

d -to-1 Hardness of Coloring 4-colorable Graphs with $O(1)$ colors

Venkatesan Guruswami* Sai Sandeep*

Abstract

The d -to-1 conjecture of Khot asserts that it is hard to satisfy an ϵ fraction of constraints of a satisfiable d -to-1 Label Cover instance, for arbitrarily small $\epsilon > 0$. We prove that the d -to-1 conjecture for any fixed d implies the hardness of coloring a 4-colorable graph with C colors for arbitrarily large integers C . Earlier, this implication was only known under the 2-to-1 conjecture, which is the strongest in the family of d -to-1 conjectures.

1 Introduction

Determining if a graph is 3-colorable is one of the classic NP-complete problems. Thus given a 3-colorable graph it is NP-hard to color it with 3 colors. The best known polynomial time algorithms for coloring 3-colorable graphs use about $n^{0.2}$ colors, where n is the number of vertices in the graph [KT17]. On the other hand, on the hardness front, we only know that 5-coloring 3-colorable graphs is NP-hard [BKO19].

This embarrassingly large gap between the hardness and algorithmic results has prompted the quest for conditional hardness results for approximate graph coloring. The canonical starting point for most strong inapproximability results is the Label Cover problem. Label Cover refers to constraint satisfaction problems of arity two over a large (but fixed) domain whose constraint relations are *functions*. Label Cover is known to be very hard to approximate even on satisfiable instances.

The Unique Games Conjecture of Khot [Kho02], which asserts strong inapproximability of Label Cover when the constraint maps are bijections, has formed the basis of numerous tight hardness results for problems which have defied NP-hardness proofs. However, the imperfect completeness inherent in the Unique Games Conjecture makes it unsuitable as the basis for hardness results for graph coloring, where we want *all* edges to be properly colored under the coloring.

In [Kho02], along with the Unique Games Conjecture, Khot introduced the d -to-1 conjecture. The d -to-1 conjecture says that given a Label Cover instance whose constraint relations are d -to-1 functions, it is NP-hard to decide if there exists a labelling that satisfies all the

*Computer Science Department, Carnegie Mellon University, Pittsburgh, USA. Email: {venkatg, spaller1}@cs.cmu.edu. Research supported in part by NSF grants CCF-1422045 and CCF-1908125.

constraints or no labelling can satisfy even an ϵ fraction of constraints, for arbitrarily small $\epsilon > 0$. (The key is that d can be held fixed and achieve soundness $\epsilon \rightarrow 0$.) Constraints similar to 2-to-1 also played an implicit role in the beautiful work of Dinur and Safra on inapproximability of vertex cover [DS05].

Based on the 2-to-1 conjecture, Dinur, Mossel and Regev [DMR09], extending the invariance principle based techniques of [KKMO07, MOO10], proved the hardness of coloring graphs that are promised to be 4-colorable with any constant number of colors. In this note, we prove that the same result can be proved under the weaker assumption of d -to-1 conjecture¹, for some (arbitrarily large) constant d .

Theorem 1. *Assume that d -to-1 conjecture is true for some constant d . Then, for every positive integer $t \geq 4$, it is NP-hard to color a 4-colorable graph G with t colors.*

We stress that the d -to-1 conjecture insists on perfect completeness (i.e., hardness on satisfiable instances), and this important feature seems necessary for its implications for coloring problems, where we seek to properly color all edges. The variant of the 2-to-1 conjecture where one settles for near-perfect completeness was recently established in a striking sequence of works [KMS17, DKK⁺18a, DKK⁺18b, KMS18].

The result of [DMR09] in fact shows hardness of finding an independent set of density ϵ in a 4-colorable graph for arbitrary $\epsilon > 0$ (which immediately implies hardness of finding a coloring with $1/\epsilon$ colors). Our result in Theorem 1 above does *not* get this stronger hardness for finding independent sets. But it is conditioned on the d -to-1 conjecture for arbitrary d rather than the specific 2-to-1 conjecture. We note that proving the d -to-1 conjecture for some large d could be significantly easier than the 2-to-1 conjecture, so Theorem 1 perhaps provides an avenue for resolving a longstanding challenge concerning the complexity of approximate graph coloring.

Our proof of Theorem 1 is a simple combination of two results. First, following the methodology of [DMR09], we prove that the d -to-1 conjecture implies that coloring a $2d$ -colorable graph with $O(1)$ colors is NP-hard. The result of [DMR09] is the $d = 2$ case of this claim. In fact, they state in a future work section that the d -to-1 conjecture should imply hardness of $O(1)$ -coloring q -colorable graphs for some large enough $q = q(d)$. However, they did not specify the details of the reduction or an explicit value of q , and mention determining the dependence of q on d as an interesting question. Here we show the conditional hardness based on d -to-1 conjecture holds for $q = 2d$ (achieving $q < 2d$ seems unlikely with the general reduction approach of [DMR09]).

