^N \rightarrow \Delta_q$, $\sum_i \text{Inf}_i^{\leq t}(f) = \sum_i \sum_j \text{Inf}_i^{\leq t}(f_j) \leq t \sum_j \|f_j\|^2 \leq t$.

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 [11] to general symmetric Markov operators T with small spectral radii.

Definition 3.19. Let $\rho \in [0, 1]$ and $\mu \in [0, 1]$. Let X and Y denote normal random variables with mean 0 and covariance matrix $\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$. We define

$$\Lambda_\rho(\mu) = \text{Prob}[X \geq t \text{ and } Y \geq t].$$

where t is chosen so that $\text{Prob}[X \geq t] = \mu$.

Following is a restatement of the famous Majority is stablest theorem:

Theorem 3.20 (Majority is Stablest, [11]). *Fix $q \geq 2$ and $\rho \in [0, 1]$. Then for any $\varepsilon > 0$ there exists a small enough $\delta = \delta(\varepsilon, \rho, q) > 0$ such that if $f : [q]^n \rightarrow [0, 1]$ is any function satisfying $\mathbb{E}[f] = \mu$ and $\text{Inf}_i(f) \leq \delta$ for all $i = 1, \dots, n$ then*

$$\mathbb{S}_\rho(f) \leq \Lambda_\rho(\mu) + \varepsilon.$$

Proposition 3.21. *For integers $Q, q \geq 3$, and a symmetric Markov operator T on $[Q]$ with spectral radius $\rho(T) \leq \frac{c}{q-1}$, for some $c > 0$, there exists $t = t(q) > 0$ and a small enough $\delta = \delta(q) > 0$ such that any function $f = (f_1, \dots, f_q) : [Q]^N \rightarrow \Delta_q$ with $\text{Inf}_i^{\leq t}(f) \leq \delta$, for all i , satisfies*

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

for some universal constant $C < \infty$.

Proof. Let $t = 4$, $f_i : [Q]^N \rightarrow [0, 1]$ denote the i^{th} coordinate function of f , and let $\mu_i = \mathbb{E}[f_i]$. Let $\alpha_0, \dots, \alpha_{Q-1}$ be an orthonormal set of eigenvectors for T with corresponding eigenvalues $\lambda_0 \geq \dots \geq \lambda_{Q-1}$, with $\rho = \rho(T) \leq \frac{c}{q-1}$ 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 [3]):

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

and hence

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

At this point, consider the Beckner operator, T_ρ on $[Q]$. Since α_0 is the uniform distribution $\alpha_0, \alpha_1, \dots, \alpha_{Q-1}$ is also an orthonormal basis for T_ρ by Observation 3.7. Thus

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

At this point, let $\tilde{f}_i(x) = \sum_{x:|x| \leq 4} (\prod_{a \neq 0} \lambda_a^{|x|_a}) \hat{f}_i(\alpha_x) \alpha_x$ 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 $\mathbb{E}[\tilde{f}_i] = \mu_i$, $\text{Inf}_i(f_j) \geq \text{Inf}_i(\tilde{f}_j) = \text{Inf}_i^{\leq 4}(f_j)$ and $\mathbb{S}_\rho(\tilde{f}_j) = \sum_{x:|x| \leq 4} \rho^{|x|} \hat{f}_j^2(\alpha_x)$. In particular, our assumption $\sum_j \text{Inf}_i^{\leq t}(f_j) = \sum_j \text{Inf}_i^{\leq 4}(f_j) \leq \delta$ implies $\sum_j \text{Inf}_i(\tilde{f}_j) \leq \delta$.

Let δ be a small enough constant such that $\mathbb{S}_{\frac{c}{q-1}}(\tilde{f}_i) \leq \Lambda_{\frac{c}{q-1}}(\mu_i) + \varepsilon$ for some small $\varepsilon \leq \frac{1}{q^3}$, from the Majority is Stablest Theorem [12]. Below, for a real x , $[x]^+$ denotes $\max\{x, 0\}$. Then

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

The last inequality is proved in the same way as Proposition 11.4 in [11]. 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

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

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

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

from which the above claim follows. □

Definition 3.22 (Moving between domains). For every $x = (x_1, \dots, x_{2N}) \in [q]^{2N}$, denote $\bar{x} \in [q^2]^N$ as

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

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

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

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 \bar{f} on $[q^2]^N$ as $\bar{f}(y) = f(\underline{y})$.

The relationship between influences of variables for functions f and \bar{f} are given by the following claim (Claim 2.7 in [3]).

