

Approximate Hypergraph Coloring under Low-discrepancy and Related Promises

Vijay V. S. P. Bhattiprolu* Venkatesan Guruswami† Euiwoong Lee‡

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213.

Abstract

A hypergraph is said to be χ -colorable if its vertices can be colored with χ colors so that no hyperedge is monochromatic. 2-colorability is a fundamental property (called Property B) of hypergraphs and is extensively studied in combinatorics. Algorithmically, however, given a 2-colorable k -uniform hypergraph, it is NP-hard to find a 2-coloring miscoloring fewer than a fraction 2^{-k+1} of hyperedges (which is trivially achieved by a random 2-coloring), and the best algorithms to color the hypergraph properly require $\approx n^{1-1/k}$ colors, approaching the trivial bound of n as k increases.

In this work, we study the complexity of approximate hypergraph coloring, for both the maximization (finding a 2-coloring with fewest miscolored edges) and minimization (finding a proper coloring using fewest number of colors) versions, when the input hypergraph is promised to have the following stronger properties than 2-colorability:

- Low-discrepancy: If the hypergraph has a 2-coloring of discrepancy $\ell \ll \sqrt{k}$, we give an algorithm to color the hypergraph with $\approx n^{O(\ell^2/k)}$ colors. However, for the maximization version, we prove NP-hardness of finding a 2-coloring miscoloring a smaller than $2^{-O(k)}$ (resp. $k^{-O(k)}$) fraction of the hyperedges when $\ell = O(\log k)$ (resp. $\ell = 2$). Assuming the Unique Games conjecture, we improve the latter hardness factor to $2^{-O(k)}$ for almost discrepancy-1 hypergraphs.
- Rainbow colorability: If the hypergraph has a $(k - \ell)$ -coloring such that each hyperedge is polychromatic with all these colors (this is stronger than a $(\ell + 1)$ -discrepancy 2-coloring), we give a 2-coloring algorithm that miscolors at most $k^{-\Omega(k)}$ of the hyperedges when $\ell \ll \sqrt{k}$, and complement this with a matching Unique Games hardness result showing that when $\ell = \sqrt{k}$, it is hard to even beat the 2^{-k+1} bound achieved by a random coloring.

*Supported by NSF CCF-1115525 vpb@cs.cmu.edu

†Supported in part by NSF grant CCF-1115525. guruswami@cmu.edu

‡Supported by a Samsung Fellowship and NSF CCF-1115525. euiwoon1@cs.cmu.edu

1. Introduction

Coloring (hyper)graphs is one of the most important and well-studied tasks in discrete mathematics and theoretical computer science. A k -uniform hypergraph $G = (V, E)$ is said to be χ -colorable if there exists a coloring $c : V \mapsto \{1, \dots, \chi\}$ such that no hyperedge is monochromatic, and such a coloring c is referred to as a proper χ -coloring. Graph and hypergraph coloring has been the focus of active research in both fields, and has served as the benchmark for new research paradigms such as the probabilistic method (Lovász local lemma [EL75]) and semidefinite programming (Lovász theta function [Lov79]).

While such structural results are targeted towards special classes of hypergraphs, given a general χ -colorable k -uniform hypergraph, the problem of reconstructing a χ -coloring is known to be a hard task. Even assuming 2-colorability, reconstructing a proper 2-coloring is a classic NP-hard problem for $k \geq 3$. Given the intractability of proper 2-coloring, two notions of *approximate coloring* of 2-colorable hypergraphs have been studied in the literature of approximation algorithms. The first notion, called *Min-Coloring*, is to minimize the number of colors while still requiring that every hyperedge be non-monochromatic. The second notion, called *Max-2-Coloring* allows only 2 colors, but the objective is to maximize the number of non-monochromatic hyperedges.¹

Even with these relaxed objectives, the promise that the input hypergraph is 2-colorable seems grossly inadequate for polynomial time algorithms to exploit in a significant way. For Min-Coloring, given a 2-colorable k -uniform hypergraph, the best known algorithm uses $O(n^{1-\frac{1}{k}})$ colors [CF96, AKMH96], which tends to the trivial upper bound n as k increases. This problem has been actively studied from the hardness side, motivating many new developments in constructions of probabilistically checkable proofs. Coloring 2-colorable hypergraphs with $O(1)$ colors was shown to be NP-hard for $k \geq 4$ in [GHS02] and $k = 3$ in [DRS05]. An exciting body of recent work has pushed the hardness beyond poly-logarithmic colors [DG13, GHH⁺14, KS14, Hua15]. In particular, [KS14] shows quasi-NP-hardness of $2^{(\log n)^{\Omega(1)}}$ -coloring a 2-colorable hypergraphs (very recently the exponent was shown to approach 1/4 in [Hua15]).

The hardness results for Max-2-Coloring show an even more pessimistic picture, wherein the naive *random assignment* (randomly give one of two colors to each vertex independently to leave a $(\frac{1}{2})^{k-1}$ fraction of hyperedges monochromatic in expectation), is shown to have the best guarantee for a polynomial time algorithm when $k \geq 4$ (see [Hås01]).

Given these strong intractability results, it is natural to consider what further relaxations of the objectives could lead to efficient algorithms. For maximization versions, Austrin and Håstad [AH13] prove that (almost²) 2-colorability is useless (in a formal sense that they define) for any Constraint Satisfaction Problem (CSP) that is a relaxation of 2-coloring [Wen14]. Therefore, it seems more natural to find a stronger promise on the hypergraph than mere 2-colorability that can be significantly exploited by polynomial time coloring algorithms for the objectives of Min-Coloring and Max 2-Coloring. This motivates our main question “*how strong a promise on the input hypergraph is required for polynomial time algorithms to perform significantly better than naive algorithms for Min-Coloring and Max-2-Coloring?*”

There is a very strong promise on k -uniform hypergraphs which makes the task of proper 2-coloring easy. If a hypergraph is k -partite (i.e., there is a k -coloring such that each hyperedge has each color exactly once), then one can properly 2-color the hypergraph in polynomial time. The same algorithm can be generalized to hypergraphs which admit a c -balanced coloring (i.e., c divides k and there is a k -coloring such that each hyperedge has each color exactly $\frac{k}{c}$ times). This can be seen by random hyperplane rounding of a simple SDP, or even simpler by solving a homogeneous linear system and iterating [Alo14], or by a random recoloring method analyzed using random walks [McD93]. In fact, a proper 2-coloring can be efficiently

¹The maximization version is also known as Max-Set-Splitting, or more specifically Max k -Set-Splitting when considering k -uniform hypergraphs, in the literature.

²We say a hypergraph is *almost* χ -colorable for a small constant $\epsilon > 0$, there is a χ -coloring that leaves at most ϵ fraction of hyperedges monochromatic.

achieved assuming that the hypergraph admits a fair partial 2-coloring, namely a pair of disjoint subsets A and B of the vertices such that for every hyperedge e , $|e \cap A| = |e \cap B| > 0$ [McD93].

The promises on structured colorings that we consider in this work are natural relaxations of the above strong promise of a perfectly balanced (partial) coloring.

- A hypergraph is said to have *discrepancy* ℓ when there is a 2-coloring such that in each hyperedge, the difference between the number of vertices of each color is at most ℓ .
- A χ -coloring ($\chi \leq k$) is called *rainbow* if every hyperedge contains each color at least once.
- A χ -coloring ($\chi \geq k$) is called *strong* if every hyperedge contains k different colors.

These three notions are interesting in their own right, and have been independently studied. Discrepancy minimization has recently seen different algorithmic ideas [Ban10, LM12, Rot14] to give constructive proofs of the classic six standard deviations result of Spencer [Spe85]. Rainbow coloring admits a natural interpretation as a partition of V into the maximum number of disjoint vertex covers, and has been actively studied for geometric hypergraphs due to its applications in sensor networks [BPRS13]. Strong coloring is closely related to graph coloring by definition, and is known to capture various other notions of coloring [AH05]. It is easy to see that ℓ -discrepancy ($\ell < k$), χ -rainbow colorability ($2 \leq \chi \leq k$), and χ -strong colorability ($k \leq \chi \leq 2k - 2$) all imply 2-colorability. For odd k , both $(k + 1)$ -strong colorability and $(k - 1)$ -rainbow colorability imply discrepancy-1, so strong colorability and rainbow colorability seem stronger than low discrepancy.

Even though they seem very strong, previous works have mainly focused on hardness with these promises. The work of Austrin et al. [AGH14] shows NP-hardness of finding a proper 2-coloring under the discrepancy-1 promise. The work of Bansal and Khot [BK10] shows hardness of $O(1)$ -coloring even when the input hypergraph is promised to be almost k -partite (under the Unique Games Conjecture); Sachdeva and Saket [SS13] establish NP-hardness of $O(1)$ -coloring when the graph is almost $k/2$ -rainbow colorable; and Guruswami and Lee [GL15a] establish NP-hardness when the graph is perfectly (not almost) $\frac{k}{2}$ -rainbow colorable, or admits a 2-coloring with discrepancy 2. These hardness results indicate that it is still a nontrivial task to exploit these strong promises and outperform naive algorithms.

1.1. Our Results

In this work, we prove that our three promises, unlike mere 2-colorability, give enough structure for polynomial time algorithms to perform significantly better than naive algorithms. We also study these promises from a hardness perspective to understand the asymptotic threshold at which beating naive algorithms goes from easy to UG/NP-Hard. In particular assuming the UGC, for Max-2-Coloring under ℓ -discrepancy or $k - \ell$ -rainbow colorability, this threshold is $\ell = \Theta(\sqrt{k})$.