The key technical ingredient necessary for such a reduction is a symmetric Markov chain on $[q]^d$ where transitions are allowed only between disjoint tuples and which has spectral radius bounded away from 1. We show the existence of such a symmetric Markov chain for $q = 2d$. We do so via a connection to matrix scaling, which enables us to deduce the necessary chain at a conceptual level without messy calculations. Specifically, we use the result [CD72], which follows from the Sinkhorn-Knopp iterative matrix scaling algorithm [SK67], that if a non-negative symmetric matrix A has *total support* then there is a symmetric doubly stochastic matrix supported on the non-zero entries of A . When A is the adjacency matrix of a graph

¹For d -to-1 Label Cover, there are two definitions possible, one where the constraint maps are *at most* d -to-1 with each element in the range having at most d pre-images, and one where the constraint maps are *exactly* d -to-1. In this paper, we stick with the exact variant.

G , the total support condition is equivalent to every edge of G belonging to a cycle cover. We describe a graph on $[q]^d$ whose edges connect disjoint tuples and where every edge belongs to a cycle cover.

Our second ingredient is a remarkable yet simple reduction due to Wrochna and Živný [WZ19], which exploits the relation between the arc-chromatic number and chromatic number of a digraph [PR81]. Let $b : \mathbb{N} \rightarrow \mathbb{N}$ be defined by $b(n) := \binom{n}{\lfloor n/2 \rfloor}$. Their result then is that $b(t)$ -coloring $b(c)$ -colorable graphs is polynomial time (in fact logspace) reducible to t -coloring c -colorable graphs. Since $b(n)$ is increasing and $b(n) > n$ for all $n \geq 4$, it follows that a NP-hardness result for $O(1)$ -coloring q -colorable graphs also implies NP-hardness of $O(1)$ -coloring 4-colorable graphs.

Overview. In Section 2, we define the Label Cover problem, and state the d -to-1 conjecture formally. We also introduce low degree influences that we need later. In Section 3, we first prove the existence of the Markov chain with required properties, and then describe the reduction from Label Cover to Approximate Coloring. We note that the reduction is in fact exactly the same one used in [DMR09], the difference being in using a different Markov Chain. We present the reduction and the preliminaries required in this paper for the sake of completeness.

2 Preliminaries

We first formally define the Label Cover problem and then state the hardness conjectures.

2.1 Label Cover

Definition 2. (*Label Cover*) In the Label Cover instance, we are given a tuple $G = ((V, E), R, \Psi)$ where

1. (V, E) is a graph on vertex set V with edge set E .
2. Each vertex in V has to be assigned a label from the set $\Sigma = [R] = \{1, 2, \dots, R\}$.
3. For every edge $e = (u, v) \in E$, there is an associated relation $\Psi_e \subseteq \Sigma \times \Sigma$. This corresponds to a constraint between u and v .

A labeling $\sigma : V \rightarrow \Sigma$ satisfies a constraint associated with the edge $e = (u, v)$ if and only if $(\sigma(u), \sigma(v)) \in \Psi_e$. Given such an instance, the goal is to distinguish if there is a labeling that can satisfy all the constraints or no labeling can satisfy a significant fraction of constraints.

We now state the d -to-1 conjecture. As is the case with [DMR09], we will state and use the exact d -to-1 variant where the constraint maps have exactly d pre-images for each element in the range. Khot's original formulation only required that there are at most d pre-images for each element in the range. The d -to-1 conjecture becomes stronger for smaller d (so that the 2-to-1 is the strongest form of the conjecture)—this is obvious for the variant where the maps are at most d -to-1. For the exact variant, if we allow the Label cover graph to have multiple edges, we can reduce d -to-1 conjecture to $(d + 1)$ -to-1 conjecture using a simple argument.

We present this reduction in Appendix A. On that note, we remark without details that our reduction indeed works with the multigraph variant of d -to-1 conjecture.

Conjecture 3. (*(Exact) d -to-1 Conjecture*) For every $\epsilon > 0$, given a bipartite Label Cover instance $G = ((V = X \cup Y, E), (dR, R), \Psi)$ satisfying the following constraints:

- (i) We refer to X as the vertices on the left, and Y as the set of vertices on the right. The vertices belonging to X are to be assigned labels from $[dR]$ while the vertices in Y are to be assigned labels from $[R]$.
- (ii) The constraints are d -to-1 i.e. for every $b \in [R]$, there are precisely d values $a \in [dR]$ such that $(a, b) \in \Psi_e$ for every relation Ψ_e in the instance.

It is NP-hard to distinguish between the following cases:

1. There is a labeling that satisfies all the constraints in G .
2. No labeling can satisfy more than ϵ fraction of constraints in G .

Similar to the d -to-1 constraints, one can consider d -to- d constraints in the Label Cover. In order to do so, we define the relation $d \leftrightarrow d$ on $[dR] \times [dR]$:

$$d \leftrightarrow d = \{(di - p + 1, di - q + 1) \mid 1 \leq i \leq R, 1 \leq p, q \leq d\} .$$

A constraint $\psi \subseteq [dR] \times [dR]$ is said to be d -to- d if there exist permutations π_1 and π_2 on $[dR]$ such that $(a, b) \in \psi$ iff $(\pi_1^{-1}(a), \pi_2^{-1}(b)) \in d \leftrightarrow d$.