Claim 3.23. For every function $f: [q]^{2N} \rightarrow \mathbb{R}$, $i \in \{1, \dots, N\}$ and any $t \geq 1$, $\text{Inf}_i^{\leq t}(\bar{f}) \leq \text{Inf}_{2i-1}^{\leq 2t}(f) + \text{Inf}_{2i}^{\leq 2t}(f)$.

3.4 PCP Verifier for Max k -Colorable Subgraph

This verifier uses ideas similar to the Max k -Cut verifier given in [11] and the 4-coloring hardness reduction in [3]. 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 3.4. 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 $\{1, \dots, n\}$ and a vector $x \in \mathbb{R}^n$, $x \circ \sigma$ denotes $(x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(n)})$. 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 3.8. Pick $x \sim [k^2]^R$ and $y \sim T^{\otimes R}x$. Let $\sigma_v, \sigma_{v'}$ be two permutations of $\{1, \dots, 2R\}$ 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 $\chi_v \circ \sigma_v(x)$ and $\chi_{v'} \circ \sigma_{v'}(y)$ are different.

Lemma 3.24 (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.*

Proof. Let $\ell : V \rightarrow \{1, \dots, 2R\}$ 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

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

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 $i \leq R$. 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. \square

Lemma 3.25 (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 .*

Proof. For each node $v \in V$, let $f^v : [k]^{2R} \rightarrow \Delta_k$ 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

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

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

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

Let these be “good” vertices. For a good vertex, by Proposition 3.21, there exist constants $\delta = \delta(k)$, $t = t(k)$ and i such that $\text{Inf}_i^{\leq t}(g^u) \geq \delta$. Let $\text{Sugg}_u = \{i | i \in \{1, \dots, R\} \wedge \text{Inf}_i^{\leq t}(g^u) \geq \delta\}$, so $|\text{Sugg}_u| \geq 1$. By Observation 3.18, $|\text{Sugg}_u| \leq t/\delta$. For a good vertex u , and $j \in \text{Sugg}_u$:

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

Therefore, for at least a fraction $\delta/2$ of neighbors v of u , $\text{Inf}_j^{\leq t}(\overline{f^v \circ \sigma_v}) \geq \delta/2$. For such v and j , by Claim 3.23, $\text{Inf}_{2j-1}^{\leq 2t}(f^v \circ \sigma_v) + \text{Inf}_{2j}^{\leq 2t}(f^v \circ \sigma_v) \geq \delta/2$. Therefore for some $j \in [2R]$, $\text{Inf}_j^{\leq 2t}(f^v) \geq \delta/4$. Let $\text{Sugg}_v = \{j | j \in \{1, \dots, 2R\} \wedge \text{Inf}_j^{\leq 2t}(f^v) \geq \delta/4\}$. Again, Sugg_v is not empty and $|\text{Sugg}_v| \leq 8t/\delta$.

Following the decoding procedure in [11], we deduce that it is possible to satisfy a fraction $\gamma' = \gamma'(\delta, t) = \gamma'(k)$ of the constraints. \square

Note that our PCP verifier makes “ k -coloring” tests. By the standard conversion from PCP verifiers to CSP hardness, and Remark 2.6 about conversion to unweighted graphs with the same inapproximability factor, we conclude the main result of this section by combining Lemmas 3.24 and 3.25.

Theorem 3.26. *For every constant $k \geq 3$, assuming the 2-to-1 Conjecture, it is NP-hard to approximate Max k -Colorable Subgraph within a factor of $1 - 1/k + O(\ln k / k^2)$.*

A Details about “tripartite” PCP

We now briefly discuss the (minor) changes in analysis needed to make the PCP in [6] have the properties claimed in Remark 2.5. The discussion assumes substantial familiarity with Håstad-style PCP constructions [8], and is intended as an aid for an expert to check our claim.

The crux in PCP constructions is a procedure to test that two tables A, B are the legal long code encodings of two labels a, b which satisfy some projection constraint $\pi(b) = a$. (In the overall construction, this projection constraint is between two nodes u, v of a Label Cover instance, and A, B supposedly encodes labels to u, v respectively.) Here $\pi : [M] \rightarrow [N]$, where we denote $[m] = \{1, 2, \dots, m\}$, $A : \mathcal{F}_N \rightarrow \{1, -1\}$ and $B : \mathcal{F}_M \rightarrow \{1, -1\}$ with \mathcal{F}_m denoting the set of Boolean functions $[m] \rightarrow \{1, -1\}$. (It is convenient to use ± 1 to denote truth values, with -1 being true.) We say A is the long code of $a \in [N]$ if $A(f) = f(a)$ for all $f \in \mathcal{F}_N$, and similarly for B being the long code of $b \in [M]$. Note that a long code has the following “folded” property: $A(-f) = -A(f)$. We can assume that the tables A, B are folded, by inferring the value $A(-f)$ from the value $A(f)$ for half the functions.