Theorem 1.1. *There is a randomized polynomial time algorithm that produces a 2-coloring of a k -uniform hypergraph H with the following guarantee. For any $0 < \varepsilon < \frac{1}{2}$ (let $\ell = k^\varepsilon$), there exists a constant $\eta > 0$ such that if H is $(k - \ell)$ -rainbow colorable or $(k + \ell)$ -strong colorable, the fraction of monochromatic edges in the produced 2-coloring is $O((\frac{1}{k})^{\eta k})$ in expectation.*

Our results indeed show that this algorithm significantly outperforms the random assignment even when ℓ approaches \sqrt{k} asymptotically. See Theorem 2.10 and Theorem 2.16 for the precise statements.

For the ℓ -discrepancy case, we observe that when $\ell < \sqrt{k}$, the framework of the second and the third authors [GL15b] yields an approximation algorithm that marginally (by an additive factor much less than 2^{-k}) outperforms the random assignment, but we do not formally prove this here.

The following hardness results suggest that this gap between low-discrepancy and rainbow/strong colorability might be intrinsic. Let the term *UG-hardness* denote NP-hardness assuming the Unique Games Conjecture.

Theorem 1.2. *For sufficiently large odd k , given a k -uniform hypergraph which admits a 2-coloring with at most a $(\frac{1}{2})^{6k}$ fraction of edges of discrepancy larger than 1, it is UG-hard to find a 2-coloring with a $(\frac{1}{2})^{5k}$ fraction of monochromatic edges.*

Theorem 1.3. *For even $k \geq 4$, given a k -uniform hypergraph which admits a 2-coloring with no edge of discrepancy larger than 2, it is NP-hard to find a 2-coloring with a $k^{-O(k)}$ fraction of monochromatic edges.*

Theorem 1.4. *For k sufficiently large, given a k -uniform hypergraph which admits a 2-coloring with no edge of discrepancy larger than $O(\log k)$, it is NP-hard to find a 2-coloring with a $2^{-O(k)}$ fraction of monochromatic edges.*

Theorem 1.5. *For k such that $\chi := k - \sqrt{k}$ is an integer greater than 1, and any $\varepsilon > 0$, given a k -uniform hypergraph which admits a χ -coloring with at most ε fraction of non-rainbow edges, it is UG-hard to find a 2-coloring with a $(\frac{1}{2})^{k-1}$ fraction of monochromatic edges.*

For Min-Coloring, all three promises lead to an $\tilde{O}(n^{\frac{1}{k}})$ -coloring that is decreasing in k . These results are also notable in the sense that our promises are helpful not only for structured SDP solutions, but also for combinatorial *degree reduction* algorithms.

Theorem 1.6. *Consider any k -uniform hypergraph $H = (V, E)$ with n vertices and m edges. For any $\ell < O(\sqrt{k})$, If H has discrepancy- ℓ , $(k - \ell)$ -rainbow colorable, or $(k + \ell)$ -strong colorable, one can color H with $\tilde{O}(\frac{m}{n} \frac{\ell^2}{k^2}) \leq \tilde{O}(n^{\frac{\ell^2}{k}})$ colors.*

These results significantly improve the current best $\tilde{O}(n^{1-\frac{1}{k}})$ colors that assumes only 2-colorability. Our techniques give slightly better results depending on the promise — see Theorem 4.1. Table 1.1 summarizes our results.

Promises	ℓ -Discrepancy	$(k - \ell)$ -Rainbow	$(k + \ell)$ -Strong
Max-2-Coloring Algorithm	$1 - (1/2)^{k-1} + \delta, \quad \ell < \sqrt{k}^\dagger$	$1 - (1/k)^{\Omega(k)}, \quad \ell \ll \sqrt{k}$	$1 - (1/k)^{\Omega(k)}, \quad \ell \ll \sqrt{k}$
Max-2-Coloring Hardness	UG: $1 - (1/2)^{5k}, \quad \ell = 1.$ NP: $1 - (1/k)^{O(k)}, \quad \ell = 2.$ NP: $1 - (1/2)^{O(k)}, \quad \ell = \Omega(\log k)$ UG: $1 - (1/2)^{k-1}, \quad \ell \geq \sqrt{k}^\dagger$	UG: $1 - (1/2)^{k-1}, \quad \ell = \Omega(\sqrt{k})$	
Min-Coloring Algorithm	$n^{\ell^2/k}, \quad \ell = O(\sqrt{k})$	$n^{\ell^2/k}, \quad \ell = O(\sqrt{k})$	$n^{\ell^2/k}, \quad \ell = O(\sqrt{k})$

Table 1.1: Summary of our algorithmic and hardness results with valid ranges of ℓ . Two results with \dagger are implied in [GL15b]. The numbers of the first row indicate lower bounds on the fraction of non-monochromatic edges in a 2-coloring produced by our algorithms. $\delta := \delta(k, \ell) > 0$ is a small constant. The second row shows upper bounds on the fraction of non-monochromatic edges achieved by polynomial time algorithms. For the UG-hardness results, note that the input hypergraph does not have all edges satisfying the promises but almost edges satisfying them. The third row shows the upper bound upto log factors, on the number of colors one can use to properly 2-color the graph.

1.2. Techniques

Our algorithms for Max-2-Coloring are straightforward applications of semidefinite programming, namely, we use natural vector relaxations of the promised properties, and round using a random hyperplane. The analysis however, is highly non-trivial and boils down to approximating a multivariate Gaussian integral. In particular, we show a (to our knowledge, new) upper bound on the Gaussian measure of simplicial cones in terms of simple properties of these cones. We should note that this upper bound is sensible only for simplicial cones that are well behaved with respect to these properties. (The cones we are interested in are those given by the intersection of hyperplanes whose normal vectors constitute a solution to our vector relaxations). We believe our analysis to be of independent interest as similar approaches may work for other k -CSP's.

1.2.1. Gaussian Measure of Simplicial Cones

As can be seen via an observation of Kneser [Kne36], the Gaussian measure of a simplicial cone is equal to the fraction of spherical volume taken up by a spherical simplex (a spherical simplex is the intersection of a simplicial cone with a ball centered at the apex of the cone). This however, is a very old problem in spherical geometry, and while some things are known, like a nice differential formula due to Schlafli (see [Sch58]), closed forms upto four dimensions (see [MY05]), and a complicated power series expansion due to Aomoto [A⁺77], it is likely hopeless to achieve a closed form solution or even an asymptotic formula for the volume of general spherical simplices.

Zwick [Zwi98] considered the performance of hyperplane rounding in various 3-CSP formulations, and this involved analyzing the volume of a 4-dimensional spherical simplex. Due to the complexity of this volume function, the analysis was tedious, and non-analytic for many of the formulations. His techniques were based on the Schlafli differential formula, which relates the volume differential of a spherical simplex to the volume functions of its codimension-2 faces and dihedral angles. However, to our knowledge not much is known about the general volume function in even 6 dimensions. This suggests that Zwick's techniques are unlikely to be scalable to higher dimensions.

On the positive side, an asymptotic expression is known in the case of symmetric spherical simplices, due to H. E. Daniels [Rog64] who gave the analysis for regular cones of angle $\cos^{-1}(1/2)$. His techniques were extended by Rogers [Rog61] and Boeroeczky and Henk [BJH99] to the whole class of regular cones.

We combine the complex analysis techniques employed by Daniels with a lower bound on quadratic forms in the positive orthant, to give an upper bound on the Gaussian measure of a much larger class of simplicial cones.

1.2.2. Column Subset Selection

Informally, the cones for which our upper bound is relevant are those that are high dimensional in a strong sense, i.e. the normal vectors whose corresponding hyperplanes form the cone, must be such that no vector is too close to the linear span of any subset of the remaining vectors.

When the normal vectors are solutions to our rainbow colorability SDP relaxation, this need not be true. However, this can be remedied. We consider the column matrix of these normal vectors, and using spectral techniques, we show that there is a reasonably large subset of columns (vectors) that are well behaved with respect to condition number. We are then able to apply our Gaussian Measure bound to the cone given by this subset, admittedly in a slightly lower dimensional space.

2. Approximate Max-2-Coloring

In this section we show how the properties of $(k + \ell)$ -strong colorability and $(k - \ell)$ -rainbow colorability in k -uniform hypergraphs allow one to 2-color the hypergraph, such that the respective fractions of monochromatic edges are small. For $\ell = o(\sqrt{k})$, these guarantees handsomely beat the naive random algorithm (color every vertex blue or red uniformly and independently at random), wherein the expected fraction of monochromatic edges is $1/2^{k-1}$.

Our algorithms are straightforward applications of semidefinite programming, namely, we use natural vector relaxations of the above properties, and round using a random hyperplane. The analysis however, is quite involved.

2.1. Semidefinite Relaxations

Our SDP relaxations for low-discrepancy, rainbow-colorability, and strong-colorability are the following. Given that $\langle v_i, v_j \rangle = \frac{-1}{\chi-1}$ when unit vectors v_1, \dots, v_χ form a χ -regular simplex centered at the origin, it is easy to show that they are valid relaxations.