In [DMR09], it is proved that Conjecture 3 implies the following conjecture.

Conjecture 4. (*d -to- d conjecture*) For every $\epsilon > 0$ and every $t \in \mathbb{N}$, there exists $R \in \mathbb{N}$ such that given a Label Cover instance $G = ((V, E), dR, \Psi)$ where all the constraints are d -to- d , it is NP-hard to distinguish between the following cases:

- (i) $\text{sat}(G) = 1$, or
- (ii) $\text{isat}_t(G) < \epsilon$

Here, $\text{sat}(G)$ denotes the maximum fraction of constraints satisfied by any labeling. Similarly, $\text{isat}(G)$ denotes the size of the largest set $S \subseteq V$ such that there exists a labeling that satisfies all the constraints induced on S . The value $\text{isat}_t(G)$ denotes the size of largest set $S \subseteq V$ such that there exists a labeling that assigns at most t labels to each vertex that satisfies all the constraints induced on S . A constraint between u, v is said to be satisfied by labeling assigning multiple labels to u and v if and only if there exists at least one pair of labels to u and v among the multiple labels that satisfy the constraint.

2.2 Low degree influences

Next, we define the low degree influences that we need later. We refer the reader to [DMR09] for a comprehensive treatment of the same.

Let $\alpha_0 = \mathbf{1}, \alpha_1, \dots, \alpha_{q-1}$ be an orthonormal basis of \mathbb{R}^q . We can define the set of functions $\alpha_x : [q]^n \rightarrow \mathbb{R}, x \in [q]^n$ as $\alpha_x(y) = (\alpha_{x_1}(y_1), \alpha_{x_2}(y_2), \dots, \alpha_{x_n}(y_n))$. Observe that these

functions form a basis for the functions from $[q]^n$ to \mathbb{R} . Let $\hat{f}(\alpha_x) = \langle f, \alpha_x \rangle$, where we define the inner product between functions $f, g : [q]^n \rightarrow \mathbb{R}$ as $\langle f, g \rangle = q^{-n} \sum_{x \in [q]^n} f(x)g(x)$. We define the low degree influence of f as follows:

Definition 5. For a function $f : [q]^n \rightarrow \mathbb{R}$, the degree k influence of the coordinate i is defined as follows:

$$I_i^{\leq k}(f) = \sum_{x: x_i \neq 0, |x| \leq k} \hat{f}^2(\alpha_x)$$

Note that the above definition is independent of the basis $\alpha_0, \alpha_1, \dots, \alpha_{q-1}$ that we start with, as long as $\alpha_0 = \mathbf{1}$. From the above definition, we can infer that for functions $f : [q]^n \rightarrow [0, 1]$, the sum of low degree influences is bounded by

$$\sum_i I_i^{\leq k}(f) \leq k$$

For a vector $x \in [q]^{dR}$, let $\bar{x} \in [q^d]^R$ be the corresponding element in $[q^d]^R$ i.e.

$$\bar{x} = ((x_1, x_2, \dots, x_d), (x_{d+1}, x_{d+2}, \dots, x_{2d}), \dots, (x_{dR-d+1}, x_{dR-d+2}, \dots, x_{dR}))$$

Similarly, for $y \in [q^d]^R$, let \underline{y} denote the inverse of above operation. We can extend this notion to functions as well: For a function $f : [q]^{dR} \rightarrow \mathbb{R}$, let the function $\bar{f} : [q^d]^R \rightarrow \mathbb{R}$ be defined naturally by

$$\bar{f}(y) = f(\underline{y})$$

Similarly, for a function $f : [q^d]^R \rightarrow \mathbb{R}$, let $\underline{f} : [q]^{dR} \rightarrow \mathbb{R}$ be defined as $\underline{f}(x) = f(\bar{x})$.

We need the following lemma:

Lemma 6. For any function $f : [q]^{dR} \rightarrow \mathbb{R}$ and any $k \in \mathbb{N}$ and $i \in [R]$,

$$I_i^{\leq k}(\bar{f}) \leq \sum_{j=1}^d I_{di-d+j}^{\leq dk}(f)$$

Proof. Fix a basis α_x of functions from $[q]^{dR} \rightarrow \mathbb{R}$ as above. The functions $\alpha_{\bar{x}}$ form a basis for functions from $[q^d]^R \rightarrow \mathbb{R}$, where $\alpha_{\bar{x}}(\bar{y}) = \alpha_x(y)$. Note that $\hat{\bar{f}}(\alpha_{\bar{x}}) = \hat{f}(\alpha_x)$. Thus we get

$$\begin{aligned} \sum_i I_i^{\leq k}(\bar{f}) &= \sum_{\bar{x}: \bar{x}_i \neq (0,0,\dots,0), |\bar{x}| \leq k} \hat{\bar{f}}^2(\alpha_{\bar{x}}) = \sum_{\bar{x}: \bar{x}_i \neq (0,0,\dots,0), |\bar{x}| \leq k} \hat{f}^2(\alpha_x) \\ &\leq \sum_{x: \bar{x}_i \neq (0,0,\dots,0), |x| \leq dk} \hat{f}^2(\alpha_x) \\ &\leq \sum_{j=1}^d \sum_{x: x_{di-d+j} \neq 0, |x| \leq dk} \hat{f}^2(\alpha_x) \\ &= \sum_{j=1}^d I_{di-d+j}^{\leq dk}(f) \end{aligned} \quad \square$$