The specific nature of the test must correspond to the predicate of the CSP for which inapproximability is sought. Specifically, for the PCP of [6], which checks the predicate

$$(A(f) \vee (B(g) = B(g_1))) \wedge (\overline{A(f)} \vee (B(g) = B(g_2))) , \quad (\text{A.1})$$

one would pick functions f, g, g_1, g_2 such that for $(a, b) \in [N] \times [M]$ with $\pi(b) = a$, we have

$$((f(a) = 1) \implies (g(b) = g_1(b))) \wedge ((f(a) = -1) \implies (g(b) = g_2(b))) .$$

This would ensure that the test has perfect completeness, i.e., accepts A, B which are long codes of a, b with $\pi(b) = a$ with probability 1. The actual test picks f, g randomly, and defines $g_1 = (f \circ \pi \wedge h)g$ and $g_2 = (-f \circ \pi \wedge h')g$ for functions h, h' being independent samples from some suitable distribution. It is easy to see that the perfect completeness of the test holds for any choice of h, h' . The soundness claim relies on picking h, h' according to a carefully chosen distribution (see [6] or [8] for details), which in particular sets those functions to -1 at each point with probability close to 1.

We note that the foldedness of A, B will imply negations in the constraint (A.1) above.

The structure of the soundness analysis in [6], which in fact is essentially identical to the analysis for the PCP for satisfiable Max 3SAT from [8], is as follows. Assuming that the test accepts with probability at least $1/2 + \varepsilon$, one shows how to decode A and B into labels a and b respectively such that $\pi(b) = a$ holds with probability $\varepsilon_1(\varepsilon) > 0$. When combined with the hardness of Label Cover, such a decoding procedure immediately implies the construction of a PCP with soundness $1/2 + \varepsilon$, which in turn implies the inapproximability of satisfying a $1/2 + \varepsilon$ fraction constraints of a satisfiable instance of the associated CSP.

The change we make to the PCP in [6] is to make the queries g_1, g_2 into a C -table, that is distinct from B and further is *not* assumed to be folded. Note that the perfect completeness of the test is maintained, as one can simply choose $C = B$ (or to be precise, the unfolded version of table B). Let us turn to how the soundness analysis is affected.

The probability that the test (A.1) accepts (when the queries $B(g_i)$ are replaced by $C(g_i)$) can be arithmetized as follows:

$$1 - \mathbb{E}_{f,g,g_1,g_2} \left[\left(\frac{1+A(f)}{2} \right) \left(\frac{1-B(g)C(g_1)}{2} \right) + \left(\frac{1-A(f)}{2} \right) \left(\frac{1-B(g)C(g_2)}{2} \right) \right].$$

Using the fact that A is folded (so $A(-f) = -A(f)$ and $\mathbb{E}_f[A(f)] = 0$), and the distribution of (f, g, g_2) is identical to that of $(-f, g, g_1)$, the above expression simplifies to

$$\frac{1}{2} + \frac{1}{2} \mathbb{E}_{f,g,g_1} [B(g)C(g_1)] + \frac{1}{2} \mathbb{E}_{f,g,g_1} [A(f)B(g)C(g_1)].$$

The analysis in [8, Thm. 6.5] upper bounds $\mathbb{E}[B(g)B(g_1)]$ by ε . Actually, this is not done for a fixed A table, but rather when expectation is also taken over the projection π corresponding to a random neighbor u of the node v whose label B is encoding. This part of the analysis is not affected by the change to a C -table, so we ignore this complication in the description here.