Discrepancy ℓ .

$$\begin{aligned} \forall e \in E, \quad & \left\| \sum_{i \in e} u_i \right\|_2 \leq \ell \\ \forall i \in [n], \quad & \|u_i\|_2 = 1 \\ \forall i \in [n], \quad & u_i \in \mathbb{R}^n \end{aligned} \tag{2.1}$$

Feasibility. For k, ℓ such that $(k - \ell) \bmod 2 \equiv 0$, consider any k -uniform hypergraph $H = (V = [n], E)$, and any 2-coloring of H of discrepancy ℓ . Pick any unit vector $w \in \mathbb{R}^n$. For each vertex of the first color in the coloring, assign the vector w , and for each vertex of the second color assign the vector $-w$. This is a feasible assignment, and hence Relaxation 2.1 is a feasible relaxation for any hypergraph of discrepancy ℓ .

$(k - \ell)$ -Rainbow Colorability.

$$\begin{aligned} \forall e \in E, \quad & \left\| \sum_{i \in e} u_i \right\|_2 \leq \ell \\ \forall e \in E, \forall i < j \in e, \quad & \langle u_i, u_j \rangle \geq \frac{-1}{k - \ell - 1} \\ \forall i \in [n], \quad & \|u_i\|_2 = 1 \\ \forall i \in [n], \quad & u_i \in \mathbb{R}^n \end{aligned} \tag{2.2}$$

Feasibility. Consider any k -uniform hypergraph $H = (V = [n], E \subseteq \binom{V}{k})$, and any $(k - \ell)$ -rainbow coloring of H . As testified by the vertices of the $(k - \ell)$ -simplex, we can always choose unit vectors $w_1 \dots w_{k-\ell} \in \mathbb{R}^n$ satisfying,

$$\forall i < j \in [k - \ell], \quad \langle w_i, w_j \rangle = \frac{-1}{k - \ell - 1},$$

It is not hard to verify that consequently,

$$\forall a_1, \dots, a_{k-\ell} \in [l], \quad \sum_{i \in [k-\ell]} a_i = k, \quad \text{we have,} \quad \left\| \sum_{i \in e} a_i w_i \right\|_2 \leq \ell$$

For each vertex of the color i , assign the vector w_i . This is a feasible assignment, and hence Relaxation 2.2 is a feasible relaxation for any hypergraph of rainbow colorability $k - \ell$.

$(k + \ell)$ -Strong Colorability.

$$\begin{aligned} \forall e \in E, \forall i < j \in e, \quad \langle u_i, u_j \rangle &= \frac{-1}{k + \ell - 1} \\ \forall i \in [n], \quad \|u_i\|_2 &= 1 \\ \forall i \in [n], \quad u_i &\in \mathbb{R}^n \end{aligned} \tag{2.3}$$

Feasibility. Consider any k -uniform hypergraph $H = (V = [n], E \subseteq \binom{V}{k})$, and any $(k + \ell)$ -strong coloring of H . As testified by the vertices of the $(k + \ell)$ -simplex, we can always choose unit vectors $w_1 \dots w_{k+\ell} \in \mathbb{R}^n$ satisfying,

$$\forall i < j \in [k + \ell], \langle w_i, w_j \rangle = -\frac{1}{k + \ell - 1},$$

It is not hard to verify that consequently,

$$\forall J \subset [k + \ell], |J| = k, \left\| \sum_{i \in J} w_i \right\|_2 = \ell$$

For each vertex of the color i , assign the vector w_i . This is a feasible assignment, and hence the Relaxation 2.3 is a feasible relaxation for any hypergraph of strong colorability $k + \ell$.

Our rounding scheme is the same for all the above relaxations.

Rounding Scheme. Pick a standard n -dimensional Gaussian random vector r . For any $i \in [n]$, if $\langle v_i, r \rangle \geq 0$, then vertex i is colored blue, and otherwise it is colored red.

2.2. Setup of Analysis

We now setup the framework for analyzing all the above relaxations.

Consider a standard n -dimensional Gaussian random vector r , i.e. each coordinate is independently picked from the standard normal distribution $\mathcal{N}(0, 1)$. The following are well known facts (the latter being due to Renyi),

Lemma 2.1. $r / \|r\|_2$ is uniformly distributed over the unit sphere in \mathbb{R}^n .

Note. Lemma 2.1 establishes that our rounding scheme is equivalent to random hyperplane rounding.

Lemma 2.2. Consider any $j < n$. The projections of r onto the pairwise orthogonal unit vectors e_1, \dots, e_j are independent and have distribution $\mathcal{N}(0, 1)$.

Next, consider any k -uniform hypergraph $H = (V = [n], E \subseteq \binom{V}{k})$ that is feasible for any of the aforementioned formulations. Our goal now, is to analyze the expected number of monochromatic edges. To obtain this expected fraction with high probability, we need only repeat the rounding scheme polynomially

many times, and the high probability of a successful round follows by Markov's inequality. Thus we are only left with bounding the probability that a particular edge is monochromatic.

To this end, consider any edge $e \in E$ and let the vectors corresponding to the vertices in e be u'_1, \dots, u'_k . Consider a k -flat \mathcal{F} (subspace of \mathbb{R}^n congruent to \mathbb{R}^k), containing u'_1, \dots, u'_k . Applying Lemma 2.2 to the standard basis of \mathcal{F} , implies that the projection of r into \mathcal{F} has the standard k -dimensional Gaussian distribution. Now since projecting r onto $\text{Span}(u'_1, \dots, u'_k)$ preserves the inner products $\{\langle r, u'_i \rangle\}_i$, we may assume without loss of generality that u'_1, \dots, u'_k are vectors in \mathbb{R}^k , and the rounding scheme corresponds to picking a random k -dimensional Gaussian vector r , and proceeding as before.

Let U be the $k \times k$ matrix whose columns are the vectors u'_1, \dots, u'_k and μ represent the Gaussian measure in \mathbb{R}^k . Then the probability of e being monochromatic in the rounding is given by,

$$\mu\left(\left\{x \in \mathbb{R}^k \mid U^T x \geq 0\right\}\right) + \mu\left(\left\{x \in \mathbb{R}^k \mid U^T x < 0\right\}\right) = 2\mu\left(\left\{x \in \mathbb{R}^k \mid U^T x \geq 0\right\}\right) \quad (2.4)$$

In other words, this boils down to analyzing the Gaussian measure of the cone given by $U^T x \geq 0$. We thus take a necessary detour.

2.3. Gaussian Measure of Simplicial Cones

In this section we show how to bound the Gaussian measure of a special class of simplicial cones. This is one of the primary tools in our analysis of the previously introduced SDP relaxations. We first state some preliminaries.

2.3.1. Preliminaries

Simplicial Cones and Equivalent Representations. A simplicial cone in \mathbb{R}^k , is given by the intersection of a set of k linearly independent halfspaces. For any simplicial cone with apex at position vector p , there is a unique set (upto changes in lengths), of k linearly independent vectors, such that the direct sum of $\{p\}$ with their positive span produces the cone. Conversely, a simplicial cone given by the direct sum of $\{p\}$ and the positive span of k linearly independent vectors, can be expressed as the intersection of a unique set of k halfspaces with apex at p . We shall refer to the normal vectors of the halfspaces above, as simply normal vectors of the cone, and we shall refer to the spanning vectors above, as simplicial vectors. We represent a simplicial cone C with apex at p , as (p, U, V) where U is a column matrix of unit vectors u_1, \dots, u_k (normal vectors), V is a column matrix of unit vectors v_1, \dots, v_k (simplicial vectors) and

$$C = \left\{x \in \mathbb{R}^k \mid u_1^T x \geq p_1, \dots, u_k^T x \geq p_k\right\} = \left\{p + x_1 v_1 + \dots + x_k v_k \mid x \geq 0, x \in \mathbb{R}^k\right\}$$

Switching Between Representations. Let $C \equiv (0, U, V)$ be a simplicial cone with apex at the origin. It is not hard to see that any v_i is in the intersection of exactly $k - 1$ of the k halfspaces determined by U , and it is thus orthogonal to exactly $k - 1$ vectors of the form u_j . We may assume without loss of generality that for any v_i , the only column vector of U not orthogonal to it, is u_i . Thus clearly $V^T U = D$ where D is some non-singular diagonal matrix. Let $A_U = U^T U$ and $A_V = V^T V$, be the gram matrices of the vectors. A_U and A_V are positive definite symmetric matrices with diagonal entries equal to one (they comprise of the pairwise inner products of the normal and simplicial vectors respectively). Also, clearly,

$$V = U^{-T} D, \quad A_V = D A_U^{-1} D \quad (2.5)$$

One then immediately obtains: $(A_V)_{ij} = \frac{a_{ij}}{\sqrt{a_{ii} a_{jj}}}$, and $(A_U)_{ij} = \frac{-a'_{ij}}{\sqrt{a'_{ii} a'_{jj}}}$. where a_{ij} and a'_{ij} are the cofactors of the $(i, j)^{th}$ entries of A_U and A_V respectively.