Using the invariance principle and Borell's inequality, [DMR09] prove the following:

Theorem 7. *Let q be a fixed integer, and T be a symmetric Markov chain on $[q]$ with $r(T) < 1$. Then for every $\epsilon > 0$, there exists a $\delta > 0$ and a positive integer k such that the following holds: For every $f, g : [q]^n \rightarrow [0, 1]$ if $\mathbb{E}[f] > \epsilon$, $\mathbb{E}[g] > \epsilon$ and $\langle f, Tg \rangle = 0$, then*

$$\exists i \in [n] : I_i^{\leq k}(f) \geq \delta, I_i^{\leq k}(g) \geq \delta$$

where $r(T)$ denotes the second largest eigenvalue (in absolute value) of T .

3 d -to-1 hardness for 4-colorable graphs

In this section, we will prove Theorem 1.

3.1 Reducing chromatic number to 4

The following lemma is implicit in [WZ19] based on a beautiful result concerning the arc-chromatic numbers of digraphs from [PR81].

Lemma 8. *Suppose there exists $q \in \mathbb{N}$ such that $O(1)$ coloring q -colorable graphs is NP-hard. Then, $O(1)$ coloring 4-colorable graphs is NP hard.*

Proof. Let $b : \mathbb{N} \rightarrow \mathbb{N}$ be the function defined $b(n) := \binom{n}{\lfloor n/2 \rfloor}$, and Graph-Coloring(t, c) denote the promise problem of distinguishing if a graph can be colored with c colors, or cannot even be colored with t colors. In [WZ19] (Lemma 3.3), it is shown that Graph-Coloring($b(t), b(c)$) reduces to Graph-Coloring(t, c) in polynomial time. Since b is an increasing function and $b(n) > n$ for all $n \geq 4$, setting $c = 4$ and t large enough proves the statement claimed in the lemma. \square

Thanks to Lemma 8, we can restrict ourselves to the weaker goal of proving that $O(1)$ coloring q -colorable graphs is NP-hard for some fixed constant q assuming Conjecture 3. In fact, following [DMR09], we prove a stronger statement showing hardness of finding independent sets of ϵ fraction of vertices for any $\epsilon > 0$. Combined with Lemma 8, this immediately gives us Theorem 1.

Theorem 9. *Suppose that Conjecture 4 is true for a constant d . Then, there exists a constant $q = q(d)$ such that for every $\epsilon > 0$, given a graph G , it is NP-hard to distinguish the following cases:*

1. G can be colored with q colors.
2. G does not have any independent set of relative size ϵ .

In fact, we can take $q = 2d$.

In the remainder of the section, we will prove Theorem 9. We next develop the main technical ingredient that we will plug into the reduction framework of [DMR09] to establish Theorem 9.

3.2 A symmetric Markov chain supported on disjoint tuples

A Markov chain T defined on a state space Ω is said to be symmetric if the transition matrix of T is symmetric, namely for all pairs of states $x, y \in \Omega$, the probability of transition from x to y is equal to the probability of transition from y to x . Symmetry of the Markov chain ensures that the uniform distribution is stationary which is essential when we compose the Label Cover-Long Code reduction with the Markov chain. We define the spectral radius $r(T)$ of a symmetric Markov chain as the second largest eigenvalue in absolute value of its transition probability matrix, i.e., if $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_q$ are the eigenvalues, then $r(T) = \max(|\lambda_2|, |\lambda_q|)$.

We now show the existence of a symmetric Markov Chain T on $[q]^d$ with $r(T) < 1$ if $d \geq 2, q \geq 2d$. Furthermore, there is a nonzero transition probability between two elements $x, y \in [q]^d$ only if the support of x and y are disjoint. In [DMR09], such a Markov Chain is shown to exist for the values $(q, d) = (3, 1), (4, 2)$.