The analysis proceeds by expanding $\mathbb{E}[B(g)B(g_1)]$ by Fourier analysis and upper bounds

$$\sum_{\beta} \hat{B}_{\beta}^2 \prod_{x \in \pi(\beta)} \left| \frac{(-1)^{s_x} + (1 - 2\varepsilon)^{s_x}}{2} \right| \quad (\text{A.2})$$

where $s_x = |\pi^{-1}(x) \cap \beta|$ (this is eq. (36) in [8]). To bound (A.2), Håstad proceeds by dividing the sum into three parts, small $|\beta|$, medium $|\beta|$, and large $|\beta|$ (see Lemma 6.7 of [8]). The sum for small and large sized β is bounded by a small δ , and the contribution of the medium terms is amortized over the different choices of the bias in the distribution of the function h in the choice of g_1 . We can mimic his analysis essentially verbatim: the only change we need to do is replace \hat{B}_{β}^2 by $|\hat{B}_{\beta}| |\hat{C}_{\beta}|$. In the analysis of the medium and large $|\beta|$ terms, wherever Håstad uses the upper bound $\sum_{\beta} \hat{B}_{\beta}^2 \leq 1$ in the analysis, we can instead use $\sum_{\beta} |\hat{B}_{\beta}| |\hat{C}_{\beta}| \leq 1$ which follows from Parseval's identity and Cauchy-Schwarz inequality. The bound for small β (Lemma 6.8 in [8]) uses the fact that terms with even $|\beta|$ contribute 0 to (A.2) by virtue of B being folded. This remains true for us — even though C may not be folded, for $|\hat{B}_{\beta}| |\hat{C}_{\beta}|$ to be nonzero, we must have $|\beta|$ to

be odd as B is still folded. Thus, in the end, we can upper bound $\mathbb{E}[B(g)C(g_1)]$ by the same bound Håstad got for $\mathbb{E}[B(g)B(g_1)]$.

Let us now turn to the $\mathbb{E}[A(f)B(g)C(g_1)]$ term. Once again the analysis is similar to that of the $\mathbb{E}[A(f)B(g)B(g_1)]$ term in [8], where it is shown that if $\mathbb{E}[A(f)B(g)B(g_1)] \geq \delta$, then one can decode labels a, b such that $\pi(b) = a$ with probability $\delta_1 = \delta_1(\delta) > 0$. For us, using Fourier expansion the assumption $\mathbb{E}[A(f)B(g)C(g_1)] \geq \delta$ gives

$$\left| \mathbb{E} \left[\sum_{\beta, \alpha \subseteq \pi(\beta)} \hat{A}_\alpha \hat{B}_\beta \hat{C}_\beta p(\alpha, \beta) \right] \right| \geq \delta ,$$

similar to Equation (42) of [8] (where $p(\alpha, \beta)$ is also defined). The only change compared to [8] is that $\hat{B}_\beta \hat{C}_\beta$ replaces \hat{B}_β^2 . One can check that the rest of the analysis in the proof of Lemma 6.11 of [8] goes through unaltered if we replace \hat{B}_β^2 by $|\hat{B}_\beta| |\hat{C}_\beta|$ everywhere. Once again all uses of the key fact $\sum_\beta \hat{B}_\beta^2 \leq 1$ can be replaced by the valid inequality $\sum_\beta |\hat{B}_\beta| |\hat{C}_\beta| \leq 1$. A minor change we make in the decoding procedure is that the “clause prover” P_1 picks a random subset β with probability *proportional* to $|\hat{B}_\beta| |\hat{C}_\beta|$ instead of probability equal to \hat{B}_β^2 , and then returns a random $y \in \beta$. As the B -table is folded, we have $\hat{B}_\emptyset = 0$, and so the decoding procedure will always pick a nonempty β and is thus well defined. Since $\sum_\beta |\hat{B}_\beta| |\hat{C}_\beta| \leq 1$, the probability of picking β is *at least* $|\hat{B}_\beta| |\hat{C}_\beta|$, so replacing \hat{B}_β^2 by $|\hat{B}_\beta| |\hat{C}_\beta|$ still gives a *lower bound* on the success of the label-decoding procedure, which is what we seek.

Thus, with these changes we can complete the soundness analysis of the tripartite PCP where the B table is split into two tables B and C , and without requiring the C table to be folded (though we necessarily require the B table to be folded). This gives the PCP underlying the inapproximability result claimed in Proposition 2.4.