Formulating the Integral. Let $C \equiv (0, U, V)$ be a simplicial cone with apex at the origin, and for $x \in \mathbb{R}^k$, let dx denote the differential of the standard k -dimensional Lebesgue measure. Then the Gaussian measure of C is given by,

$$\begin{aligned}
& \frac{1}{\pi^{k/2}} \int_{U^T x \geq 0} e^{-\|x\|_2^2} dx &= \frac{\det(V)}{\pi^{k/2}} \int_{\mathbb{R}_+^k} e^{-\|Vx\|_2^2} dx && \text{Subst. } x \leftarrow Vx \\
&= \frac{\det(V)}{\pi^{k/2}} \int_{\mathbb{R}_+^k} e^{-\|U^{-T}Dx\|_2^2} dx && \text{By Eq. (2.5)} &= \frac{\det(V)}{\pi^{k/2} \det(D)} \int_{\mathbb{R}_+^k} e^{-\|U^{-T}x\|_2^2} dx && \text{Subst. } x \leftarrow Dx \\
&= \frac{1}{\pi^{k/2} \det(U)} \int_{\mathbb{R}_+^k} e^{-\|U^{-T}x\|_2^2} dx &= \frac{1}{\pi^{k/2} \sqrt{\det(AU)}} \int_{\mathbb{R}_+^k} e^{-x^T A_U^{-1} x} dx
\end{aligned}$$

For future ease of use, we give a name to some properties.

Definition 2.3. The para-volume of a set of vectors (resp. a matrix U), is the volume of the parallelotope determined by the set of vectors (resp. the column vectors of U).

Definition 2.4. The sum-norm of a set of vectors (resp. a matrix U), is the length of the sum of the vectors (resp. the sum of the column vectors of U).

Walkthrough of Symmetric Case Analysis. We next state some simple identities that can be found in say, [Rog64], some of which were originally used by Daniels to show that the Gaussian measure of a symmetric cone in \mathbb{R}^k of angle $\cos^{-1}(1/2)$ (between any two simplicial vectors) is $\frac{(1+o(1))e^{k/2-1}}{\sqrt{2}^{k+1} \sqrt{k}^{k-1} \sqrt{\pi}^k}$. We state these identities, while loosely describing the analysis of the symmetric case, to give the reader an idea of their purpose.

First note that the gram matrices S_U and S_V , of the symmetric cone of angle $\cos^{-1}(1/2)$ are given by:

$$S_U = (1 + 1/k)\mathbf{I} - \mathbf{1}\mathbf{1}^T/k \quad S_V = (\mathbf{I} + \mathbf{1}\mathbf{1}^T)/2$$

Thus $x^T S_U^{-1} x$ is of the form,

$$\alpha \|x\|_1^2 + \beta \|x\|_2^2 \tag{2.6}$$

The key step is in linearizing the $\|x\|_1^2$ term in the exponent, which allows us to separate the terms in the multivariate integral into a product of univariate integrals, and this is easier to analyze.

Lemma 2.5 (Linearization). $\sqrt{\pi} e^{-s^2} = \int_{-\infty}^{\infty} e^{-t^2 + 2its} dt$

Observation 2.6. Let $f : (-\infty, \infty) \mapsto \mathbb{C}$ be a continuous complex function. Then, $\left| \int_{-\infty}^{\infty} f(t) dt \right| \leq \int_{-\infty}^{\infty} |f(t)| dt$.

On applying Lemma 2.5 to Eq. (2.3.1) in the symmetric case, one obtains a product of identical univariate complex integrals. Specifically, by Eq. (2.3.1), Eq. (2.6), and Lemma 2.5, we have the expression,

$$\int_{\mathbb{R}_+^k} e^{-\beta \|x\|_2^2 - \alpha \|x\|_1^2} dx = \int_{-\infty}^{\infty} e^{-t^2} \int_{\mathbb{R}_+^k} e^{-\beta(x_1^2 + \dots + x_k^2) + 2it\sqrt{\alpha}(x_1 + \dots + x_k)} dx dt = \int_{-\infty}^{\infty} e^{-t^2} \left(\int_0^{\infty} e^{-\beta s^2 + 2it\sqrt{\alpha}s} ds \right)^k$$

The inner univariate complex integral is not readily evaluable. To circumvent this, one can change the line of integration so as to shift mass from the inner integral to the outer integral. Then we can apply the crude upper bound of Observation 2.6 to the inner integral, and by design, the error in our estimate is small.

Lemma 2.7 (Changing line of integration). Let $g(t)$ be a real valued function for real t . If, when interpreted as a complex function in the variable $t = a + ib$, $g(a + ib)$ is an entire function, and furthermore, $\lim_{a \rightarrow \infty} g(a + ib) = 0$ for some fixed b , then we have, $\int_{-\infty}^{\infty} g(t) dt = \int_{-\infty}^{\infty} g(a + ib) da$.

Squared L_1 Inequality. Motivated by the above linearization technique, we prove the following lower bound on quadratic forms in the positive orthant:

Lemma 2.8. Consider any $k \times k$ matrix A , and $x \in \mathbb{R}_+^k$, such that x is in the column space of A . Let A^\dagger denote the Moore-Penrose pseudo-inverse of A . Then, $x^T A^\dagger x \geq \frac{\|x\|_1^2}{\text{sum}(A)}$.

Proof. Consider any x in the positive orthant and column space of A . Let v_1, \dots, v_q be the eigenvectors of A corresponding to its non-zero eigenvalues. We may express x in the form $x = \sum_i \beta_i v_i$, so that

$$\|x\|_1 = \langle \mathbf{1}, x \rangle = \sum_{i \in [q]} \beta_i \langle \mathbf{1}, v_i \rangle \Rightarrow \|x\|_1^2 = \left(\sum_{i \in [q]} \beta_i \langle \mathbf{1}, v_i \rangle \right)^2.$$

We also have

$$x^T A^\dagger x = x^T \left(\sum_{i \in [q]} \lambda_i^{-1} v_i v_i^T \right) x = \sum_{i \in [q]} \lambda_i^{-1} \beta_i^2.$$

Now by Cauchy-Schwartz,

$$\left(\sum_i \lambda_i \langle \mathbf{1}, v_i \rangle^2 \right) \left(\sum_{i \in [q]} \lambda_i^{-1} \beta_i^2 \right) \geq \|x\|_1^2.$$

Therefore, we have

$$x^T A^\dagger x \geq \frac{\|x\|_1^2}{\sum_{i \in [q]} \lambda_i \langle \mathbf{1}, v_i \rangle^2} = \frac{\|x\|_1^2}{\mathbf{1}^T A \mathbf{1}} = \frac{\|x\|_1^2}{\text{sum}(A)}. \quad \blacksquare$$

Equipped with all necessary tools, we may now prove our result.

2.3.2. Our Gaussian Measure Bound

Let $C \equiv (0, U, V)$ be a simplicial cone with apex at the origin. We now show an upper bound on the Gaussian measure of C that depends surprisingly on only the para-volume and sum-norm of U . Since Gaussian measure is at most 1, it is evident when viewing our bound that it can only be useful for simplicial cones wherein the sum-norm of their normal vectors is $O(\sqrt{k})$, and the para-volume of their normal vectors is not too small.

Theorem 2.9. Let $C \equiv (0, U, V)$ be a simplicial cone with apex at the origin. Let $\ell = \|\sum_i u_i\|_2$ (i.e. sum-norm of the normal vectors), then the Gaussian measure of C is at most $\left(\frac{e}{2\pi k}\right)^{k/2} \frac{\ell^k}{\sqrt{\det(A_U)}}$

Proof. By the sum-norm property, the sum of entries of A_U is ℓ^2 . Also by the definition of a simplicial cone, U , and consequently A_U , must have full rank. Thus we may apply Lemma 2.8 over the entire positive orthant. We proceed to analyze the multivariate integral in Eq. (2.3.1), by first applying Lemma 2.8 and then linearizing the exponent using Lemma 2.5. Post-linearization, our approach is similar to the presentation of Boeroeczky and Henk [BJH99]. We have,

$$I \leftarrow \int_{\mathbb{R}_+^k} e^{-x^T A_U^{-1} x} dx \leq \int_{\mathbb{R}_+^k} e^{-\|x\|_1^2 / \ell^2} dx \quad (\text{by Lemma 2.8}) = \ell^k \int_{\mathbb{R}_+^k} e^{-\|y\|_1^2} dy \quad (\text{Subst. } y \leftarrow x/\ell)$$

$$\begin{aligned}
&= \frac{\ell^k}{\sqrt{\pi}} \int_{\mathbb{R}_+^k} \int_{-\infty}^{\infty} e^{-t^2 + 2it \sum_{i \in [k]} y_i} dt dy \quad (\text{by Lemma 2.5}) = \frac{\ell^k}{\sqrt{\pi}} \int_{-\infty}^{\infty} e^{-t^2} \prod_{i \in [k]} \left(\int_0^{\infty} e^{2ity_i} dy_i \right) dt \\
&= \frac{\ell^k}{\sqrt{\pi}} \int_{-\infty}^{\infty} e^{-t^2} \left(\int_0^{\infty} e^{2its} ds \right)^k dt && \text{Let } g(t) = e^{-t^2} \left(\int_0^{\infty} e^{2its} ds \right)^k \\
& && |g(a+ib)| \leq e^{-a^2+b^2} \left(\int_0^{\infty} e^{-2bs} ds \right)^k \\
& && \Rightarrow \lim_{a \rightarrow \infty} g(a+ib) \rightarrow 0, \forall b > 0 \\
&= \frac{\ell^k}{\sqrt{\pi}} \int_{-\infty}^{\infty} e^{-a^2+b^2-2abi} \left(\int_0^{\infty} e^{-2bs+2asi} ds \right)^k da && (\text{for } b > 0) \text{ by Lemma 2.7} \\
&= \frac{e^{k/2} \ell^k}{\sqrt{\pi}(2k)^{k/2}} \int_{-\infty}^{\infty} e^{-a^2} \left(2be^{-ia/b} \int_0^{\infty} e^{-2bs+2asi} ds \right)^k da && \text{Fixing } b = \sqrt{k/2} \\
&= \frac{e^{k/2} \ell^k}{\sqrt{\pi}(2k)^{k/2}} \left| \int_{-\infty}^{\infty} e^{-a^2} \left(2be^{-ia/b} \int_0^{\infty} e^{-2bs+2asi} ds \right)^k da \right| && \text{Since expr. is real and +ve} \\
&\leq \frac{e^{k/2} \ell^k}{\sqrt{\pi}(2k)^{k/2}} \int_{-\infty}^{\infty} e^{-a^2} \left(2b \int_0^{\infty} e^{-2bs} ds \right)^k da && \text{By Observation 2.6} \\
&= \frac{e^{k/2} \ell^k}{\sqrt{\pi}(2k)^{k/2}} \int_{-\infty}^{\infty} e^{-a^2} da = \frac{e^{k/2} \ell^k}{(2k)^{k/2}}
\end{aligned}$$

Lastly, the claim follows by substituting the above in Eq. (2.3.1). ■

2.4. Analysis of Hyperplane Rounding given Strong Colorability

In this section we analyze the performance of random hyperplane rounding on k -uniform hypergraphs that are $(k + \ell)$ -strongly colorable.