Lemma 10. *Suppose that $q, d \in \mathbb{N}, q \geq 2d, d \geq 2$. There exists a symmetric Markov chain T on $[q]^d$ such that $r(T) < 1$. Furthermore, if the transition $\{x_1, x_2, \dots, x_d\} \leftrightarrow \{y_1, y_2, \dots, y_d\}$ has positive probability in T , then $\{x_1, x_2, \dots, x_d\} \cap \{y_1, y_2, \dots, y_d\} = \phi$.*

Proof. We first construct an undirected graph G on $[q]^d$ such that there is an edge between $x, y \in [q]^d$ only if the support of x and y are disjoint. We then use a matrix scaling algorithm to obtain a symmetric Markov chain T from the adjacency matrix of G . For the resulting Markov chain to have $r(T) < 1$, we need that the underlying graph G is connected, and is not bipartite. Furthermore, for the scaling algorithm to produce a valid Markov chain, we need that every edge of G is present in a cycle cover, where a cycle cover of a graph is a disjoint union of cycles that covers every vertex in the graph. Note that we allow trivial 2-cycles in a cycle cover, where we just take an edge twice.

We say that two multisets $x = (x_1, x_2, \dots, x_d), y = (y_1, y_2, \dots, y_d) \in [q]^d$ are of the same *type* if the following condition holds: for all pairs of indices $i, j \in [d]$, $x_i = x_j$ if and only if $y_i = y_j$ and $(x_i - x_j)(y_i - y_j) \geq 0$. Note that this is an equivalence relation, and thus each element $x \in [q]^d$ uniquely determines its type.

Consider the graph $G = (V, E)$ where the vertex set is $V = [q]^d$. We add two kinds of edges in this graph. We add an edge between every pair of $x, y \in [q]^d$ that are of the same type, and have disjoint support. Let the subset of $[q]^d$ of elements that are supported on single element be denoted by S , i.e.,

$$S = \{(1, 1, \dots, 1), (2, 2, \dots, 2), \dots, (q, q, \dots, q)\}.$$

We also add edges between x and y if their support is disjoint, and at least one of x and y belongs to S .

First, we claim that G is connected. This follows from the fact that the set of nodes in S are connected to each other, and every vertex in V is adjacent to at least one vertex in S . As $q \geq 4$, the graph is not bipartite (indeed S induces a q -clique). We will now prove that every edge in this graph is part of a cycle cover. Given an undirected graph on vertex set V , a cycle cover of it is a function $\sigma : V \rightarrow V$ that is bijective, and $\sigma(u) = v$ only when u and v are adjacent in the underlying graph.

Towards this, we first prove that for every edge in G between multisets of the same type, there is a cycle cover that uses that edge. For each type, consider the graph obtained by taking the vertices as multisets of that type, and with edges between two multisets of the same type if they are disjoint. Note that for every type, this graph is isomorphic to a Kneser graph $KG(q, k)$ (for some $k \leq d$), whose vertex set corresponds to k -element subsets of $[q]$ and there is an edge between two subsets if they are disjoint.

By symmetry across the subsets, we can infer that the Kneser graphs are regular. Note that every regular graph contains a cycle cover: For a regular graph H , consider a bipartite graph H' which contains a copy of H on both the left side L , and right side R . There is an edge between $x \in L, y \in R$ of H' if and only if x, y are adjacent in H . As H is a regular graph, H' is a regular bipartite graph, and thus, contains a perfect matching. This perfect matching in H' directly gives a cycle cover of H . Furthermore, as Kneser graphs are also vertex-transitive, every edge in these graphs is part of a cycle cover.

Next, we consider edges of G that are between multisets of different types i.e. edges between multisets x, y where exactly one of x and y is in S . Consider an edge between $s \in S$ and $x \in V \setminus S$. As $q \geq 2d$, every multiset in G is adjacent to at least one multiset of the same type. Let y be a multiset that is adjacent to x in G and is of the same type as x . Let $s' \in S$ be chosen such that it is adjacent to y in G . As S is a complete subgraph of G , s and s' are adjacent in G . From the previous argument about edges between multisets of the same type, we can infer that there is a cycle cover of G where y is mapped to x , and s is mapped to s' . We can modify this cycle cover by transforming it as follows - $(s \rightarrow x)$ can be made part of cycle cover by transforming $(s \rightarrow s'), (y \rightarrow x)$ to $(s \rightarrow x), (y \rightarrow s')$ and keeping rest of the cycle cover intact. Thus, we have proved that every edge of G is part of a cycle cover.

Let A denote the adjacency matrix of the above graph G . Using the Sinkhorn Knopp iterative algorithm, it is proved in [CD72] that if a non-negative symmetric matrix A has total support, then there exists a diagonal matrix D such that DAD is a doubly stochastic matrix. A square matrix $A = (a_{ij})$ of order n is said to have total support if $A \neq 0$, and for every nonzero entry a_{ij} of A , there exists a permutation σ of $[n]$ such that $\sigma(i) = j$ and for all $e \in [n], a_{e, \sigma(e)} \neq 0$. When the matrix A is an adjacency matrix of a graph G , the total support condition translates to the requirement that every edge in G is part of a cycle cover, a property we have already shown to hold for the graph G .