References

- [1] PER AUSTRIN, RYAN O’DONNELL, LI-YANG TAN, AND JOHN WRIGHT: New NP-hardness results for 3-coloring and 2-to-1 label cover. *CoRR*, abs/1210.5648, 2012. 4, 7
- [2] PIERLUIGI CRESCENZI, RICCARDO SILVESTRI, AND LUCA TREVISAN: On weighted vs unweighted versions of combinatorial optimization problems. *Inf. Comput.*, 167(1):10–26, 2001. 5
- [3] IRIT DINUR, ELCHANAN MOSSEL, AND ODED REGEV: Conditional hardness for approximate coloring. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pp. 344–353, 2006. 3, 10, 13, 14, 15
- [4] ALAN M. FRIEZE AND MARK JERRUM: Improved approximation algorithms for max k-cut and max bisection. *Algorithmica*, 18(1):67–81, 1997. 2
- [5] VENKATESAN GURUSWAMI: Inapproximability results for set splitting and satisfiability problems with no mixed clauses. *Algorithmica*, 38(3):451–469, 2003. 4
- [6] VENKATESAN GURUSWAMI, DANIEL LEWIN, MADHU SUDAN, AND LUCA TREVISAN: A tight characterization of NP with 3 query PCPs. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pp. 8–17, 1998. 3, 4, 5, 16, 17
- [7] VENKATESAN GURUSWAMI AND ALI KEMAL SINOP: Improved inapproximability results for maximum k-colorable subgraph. In *APPROX-RANDOM*, pp. 163–176, 2009. 1

- [8] JOHAN HÅSTAD: Some optimal inapproximability results. *Journal of the ACM*, 48:798–859, 2001. 5, 16, 17, 18
- [9] VIGGO KANN, SANJEEV KHANNA, JENS LAGERGREN, AND ALESSANDRO PANCONESI: On the hardness of approximating max k-cut and its dual. *Chicago J. Theor. Comput. Sci.*, 1997, 1997. 2, 3, 7, 9
- [10] SUBHASH KHOT: On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pp. 767–775, 2002. 10
- [11] 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, 12, 13, 14, 15, 16
- [12] ELCHANAN MOSSEL, RYAN O’DONNELL, AND KRZYSZTOF OLESZKIEWICZ: Noise stability of functions with low influences: invariance and optimality. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pp. 21–30, 2005. 3, 14
- [13] RYAN O’DONNELL AND YI WU: Conditional hardness for satisfiable CSPs. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, 2009. 3
- [14] CHRISTOS H. PAPADIMITRIOU AND MIHALIS YANNAKAKIS: Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991. 2
- [15] EREZ PETRANK: The hardness of approximation: Gap location. *Computational Complexity*, 4:133–157, 1994. 2, 3, 4
- [16] PRASAD RAGHAVENDRA: Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th ACM Symposium on Theory of Computing*, pp. 245–254, 2008. 3
- [17] ALISTAIR SINCLAIR: Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1:351–370, 1992. 11

AUTHORS

Venkatesan Guruswami
 Associate professor
 Carnegie Mellon University, Pittsburgh, PA
guruswami@cmu.edu
<http://www.cs.cmu.edu/~venkatg>

Ali Kemal Sinop
 Postdoctoral Research Associate
 Princeton University, Princeton, NJ
asinop@cs.cmu.edu
<http://www.cs.cmu.edu/~asinop>

ABOUT THE AUTHORS

VENKATESAN GURUSWAMI, or Venkat, as his colleagues call him, received his Bachelor's degree from the [Indian Institute of Technology at Madras](#) in 1997. He is grateful to his undergraduate research advisor [C. Pandu Rangan](#) for initiating him to the joys of theoretical computer science. Venkat received his Ph. D. from the [Massachusetts Institute of Technology](#) in 2001 under the excellent advisorship of [Madhu Sudan](#).

Since 2008, Venkat has been on the faculty of the [Computer Science Department](#) at [Carnegie Mellon University](#). Earlier, he was a faculty member at the [University of Washington](#). Venkat was a [Miller Research Fellow](#) at [UC Berkeley](#) during 2001-02, and was a member in the [School of Mathematics, Institute for Advanced Study](#) during 2007-08.

Venkat is interested in several topics in theoretical computer science, including algorithmic and algebraic coding theory, approximability of fundamental optimization problems, pseudorandomness, probabilistically checkable proofs, computational complexity theory, and algebraic algorithms.

Venkat currently serves on the editorial boards of the [SIAM Journal on Computing](#), [IEEE Transactions on Information Theory](#), and the [ACM Transactions on Computation Theory](#). He was recently program committee chair for the 2012 [Computational Complexity conference](#). Venkat is a recipient of the [Presburger award](#), [Packard Fellowship](#), [Sloan Fellowship](#), NSF CAREER award, [ACM's Doctoral Dissertation Award](#), and the [IEEE Information Theory Society Paper Award](#).

In his (sadly limited) spare time, Venkat enjoys traveling, ethnic vegetarian food, racquet sports, hiking/running within his humble limits, and listening to Carnatic music.

ALI KEMAL SINOP graduated from [C.M.U.](#) in 2012; his advisor was [Venkatesan Guruswami](#). His thesis focused on the use of semi-definite programming hierarchies for various graph partitioning problems.