Theorem 2.10. *Consider any $(k + \ell)$ -strongly colorable k -uniform hypergraph $H = (V, E)$. The expected fraction of monochromatic edges obtained by performing random hyperplane rounding on the solution of Relaxation 2.3, is $O\left(\ell^{k-1/2} \left(\frac{e}{2\pi}\right)^{k/2} \frac{1}{k^{(k-1)/2}}\right)$.*

Proof. Let U be any $k \times k$ matrix whose columns are unit vectors $u_1, \dots, u_k \in \mathbb{R}^k$ that satisfy the edge constraints in Relaxation 2.3. Recall from Section 2.2, that to bound the probability of a monochromatic edge we need only bound the expression in Eq. (2.4) for U of the above form. By Relaxation 2.3, the gram matrix $A_U = U^T U$, is exactly, $A_U = (1 + \alpha)\mathbf{I} - \alpha \mathbf{1}\mathbf{1}^T$ where $\alpha = \frac{1}{k+\ell-1}$. By matrix determinant lemma (determinant formula for rank one updates), we know

$$\det(A_U) = (1 + \alpha)^k \left(1 - \frac{k\alpha}{1 + \alpha}\right) \geq \left(\frac{\ell}{k + \ell}\right) = \Omega\left(\frac{\ell}{k}\right)$$

Further, Relaxation 2.3 implies the length of $\sum_i u_i$, is at most ℓ . The claim then follows by combining Eq. (2.4) with Theorem 2.9. ■

Note. Being that any edge in the solution to the strong colorability relaxation corresponds to a symmetric cone, Theorem 2.10 is directly implied by prior work on the volume of symmetric spherical simplices. It is in the next section, where the true power of Theorem 2.9 is realized.

Remark. As can be seen from the asymptotic volume formula of symmetric spherical simplices, $\sqrt{\pi k/(2e)}$ is a sharp threshold for ℓ , i.e. when $\ell > (1 + o(1))\sqrt{\pi k/(2e)}$, hyperplane rounding does worse than the naive random algorithm, and when $\ell < (1 - o(1))\sqrt{\pi k/(2e)}$, hyperplane rounding beats the naive random algorithm.

2.5. Analysis of Hyperplane Rounding given Rainbow Colorability

In this section we analyze the performance of random hyperplane rounding on k -uniform hypergraphs that are $(k - \ell)$ -rainbow colorable.

Let U be the $k \times k$ matrix whose columns are unit vectors $u_1, \dots, u_k \in \mathbb{R}^k$ satisfying the edge constraints in Relaxation 2.2. We need to bound the expression in Eq. (2.4) for U of the above form. While we'd like to proceed just as in Section 2.4, we are limited by the possibility of U being singular or the parallelotope determined by U having arbitrarily low volume (as u_1 can be chosen arbitrarily close to the span of u_2, \dots, u_k while still satisfying $\|\sum_i u_i\|_2 \leq \ell$).

While U can be bad with respect to our properties of interest, we will show that some subset of the vectors in U are reasonably well behaved with respect to para-volume and sum-norm.

2.5.1. Finding a Well Behaved Subset

We'd like to find a subset of U with high para-volume, or equivalently, a principal sub-matrix of A_U with reasonably large determinant. To this end, we express the gram matrix $A_U = U^T U$ as the sum of a symmetric skeleton matrix B_U and a residue matrix E_U . Formally, $E_U = A_U - B_U$ and $B_U = (1 + \beta)I - \beta \mathbf{1}\mathbf{1}^T$ where $\beta = \frac{1}{k-\ell-1}$. We have (assuming $\ell = o(k)$), $\text{sum}(A_U) \leq \ell^2$ and $\text{sum}(B_U) = k - k(k-1)\beta = \frac{-\ell}{1-o(1)}$. Let $s \leftarrow \text{sum}(E_U) \leq \ell^2 - \text{sum}(B_U) = \ell^2 + \frac{\ell}{1-o(1)}$.

We further observe that E_U is symmetric, with all diagonal entries zero. Also since u_1, \dots, u_k satisfy Relaxation 2.2, all entries of E_U are non-negative.

By an averaging argument, at most ck^δ columns of E_U have column sums greater than $s/(ck^\delta)$ for some parameters δ, c to be determined later. Let $S \subseteq [k]$ be the set of indices of the columns having the lowest $k - ck^\delta$ column sums. Let $\tilde{k} \leftarrow |S| = k - ck^\delta$, and let A_S, B_S, E_S be the corresponding matrices restricted to S (in both columns and rows).

Spectrum of B_S and E_S .

Observation 2.11. For a square matrix X , let $\lambda_{\min}(X)$ denote its minimum eigenvalue. The eigenvalues of B_S are exactly $(1 + \beta)$ with multiplicity $(\tilde{k} - 1)$, and $(1 + \beta - \tilde{k}\beta)$ with multiplicity 1. Thus $\lambda_{\min}(B_S) = 1 + \beta - \tilde{k}\beta$. This is true since B_S merely shifts all eigenvalues of $-\beta \mathbf{1}\mathbf{1}^T$ by $1 + \beta$.

While we don't know as much about the spectrum of E_S , we can still say some useful things.

Observation 2.12. Since E_S is non-negative, by Perron-Frobenius theorem, its spectral radius is equal to its max column sum, which is at most $s/(ck^\delta)$. Thus $\lambda_{\min}(E_S) \geq -s/(ck^\delta)$.

Now that we know some information about the spectra of B_S and E_S , the next natural step is to consider the behaviour of spectra under matrix sums.

Spectral properties of Matrix sums.

The following identity is well known.

Observation 2.13. *If X and Y are symmetric matrices with eigenvalues $x_1 > x_2 > \dots > x_m$ and $y_1 > y_2 > \dots > y_m$ and the eigenvalues of $A + B$ are $z_1 > z_2 > \dots > z_m$, then*

$$\forall 0 \leq i + j \leq m, \quad z_{m-i-j} \geq x_{m-i} + y_{m-j}.$$

In particular, this implies $\lambda_{\min}(X + Y) \geq \lambda_{\min}(X) + \lambda_{\min}(Y)$.

We may finally analyze the spectrum of A_S .

Properties of A_S .

Observation 2.14 (Para-Volume). *Let the eigenvalues of A_S be $a_1 > a_2 > \dots > a_{\tilde{k}}$. By Observation 2.11, Observation 2.12, and Observation 2.13 we have (Assuming $\ell < ck^{\delta/2}$),*

$$\begin{aligned} \lambda_{\min}(A_S) = a_{\tilde{k}} &\geq 1 + \beta - \tilde{k}\beta - \frac{s}{ck^{\delta}} = \frac{c}{k^{1-\delta}} - \frac{\ell^2}{ck^{\delta}} - o(1) \\ a_2, a_3, \dots, a_{\tilde{k}-1} &\geq 1 + \beta - \frac{s}{ck^{\delta}} = 1 - \frac{\ell^2}{ck^{\delta}} - o(1) \end{aligned}$$

Consequently,

$$\det(A_S) \geq \left(\frac{c}{k^{1-\delta}} - \frac{\ell^2}{ck^{\delta}} - o(1) \right) \left(1 - \frac{\ell^2}{ck^{\delta}} - o(1) \right)^{\tilde{k}} \geq \left(\frac{c}{k^{1-\delta}} - \frac{\ell^2}{ck^{\delta}} - o(1) \right) e^{-k}$$

In particular, note that A_S is non-singular and has non-negligible para-volume when

$$\frac{\ell^2}{ck^{\delta}} = \frac{c}{2k^{1-\delta}}, \quad \text{i.e. } \ell \approx ck^{\delta-1/2} \quad \text{or, } \delta \approx \frac{1 \log(\ell/c)}{2 \log k}$$

Observation 2.15 (Sum-Norm). *Since E_U is non-negative, $\text{sum}(E_S) \leq \text{sum}(E_U) = s$. Also we know that the sum of entries of A_S is*

$$\text{sum}(B_S) + \text{sum}(E_S) = \tilde{k}(1 + \beta) - \tilde{k}(\tilde{k} - 1)\beta + s \leq ck^{\delta} + s \quad (2.7)$$

2.5.2. The Result.

We are now equipped to prove our result.