Thus, we can apply the above scaling result, and view the resulting matrix $B = DAD$ as the transition matrix of a Markov chain T . As A and D are symmetric, B is symmetric, i.e., T is symmetric. As A is connected and no principal diagonal element of D is zero, T is connected as well. Note that every nonzero element of A stays nonzero in T , and A is not bipartite. The above two facts combined ensure that the spectral radius $r(T)$ of T is strictly less than 1. We conclude that there exists a symmetric Markov chain T on state space $[q]^d$ that has both the properties: (i) $r(T) < 1$, and (ii) there is nonzero probability of transition between two multisets only when their support is disjoint. \square

3.3 Proof of Theorem 9

Let d be the constant for which Conjecture 3 is true. Thus, Conjecture 4 is true for the same value d as well. Choose q, T from Lemma 10 such that T is a symmetric Markov chain on $[q]^d$ such that $r(T) < 1$.

We now reduce the given d -to- d Label Cover instance to the problem of finding independent sets in q -colorable graphs. To be precise, given a Label Cover instance $G = ((V, E), dR, \Psi)$, we output a graph $G' = (V', E')$ such that

1. Completeness: If G is satisfiable, G' can be colored with q colors.
2. Soundness: If $isat_t(G) < \epsilon'$, then G' does not have any independent set of size ϵ .

The parameters t and ϵ' will be set later.

Reduction. Our reduction follows the standard Label Cover Long Code paradigm, and in particular closely mirrors [DMR09]. We replace each vertex $w \in V$ of the Label Cover with a set f_w of $[q]^{dR}$ nodes, each corresponding to a vertex in G' . Consider an edge $e = (u, v)$ where Ψ_e is an associated constraint with permutations π_1, π_2 on $[dR]$ such that $(a, b) \in \Psi_e$ if and only if $(\pi_1^{-1}(a), \pi_2^{-1}(b)) \in d \leftrightarrow d$.

We add an edge between $(x_1, x_2, \dots, x_{dR}) \in f_u$ and $(y_1, y_2, \dots, y_{dR}) \in f_v$ to E' if and only if

$$\forall i \in [R], T((x_{\pi_1(di-d+1)}, x_{\pi_1(di-d+2)}, \dots, x_{\pi_1(di)}) \leftrightarrow (y_{\pi_2(di-d+1)}, y_{\pi_2(di-d+2)}, \dots, y_{\pi_2(di)})) > 0$$

Completeness. Suppose $\sigma : V \rightarrow [dR]$ be a labeling satisfying all the constraints of the Label Cover instance G . We color the node $(x_1, x_2, \dots, x_{dR}) \in f_w$ with $x_{\sigma(w)} \in [q]$. We claim that this is a legit q -coloring of G' . Suppose that we added an edge between $x \in f_u$ and $y \in f_v$. Let x be colored with x_a and y be colored with y_b . As $(a, b) \in \Psi_{(u,v)}$, we have $(\pi_1^{-1}(a), \pi_2^{-1}(b)) \in d \leftrightarrow d$. Thus, there exist $i \in [R], 1 \leq p, q \leq d$ such that $a = \pi_1(di - d + p)$ and $b = \pi_2(di - d + q)$. As we have added an edge between $x \in f_u$ and $y \in f_v$, $x_a \neq y_b$ as the Markov chain T has nonzero probability only between two elements of $[q]^d$ with disjoint support. Thus, there exists a q -coloring of G' when G is satisfiable.

Soundness. We prove the contrapositive that if G' has an independent set of relative size ϵ , then there exists a labeling of G with $isat_t(G) \geq \epsilon'$. Let $S \subseteq V'$ be the largest independent set of G' . We know that $|S| \geq \epsilon|V'|$. This implies that in at least $\epsilon' = \frac{\epsilon}{2}$ fraction of the long code blocks, at least $\frac{\epsilon}{2}$ fraction of nodes belong to S . Let this subset of V be denoted by Z . Our goal is to show that there exists a small set of labels $\tau : Z \rightarrow 2^{[dR]}$ to which we can decode the vertices in Z such that all the constraints induced in Z are satisfied by τ .

For every vertex $w \in Z$, we define functions $g_w : [q]^{dR} \rightarrow \{0, 1\}$ to be the indicator functions of set S inside the long code blocks corresponding to w i.e. $g_w(x) = 1$ if and only if $x \in S$. Consider an edge $e = (u, v)$ corresponding to the constraint Ψ_e induced in Z . Let the functions $f : [q]^{dR} \rightarrow \{0, 1\}$ and $g : [q]^{dR} \rightarrow \{0, 1\}$ be defined such that $f(x^{\pi_1}) = g_u(x)$ and $g(y^{\pi_2}) = g_v(y)$, where π_1 and π_2 are the permutations underlying the relation Ψ_e i.e. $(a, b) \in \Psi_e$ if and only if $(\pi_1^{-1}(a), \pi_2^{-1}(b)) \in d \leftrightarrow d$.