Theorem 2.16. *For $\ell < \sqrt{k}/100$, consider any $(k - \ell)$ -rainbow colorable k -uniform hypergraph $H = (V, E)$. Let $\theta = 1/2 + \log(\ell)/\log(k)$ and $\eta = 19(1 - \theta)/40$. The expected fraction of monochromatic edges obtained by performing random hyperplane rounding on the solution of Relaxation 2.2, is at most*

$$\frac{1}{2.1^k k^{\eta k}}$$

Proof. Let U be any $k \times k$ matrix whose columns are unit vectors $u_1, \dots, u_k \in \mathbb{R}^k$ that satisfy the edge constraints in Relaxation 2.3. Recall from Section 2.2, that to bound the probability of a monochromatic edge we need only bound the expression in Eq. (2.4) for U of the above form.

By Section 2.5.1, we can always choose a matrix U_S whose columns $\tilde{u}_1, \dots, \tilde{u}_{\tilde{k}}$ are from the set $\{u_1, \dots, u_k\}$, such that the gram matrix $A_S = U_S^T U_S$ satisfies Eq. (2.7) and Observation 2.14. Clearly the probability of all vectors in U being monochromatic is at most the probability of all vectors in U_S being monochromatic.

Thus just as in Section 2.2, to find the probability of U_S being monochromatic, we may assume without loss of generality that we are performing random hyperplane rounding in $\mathbb{R}^{\tilde{k}}$ on any \tilde{k} -dimensional vectors $\tilde{u}_1, \dots, \tilde{u}_{\tilde{k}}$ whose gram (pairwise inner-product) matrix is the aforementioned A_S .

Specifically, by combining Eq. (2.7) and Observation 2.14 with Theorem 2.9, our expression is at most:

$$\left(\frac{e}{2\pi}\right)^{\tilde{k}/2} \left(\frac{ck^\delta + s}{k}\right)^{\tilde{k}/2} \frac{1}{\sqrt{\det(A_U)}} \leq 3 \cdot 2^{\tilde{k}/2} \left(\frac{(1-o(1))c}{k^{1-\delta}}\right)^{\tilde{k}/2} \leq \frac{1}{2 \cdot 1^k k^{(1-c)(1-\delta)k}}$$

assuming $c = 1/20$, $\delta \geq 1/2$ and $\ell < \sqrt{k}/100$ (constraint on ℓ ensures that non-singularity conditions of Observation 2.14 are satisfied). The claim follows. \blacksquare

Remark. Yet again we see a threshold for ℓ , namely, when $\ell < \sqrt{k}/100$, hyperplane rounding beats the naive random algorithm, and for $\ell = \Omega(\sqrt{k})$, it fails to do better. In fact, as we'll see in the next section, assuming the UGC, we show a hardness result when $\ell = \Omega(\sqrt{k})$.

3. Hardness of Max-2-Coloring under Low Discrepancy

In this section we consider the hardness of Max-2-Coloring when promised discrepancy as low as one. As noted in Section 2.5, our analysis requires the configuration of vectors in an edge to be well behaved with respect to sum-norm and para-volume. While in the discrepancy case, we can ensure good sum-norm, the vectors in an edge can have arbitrarily low para-volume. While in the rainbow case we can remedy this by finding a reasonably large well behaved subset of vectors, this is not possible in the case of discrepancy.

Indeed, consider the following counterexample: Start in 2 dimensions with $k/3$ copies each of any u_1, u_2, u_3 such that $u_1 + u_2 + u_3 = 0$. Lift all vectors to 3-dimensions by assigning every vector a third coordinate of value exactly $1/k$. This satisfies Relaxation 2.1, yet every superconstant sized subset has para-volume zero.

Confirming that this is not an artifact of our techniques and the problem is in fact hard, we show in this section via a reduction from Max-Cut, that assuming the Unique Games conjecture, it is NP-Hard to Max-2-Color much better than the naive random algorithm that miscolors 2^{-k+1} fraction of edges, even in the case of discrepancy-1 hypergraphs.

3.1. Reduction from Max-Cut

Let $k = 2t + 1$. Let $G = (V, E)$ be an instance of Max-Cut, where each edge has weight 1. Let $n = |V|$ and $m = |E|$. We produce a hypergraph $H = (V', E')$ where $V' = V \times [k]$. For each $u \in V$, let $\text{cloud}(u) := \{u\} \times [k]$. For each edge $(u, v) \in E$, we add $N := 2 \binom{k}{t} \binom{k}{t+1}$ hyperedges

$$\{U \cup V : U \subseteq \text{cloud}(u), V \subseteq \text{cloud}(v), |U| + |V| = k, ||U| - |V|| = 1\},$$

each with weight $\frac{1}{N}$. Call these hyperedges *created by* (u, v) . The sum of weights is m for both G and H .

3.1.1. Completeness

Given a coloring $C : V \mapsto \{B, W\}$ that cuts at least $(1 - \alpha)m$ edges of G , we color H so that for every $v \in V$, each vertex in $\text{cloud}(v)$ is given the same color as v . If $(u, v) \in E$ is cut, all hyperedges created by (u, v) will have discrepancy 1. Therefore, the total weight of hyperedges with discrepancy 1 is at least $(1 - \alpha)m$.

3.1.2. Soundness

Given a coloring $C' : V' \mapsto \{B, W\}$ such that the total weight of non-monochromatic hyperedges is $(1 - \beta)m$, $v \in V$ is given the color that appears the most in its cloud (k is odd, so it is well-defined). Consider $(u, v) \in E$. If no hyperedge created by (u, v) is monochromatic, it means that u and v should be given different colors by the above majority algorithm (if they are given the same color, say white, then there are at least $t + 1$ white vertices in both clouds, so we have at least one monochromatic hyperedge).

This means that for each $(u, v) \in E$ that is uncut by the above algorithm (lost weight 1 for Max-Cut objective), at least one hyperedge created by (u, v) is monochromatic, and we lost weight at least $\frac{1}{N}$ there for our problem. This means that the total weight of cut edges for Max-Cut is at least $(1 - \beta N)m$.

3.1.3. The Result

Theorem 3.1 ([KKMO07]). *Let $G = (V, E)$ be a graph with $m = |E|$. For sufficiently small $\varepsilon > 0$, it is UG-hard to distinguish the following cases.*

- *There is a 2-coloring that cuts at least $(1 - \varepsilon)|E|$ edges.*
- *Every 2-coloring cuts at most $(1 - (2/\pi)\sqrt{\varepsilon})|E|$ edges.*

Our reduction shows that

Theorem 3.2. *Given a hypergraph $H = (V, E)$, it is UG-hard to distinguish the following cases.*

- *There is a 2-coloring where at least $(1 - \varepsilon)$ fraction of hyperedges have discrepancy 1.*
- *Every 2-coloring cuts (in a standard sense) at most $(1 - (2/\pi)\frac{\sqrt{\varepsilon}}{N})$ fraction of hyperedges.*

$N = 2 \binom{k}{t} \binom{k}{t+1} \leq (2/\pi)2^k \cdot 2^k \leq (2/\pi)2^{2k}$. If we take $\varepsilon = 2^{-6k}$ for large enough k , we cannot distinguish

- *There is a 2-coloring where at least $(1 - 2^{-6k})$ fraction of hyperedges have discrepancy 1.*
- *Every 2-coloring cuts (in a standard sense) at most $(1 - 2^{-5k})$ fraction of hyperedges.*

This proves Theorem 1.2.

3.2. NP-Hardness

In this subsection, we show that given a hypergraph which admits a 2-coloring with discrepancy at most 2, it is NP-hard to find a 2-coloring that has less than $k^{-O(k)}$ fraction of monochromatic hyperedges. Note that while the inapproximability factor is worse than the previous subsection, we get NP-hardness and it holds when the input hypergraph is promised to have *all* hyperedges have discrepancy at most 2. The reduction and the analysis closely follow from the more general framework of Guruswami and Lee [GL15a] except that we prove a better reverse hypercontractivity bound for our case.

3.2.1. Q -Hypergraph Label Cover

An instance of Q -Hypergraph Label Cover is based on a Q -uniform hypergraph $H = (V, E)$. Each hyperedge-vertex pair (e, v) such that $v \in e$ is associated with a projection $\pi_{e,v} : [R] \rightarrow [L]$ for some positive integers R and L . A labeling $l : V \rightarrow [R]$ *strongly satisfies* $e = \{v_1, \dots, v_Q\}$ when $\pi_{e,v_1}(l(v_1)) = \dots = \pi_{e,v_Q}(l(v_Q))$. It *weakly satisfies* e when $\pi_{e,v_i}(l(v_i)) = \pi_{e,v_j}(l(v_j))$ for some $i \neq j$. The following are two desired properties of instances of Q -Hypergraph Label Cover.

- Regular: every projection is d -to-1 for $d = R/L$.
- Weakly dense: any subset of V of measure at least ε vertices induces at least $\frac{\varepsilon Q}{2}$ fraction of hyperedges.
- T -smooth: for all $v \in V$ and $i \neq j \in [R]$, $\Pr_{e \in E: e \ni v}[\pi_{e,v}(i) = \pi_{e,v}(j)] \leq \frac{1}{T}$.

The following theorem asserts that it is NP-hard to find a good labeling in such instances.

Theorem 3.3 ([GL15a]). *For all integers $T, Q \geq 2$ and $\eta > 0$, the following is true. Given an instance of Q -Hypergraph Label Cover that is regular, weakly-dense and T -smooth, it is NP-hard to distinguish between the following cases.*

- *Completeness: There exists a labeling l that strongly satisfies every hyperedge.*
- *Soundness: No labeling l can weakly satisfy η fraction of hyperedges.*

3.2.2. Distributions

We first define the distribution $\bar{\mu}'$ for each *block*. $2Q$ points $x_{q,i} \in \{1, 2\}^d$ for $1 \leq q \leq Q$ and $1 \leq i \leq 2$ are sampled by the following procedure.

- Sample $q' \in [Q]$ uniformly at random.
- Sample $x_{q',1}, x_{q',2} \in \{1, 2\}^d$ i.i.d.
- For $q \neq q', 1 \leq j \leq d$, sample a permutation $((x_{q,1})_j, (x_{q,2})_j) \in \{(1, 2), (2, 1)\}$ uniformly at random.

3.2.3. Reduction and Completeness

We now describe the reduction from Q -Hypergraph Label Cover. Given a Q -uniform hypergraph $H = (V, E)$ with Q projections from $[R]$ to $[L]$ for each hyperedge (let $d = R/L$), the resulting instance of $2Q$ -Hypergraph Coloring is $H' = (V', E')$ where $V' = V \times \{1, 2\}^R$. Let $\text{cloud}(v) := \{v\} \times \{1, 2\}^R$. The set E' consists of hyperedges generated by the following procedure.

- Sample a random hyperedge $e = (v_1, \dots, v_Q) \in E$ with associated projections $\pi_{e,v_1}, \dots, \pi_{e,v_Q}$ from E .
- Sample $(x_{q,i})_{1 \leq q \leq Q, 1 \leq i \leq 2} \in \{1, 2\}^R$ in the following way. For each $1 \leq j \leq L$, independently sample $((x_{q,i})_{\pi_{e,v_q}^{-1}(j)})_{q,i}$ from $(\{1, 2\}^d)_{2Q, \bar{\mu}'}$.
- Add a hyperedge between $2Q$ vertices $\{(v_q, x_{q,i})\}_{q,i}$ to E' . We say this hyperedge is *formed from* $e \in E$.

Given the reduction, completeness is easy to show.

Lemma 3.4. *If an instance of Q -Hypergraph Label Cover admits a labeling that strongly satisfies every hyperedge $e \in E$, there is a coloring $c : V' \rightarrow \{1, 2\}$ of the vertices of H' such that every hyperedge $e' \in E'$ has at least $(Q - 1)$ vertices of each color.*

Proof. Let $l : V \rightarrow [R]$ be a labeling that strongly satisfies every hyperedge $e \in E$. For any $v \in V, x \in \{1, 2\}^R$, let $c(v, x) = x_{l(v)}$. For any hyperedge $e' = \{(v_q, x_{q,i})\}_{q,i} \in E'$, $c(v_q, x_{q,i}) = (x_{q,i})_{l(v_q)}$, and all but one q satisfies $\{(x_{q,1})_{l(v_q)}, (x_{q,2})_{l(v_q)}\} = \{1, 2\}$. Therefore, the above strategy ensures that every hyperedge of E' contains at least $(Q - 1)$ vertices of each color. ■

3.2.4. Soundness

Lemma 3.5. *There exists $\eta := \eta(Q)$ such that if $I \subseteq V'$ of measure $\frac{1}{2}$ induces less than $Q^{-O(Q)}$ fraction of hyperedges in H' , the corresponding instance of Q -Hypergraph Label Cover admits a labeling that weakly satisfies a fraction η of hyperedges.*

Proof. Consider a vertex v and hyperedge $e \in E$ that contains v with a permutation $\pi = \pi_{e,v}$. Let $f : \{1, 2\}^R \mapsto [0, 1]$ be a *noised* indicator function of $I \cap \text{cloud}(v)$ with $\mathbb{E}_{x \in \{1, 2\}^R} [f(x)] \geq \frac{1}{2} - \varepsilon$ for small $\varepsilon > 0$ that will be determined later. We define the inner product

$$\langle f, g \rangle = \mathbb{E}_{x \in \{1, 2\}^R} [f(x)g(x)].$$

f admits the Fourier expansion

$$\sum_{S \subseteq [R]} \hat{f}(S) \chi_S$$

where

$$\chi_S(x_1, \dots, x_k) = \prod_{i \in S} (-1)^{x_i}, \quad \hat{f}(S) = \langle f, \chi_S \rangle.$$

In particular, $\hat{f}(\emptyset) = \mathbb{E}[f(x)]$, and

$$\sum_S \hat{f}(S)^2 = \mathbb{E}[f(x)^2] \leq \mathbb{E}[f(x)] \quad (3.1)$$

A subset $S \subseteq [R]$ is said to be *shattered* by π if $|S| = |\pi(S)|$. For a positive integer J , we decompose f as the following:

$$\begin{aligned} f^{\text{good}} &= \sum_{s: \text{shattered}} \hat{f}(S) \chi_S \\ f^{\text{bad}} &= f - f^{\text{good}}. \end{aligned}$$

By adding a suitable noise and using smoothness of Label Cover, for any $\delta > 0$, we can assume that $\|f^{\text{bad}}\|^2 \leq \delta$. See [GL15a] for the details.

Each time a $2Q$ -hyperedge is sampled is formed from e , two points are sampled from each cloud. Let x, y be the points in $\text{cloud}(v)$. Recall that they are sampled such that for each $1 \leq j \leq L$,

- With probability $\frac{1}{Q}$, for each $i \in \pi^{-1}(j)$, x_i and y_i are independently sampled from $\{1, 2\}$.
- With probability $\frac{Q-1}{Q}$, for each $i \in \pi^{-1}(j)$, (x_i, y_i) are sampled from $\{(1, 2), (2, 1)\}$.

We can deduce the following simple properties.

1. $\mathbb{E}_{x, y} [\chi_{\{i\}}(x) \chi_{\{i\}}(y)] = -\frac{Q-1}{Q}$. Let $\rho := -\frac{Q-1}{Q}$.
2. $\mathbb{E}_{x, y} [\chi_{\{i\}}(x) \chi_{\{j\}}(y)] = 0$ if $i \neq j$.
3. $\mathbb{E}_{x, y} [\chi_S(x) \chi_T(y)] = 0$ unless $\pi(S) = \pi(T) = \pi(S \cap T)$.

We are interested in lower bounding

$$\mathbb{E}_{x, y} [f(x)f(y)] \geq \mathbb{E}[f^{\text{good}}(x)f^{\text{good}}(y)] - 3\|f^{\text{bad}}(x)\|^2\|f\|^2 \geq \mathbb{E}[f^{\text{good}}(x)f^{\text{good}}(y)] - 3\delta.$$

By the property 3.,

$$\mathbb{E}[f^{\text{good}}(x)f^{\text{good}}(y)] = \sum_{s: \text{shattered}} \hat{f}(S)^2 \rho^{|S|}$$

$$\begin{aligned}
&= \mathbb{E}[f]^2 + \sum_{s: \text{shattered}} \hat{f}(s)^2 \rho^{|s|} \\
&\geq \mathbb{E}[f]^2 + \rho \left(\sum_{|s|>1} \hat{f}(s)^2 \right) \quad \text{since } \rho \text{ is negative} \\
&\geq \mathbb{E}[f]^2 + \rho (\mathbb{E}[f] - \mathbb{E}[f]^2) \quad \text{by (3.1)} \\
&\geq \mathbb{E}[f]^2 (1 + \rho) - \varepsilon \quad \text{since } \mathbb{E}[f] \geq \frac{1}{2} - \varepsilon \Rightarrow \mathbb{E}[f] - \mathbb{E}[f]^2 \leq \mathbb{E}[f]^2 + \varepsilon \\
&\geq \frac{\mathbb{E}[f]^2}{Q} - \varepsilon.
\end{aligned}$$

By taking ε and δ small enough, we can ensure that

$$\mathbb{E}[f(x)f(y)] \geq \zeta := \frac{1}{5Q}. \quad (3.2)$$

The soundness analysis of Guruswami and Lee [GL15a] ensures ((3.2) replaces their Step 2) that there exists $\eta := \eta(Q)$ such that if the fraction of hyperedges induced by I is less than $Q^{-O(Q)}$, the Hypergraph Label Cover instance admits a solution that satisfies η fraction of constraints. We omit the details. \blacksquare

3.2.5. Corollary to Max-2-Coloring under discrepancy $O(\log k)$

The above NP-hardness, combined with the reduction technique from Max-Cut in Section 3.1, shows that given a k -uniform hypergraph, it is NP-hard to distinguish whether it has discrepancy at most $O(\log k)$ or any 2-coloring leaves at least $2^{-O(k)}$ fraction of hyperedges monochromatic. Even though the direction reduction from Max-Cut results in a similar inapproximability factor with discrepancy even 1, this result does not rely on the UGC and hold even *all* edges (compared to *almost* in Section 3.1) have discrepancy $O(\log k)$.

Let $r = \Theta(\frac{k}{\log k})$ so that $s = \frac{k}{r} = \Theta(\log k)$ is an integer. Given a r -uniform hypergraph, it is NP-hard to distinguish whether it has discrepancy at most 2 or any 2-coloring leaves at least $r^{-O(r)}$ fraction of hyperedges monochromatic. Given a r -uniform hypergraph, the reduction replaces each vertex v with $\text{cloud}(v)$ that contains $(2s - 1)$ new vertices. Each hyperedge (v_1, \dots, v_r) is replaced by $d := \binom{2s-1}{s}^r \leq (2^s)^r = 2^k$ hyperedges

$$\{\cup_{i=1}^r V_i : V_i \subset \text{cloud}(v_i), |V_i| = s\}.$$

If the given r -uniform hypergraph has discrepancy at most 2, the resulting k -uniform hypergraph has discrepancy at most $2s = O(\log k)$.