We note that $\langle f, Tg \rangle$ is equal to zero. In other words, suppose that $x, y \in [q]^{dR}$, $x \in f_u, y \in f_v$ are such that

$$\forall i \in [R], T((x_{di-d+1}, x_{di-d+2}, \dots, x_{di}) \leftrightarrow (y_{di-d+1}, y_{di-d+2}, \dots, y_{di})) > 0. \quad (1)$$

Then, $f(x)g(y) = 0$. Suppose for contradiction that there exist $x, y \in [q]^{dR}$ satisfying the above condition, and $f(x) = g(y) = 1$. Let $x' \in f_u, y' \in f_v$ be such that $(x')^{\pi_1} = x, (y')^{\pi_2} = y$. We have $g_u(x') = g_v(y') = 1$. That is, both $x' \in f_u, y' \in f_v$ are in the independent set S . However, Equation (1) can be rewritten as the following:

$$\forall i \in [R], T((x'_{\pi_1(di-d+1)}, x'_{\pi_1(di-d+2)}, \dots, x'_{\pi_1(di)}) \leftrightarrow (y'_{\pi_2(di-d+1)}, y'_{\pi_2(di-d+2)}, \dots, y'_{\pi_2(di)})) > 0. \quad (2)$$

Note that this is precisely the condition for adding edges in G' . Thus, Equation (2) implies that $x' \in f_u$ and $y' \in f_v$ are adjacent in E' , and thus cannot both be part of the independent set S . This completes the proof that $\langle f, Tg \rangle = 0$.

Thus, $\langle \bar{f}, T\bar{g} \rangle$ is also equal to zero, where $\bar{f} : [q^d]^R \rightarrow \{0, 1\}$ and $\bar{g} : [q^d]^R \rightarrow \{0, 1\}$ are the corresponding functions in $[q^d]^R$ of f, g . From the definition of Z , $\mathbb{E}(\bar{f}) \geq \frac{\epsilon}{2}$ and $\mathbb{E}(\bar{g}) \geq \frac{\epsilon}{2}$. We apply Theorem 7 to \bar{f} and \bar{g} to deduce that there exists $i \in [R]$, a positive integer $k = k(\epsilon)$ and $\delta = \delta(\epsilon)$ such that $I_i^{\leq k}(\bar{f}) \geq \delta$ and $I_i^{\leq k}(\bar{g}) \geq \delta$. This motivates us to define the label set of vertex $w \in Z$, $L(w)$ as the following -

$$L(w) := \{i \in [dR] : I_i^{\leq dk}(g_w) \geq \frac{\delta}{d}\}$$

As the sum of k degree influences of all variables is at most k , the size of $L(w)$ is upper bounded by $\frac{kd}{\delta}$ for every v . Thus, we set the parameter t to be $\frac{kd}{\delta}$.

The final step is to prove that the labeling L is indeed a valid labeling inside edges induced in Z . Consider an edge $e = (u, v)$ induced in Z with the constraint relation being Ψ_e such that $(a, b) \in \Psi_e$ if and only if $(\pi_1(a), \pi_2(b)) \in d \leftrightarrow d$. Our goal is to show that there exist indices $\sigma_1, \sigma_2 \in [dR]$ such that $\sigma_1 \in L(u), \sigma_2 \in L(v)$ and $(\sigma_1, \sigma_2) \in \Psi_e$. Using Theorem 7, we can deduce that there exists $i \in [R]$ such that $I_i^{\leq k}(\bar{f}) \geq \delta$ and $I_i^{\geq k}(\bar{g}) \geq \delta$. Using Lemma 6, we can conclude that there exist $i_1, i_2 \in [dR]$ such that $I_{i_1}^{\leq dk}(f) \geq \frac{\delta}{d}$ and $I_{i_2}^{\leq dk}(g) \geq \frac{\delta}{d}$ such that $(i_1, i_2) \in d \leftrightarrow d$. Let $\sigma_1, \sigma_2 \in [dR]$ be such that $i_1 = \pi_1(\sigma_1), i_2 \in \sigma_2$. As $f(x^{\pi_1}) = g_u(x)$, $I_{\pi_1^{-1}(i_1)}^{\leq dk}(g_u) \geq \frac{\delta}{d}$. And thus, $\sigma_1 \in L(u)$, and similarly $\sigma_2 \in L(v)$. As $(i_1, i_2) \in d \leftrightarrow d$, $(\sigma_1, \sigma_2) \in \Psi_e$, which completes the proof.

4 Conclusion

In this paper, we prove that the d -to-1 conjecture, for arbitrarily large d , implies the NP-hardness of the longstanding and elusive problem of coloring 4-colorable graphs with constantly many colors. Note that the d -to-1 conjecture requires the soundness parameter to be arbitrarily small, independent of d . Currently, the best NP-hardness of d -to-1 Label Cover achieves a soundness of $d^{-\Omega(1)}$. This follows from the PCP Theorem [ALM⁺98, AS98] combined with Raz's parallel repetition [Raz98]. However, this does not yield any explicit constant in the exponent, obtaining which is an interesting open question. One can also investigate whether improving the soundness of d -to-1 Label Cover to something quantitatively much stronger, say inverse exponential in d , would have some implications for inapproximability of graph coloring.

References

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
- [BKO19] Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 602–613, 2019.
- [CD72] J Csima and B.N Datta. The DAD theorem for symmetric non-negative matrices. *Journal of Combinatorial Theory, Series A*, 12(1):147 – 152, 1972.
- [DKK⁺18a] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 940–951, 2018.
- [DKK⁺18b] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 376–389, 2018.
- [DMR09] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009.
- [DS05] Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–485, 2005.
- [Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 767–775, 2002.
- [KKMO07] 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.
- [KMS17] Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and Grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 576–589, 2017.
- [KMS18] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in Grassmann graph have near-perfect expansion. In *59th IEEE Annual Symposium on Foundations of Computer Science*, pages 592–601, 2018.
- [KT17] Ken-Ichi Kawarabayashi and Mikkel Thorup. Coloring 3-colorable graphs with less than $n^{1/5}$ colors. *J. ACM*, 64(1), March 2017.

- [MOO10] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, 171(1):295–341, 2010.
- [OW09] Ryan O’Donnell and Yi Wu. Conditional hardness for satisfiable 3-CSPs. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 493–502, 2009.
- [PR81] S. Poljak and V. Rödl. On the arc-chromatic number of a digraph. *Journal of Combinatorial Theory, Series B*, 31(2):190 – 198, 1981.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [SK67] Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.
- [WZ19] Marcin Wrochna and Stanislav Zivny. Improved hardness for H-colourings of G-colourable graphs. *CoRR*, abs/1907.00872, 2019.

A Reducing multigraph (exact) d -to-1 to $(d + 1)$ -to-1 conjecture

For the version of d -to-1 conjecture where we only require the constraint maps to be at most d -to-1, the d -to-1 conjecture trivially implies the $(d + 1)$ -to-1 conjecture. O’Donnell and Wu [OW09] remark that no such reduction appears to be known for the exact d -to-1 conjecture. Here we prove that the exact d -to-1 conjecture implies the exact $(d + 1)$ -to-1 conjecture when the underlying Label Cover instances are allowed to have parallel edges. We remark that multigraph version of exact d -to-1 conjecture, which is implied by the simple graph version, also suffices for our reduction to graph coloring (and indeed all known reductions from d -to-1 Label Cover).

Let $G = ((V = X \cup Y, E), (dR, R), \Psi)$ be a Label Cover instance such that every constraint is of d -to-1 structure. We reduce it to $G' = ((V = X \cup Y, E'), ((d + 1)R, R), \Psi')$ such that

1. If G is satisfiable, G' is satisfiable as well.
2. If every labeling violates at least ϵ fraction of constraints in G , then every labeling violates at least $\epsilon' = 2\epsilon$ fraction of constraints in G' .

Reduction. We first change the label set of X from $[dR]$ to $[(d + 1)R]$. For every constraint ψ in G between nodes $u \in X$ and $v \in Y$, we replace it with R constraints $\psi_1, \psi_2, \dots, \psi_R$ between u and v in the following way: the relation between old labels is the same as ψ i.e. when $x \leq dR$, $(x, y) \in \psi_j$ for $j = 1, 2, \dots, R$ if and only if $(x, y) \in \psi$. When $x > dR$, $(x, y) \in \psi_j$ if and only if R divides $(x + j - y)$. This ensures that each new label is mapped to a different label in each of the R new constraints. The constraints are clearly of $(d + 1)$ -to-1 form.

Completeness. If there is a labeling satisfying all the constraints of G , the same labeling satisfies all the constraints in G' as well.

Soundness. Suppose that there is no labeling satisfying at least ϵ fraction of constraints in G . Note that this implies that R is at least $\frac{1}{\epsilon}$ as there is always a labeling satisfying at least $\frac{1}{R}$ fraction of constraints: fix a labeling to the vertices on the left, and assign a label to the vertices in R uniformly at random from $[R]$. We claim that there is no labeling satisfying more than 2ϵ fraction of constraints in G' . Consider an arbitrary labeling of G , $\sigma : V \rightarrow [(d+1)R]$. We can divide the set of edges E' of G' into two parts: the edges (u, v) such that $\sigma(u) \leq dR$ and the edges (u, v) such that $\sigma(u) > dR$. Let the set of first type of edges where the left vertex is assigned the new label be denoted by E_1 , and the set of second type of edges be denoted by E_2 . In E_1 , the fraction of constraints that can be satisfied by σ is at most $\frac{1}{R} \leq \epsilon$. Note that we can get a labeling σ' of G by replacing labels of vertices in X with label greater than dR with an arbitrary label in $[dR]$, and keeping rest of the labels intact. For the edges in E_2 , the labelings σ and σ' coincide. As σ' can satisfy at most ϵ fraction of constraints of G , σ can only satisfy at most ϵ fraction of overall edges in E' . Thus, overall σ satisfies at most $\epsilon + \frac{1}{R} \leq 2\epsilon$ fraction of constraints in E' , which proves the required soundness claim.