If the resulting k -uniform hypergraph admits a coloring that leaves α fraction of hyperedges monochromatic, giving v the color that appears more in $\text{cloud}(v)$ is guaranteed to leaves at most $d\alpha$ fraction of hyperedges monochromatic. Therefore, if any 2-coloring of the input r -uniform hypergraph leaves at least $r^{-O(r)}$ fraction of hyperedges monochromatic, any 2-coloring of the resulting k -uniform hypergraph leaves at least $\frac{r^{-O(r)}}{d} = 2^{-O(k)}$ fraction of hyperedges.

3.3. Hardness of Max-2-Coloring under almost $(k - \sqrt{k})$ -colorability

Let k be such that $\ell := \sqrt{k}$ be an integer and let $\chi := k - \ell$. We prove the following hardness result for any $\varepsilon > 0$ assuming the Unique Games Conjecture: given a k -uniform hypergraph such that there is a χ -coloring that have at least $(1 - \varepsilon)$ fraction of hyperedges rainbow, it is NP-hard to find a 2-coloring that leaves at most $(\frac{1}{2})^{k-1}$ fraction of hyperedges monochromatic.

The main technique for this result is to show the existence of a *balanced pairwise independence distribution* with the desired support. Let μ be a distribution on $[\chi]^k$. μ is called balanced pairwise independent

if for any $i \neq j \in [k]$ and $a, b \in [\chi]$,

$$\Pr_{(x_1, \dots, x_k) \sim \mu} [x_i = a, x_j = b] = \frac{1}{\chi^2}.$$

For example, the uniform distribution on $[\chi]^k$ is a balanced pairwise distribution. We now consider the following distribution μ to sample $(x_1, \dots, x_k) \in [\chi]^k$.

- Sample $S \subseteq [k]$ with $|S| = \chi$ uniformly at random. Let $S = \{s_1 < \dots < s_\chi\}$.
- Sample a permutation $\pi : [\chi] \mapsto [\chi]$.
- Sample $y \in [\chi]$.
- For each $i \in [k]$, if $i = s_j$ for some $j \in [\chi]$, output $x_i = \pi(\chi)$. Otherwise, output $x_i = y$.

Note that for any supported by (x_1, \dots, x_k) , we have $\{x_1, \dots, x_k\} = [\chi]$. Therefore, μ is supported on *rainbow strings*. We now verify pairwise independence. Fix $i \neq j \in [k]$ and $a, b \in [\chi]$.

- If $a = b$, by conditioning on whether i, j are in S or not,

$$\begin{aligned} \Pr_{\mu} [x_i = a, x_j = b] &= \Pr[x_i = a, x_j = b | i, j \in S] \Pr[i, j \in S] + \\ &\quad \Pr[x_i = a, x_j = b | i \in S, j \notin S] \Pr[i \in S, j \notin S] + \\ &\quad \Pr[x_i = a, x_j = b | i \notin S, j \in S] \Pr[i \notin S, j \in S] + \\ &\quad \Pr[x_i = a, x_j = b | i, j \notin S] \Pr[i, j \notin S] \\ &= 0 \cdot \left(\frac{\chi(\chi-1)}{k(k-1)}\right) + 2 \cdot \left(\frac{1}{\chi^2}\right) \cdot \left(\frac{\ell\chi}{k(k-1)}\right) + \left(\frac{1}{\chi}\right) \cdot \left(\frac{\ell(\ell-1)}{k(k-1)}\right) \\ &= \frac{2\ell\chi + \chi(\ell^2 - \ell)}{\chi^2 k(k-1)} = \frac{\chi k + \chi\sqrt{k}}{\chi^2 k(k-1)} = \frac{\sqrt{k}(\sqrt{k}+1)}{\chi k(\sqrt{k}+1)(\sqrt{k}-1)} \\ &= \frac{1}{\chi(k-\sqrt{k})} = \frac{1}{\chi^2}. \end{aligned}$$

- If $a \neq b$, by the same conditioning,

$$\begin{aligned} \Pr_{\mu} [x_i = a, x_j = b] &= \left(\frac{1}{\chi(\chi-1)}\right) \cdot \left(\frac{\chi(\chi-1)}{k(k-1)}\right) + 2 \cdot \left(\frac{1}{\chi^2}\right) \cdot \left(\frac{\ell\chi}{k(k-1)}\right) + 0 \cdot \left(\frac{\ell(\ell-1)}{k(k-1)}\right) \\ &= \frac{\chi^2 + 2\ell\chi}{\chi^2 k(k-1)} = \frac{\chi + 2\ell}{\chi k(k-1)} = \frac{k + \sqrt{k}}{\chi k(k-1)} = \frac{1}{\chi^2}. \end{aligned}$$

Given such a balanced pairwise independent distribution supported on rainbow strings, a standard procedure following the work of Austrin and Mossel [AM09] shows that it is UG-hard to outperform the random 2-coloring. We omit the details.

4. Approximate Min-Coloring

In this section, we provide approximation algorithms for the Min-Coloring problem under strong colorability, rainbow colorability, and low discrepancy assumptions. Our approach is standard, namely, we first

apply degree reduction algorithms followed by the usual paradigm pioneered by Karger, Motwani and Sudan [KMS98], for coloring bounded degree (hyper)graphs. Consequently, our exposition will be brief and non-linear.

In the interest of clarity, all results henceforth assume the special cases of Discrepancy 1, or $(k-1)$ -rainbow colorability, or $(k+1)$ -strong colorability. All arguments generalize easily to the cases parameterized by l .

4.1. Approximate Min-Coloring in Bounded Degree Hypergraphs

4.1.1. The Algorithm

INPUT: k -uniform hypergraph $H = ([n], E)$ with max-degree t and m edges, having Discrepancy 1, or being $(k-1)$ -rainbow colorable, or being $(k+1)$ -strong colorable.

1. Let u_1, \dots, u_n be a solution to the SDP relaxation from Section 2.1 corresponding to the assumption on the hypergraph.
2. Let H_1 be a copy of H , and let γ, τ be parameters to be determined shortly.
3. Until no vertex remains in the hypergraph, Repeat:
 - Find an independent set \mathcal{S} in the residual hypergraph, of size at least γn by repeating the below process until $|\mathcal{S}| \geq \gamma n$:
 - (A) Pick a random vector r from the standard multivariate normal distribution.
 - (B) For all i , if $\langle u_i, r \rangle \geq \tau$, add vertex i to \mathcal{S} .
 - (C) For every edge e completely contained in \mathcal{S} , delete any single vertex in e , from \mathcal{S} .

Color \mathcal{S} with a new color and remove \mathcal{S} and all edges involving vertices in \mathcal{S} , from H_1 .

4.1.2. Analysis

First note that by Lemma 2.2, for any fixed vector a , $\langle a, r \rangle$ has the distribution $\mathcal{N}(0, 1)$. Note that all SDP formulations in Section 2.1 satisfy,

$$\left\| \sum_{j \in [k]} u_{i_j} \right\|_2 \leq 1 \quad (4.1)$$

Now consider any edge $e = (i_1, \dots, i_k)$. In any fixed iteration of the inner loop, the probability of e being contained in \mathcal{S} at Step (B), is at most the probability of

$$\langle r, \sum_{j \in [k]} u_{i_j} \rangle \geq k\tau$$

However, by Lemma 2.2 and Eq. (4.1), the inner product above is dominated by the distribution $\mathcal{N}(0, 1)$. Thus in any fixed iteration of the inner loop, let H_1 have n_1 vertices and m_1 edges, we have

$$\begin{aligned} \mathbf{E}[|\mathcal{S}|] &\geq n_1 \Phi(\tau) - m_1 \Phi(k\tau) \\ &\geq n_1 e^{-\tau^2/2} - \frac{n_1 t}{k} e^{-k^2 \tau^2/2} \\ &= \Omega(\gamma n_1) \end{aligned} \quad \text{setting, } \tau^2 = \frac{2 \log t}{k^2 - 1}, \text{ and } \gamma = t^{-1/(k^2-1)}$$

Now by applying Markov's inequality to the vertices not in \mathcal{S} , we have, $\Pr[|\mathcal{S}| < \gamma n_1] \leq 1 - \Omega(\gamma)$. Thus for a fixed iteration of the outer loop, with high probability, the inner loop doesn't repeat more than $O(\log n_1 / \gamma)$ times.

Lastly, the outermost loop repeats $O(\log n/\gamma)$ times, using one color at each iteration. Thus with high probability, in polynomial time, the algorithm colors H with

$$t^{\frac{1}{k^2-1}} \log n \text{ colors.}$$

Important Note. We can be more careful in the above analysis for the rainbow and strong colorability cases. Specifically, the crux boils down to finding the gaussian measure of the cone given by $\{x \mid U^T x \geq \tau\}$ instead of zero. Indeed, on closely following the proof of Theorem 2.9 we obtain for strong and rainbow coloring respectively (assuming max-degree n^k),

$$n^{\frac{1}{k}(1-\frac{3\beta}{2})} \log n \quad \text{and} \quad n^{\frac{1}{k}(1-\frac{5\beta}{4})} \log n, \quad \text{where } \beta = \frac{\log k}{\log n}$$

While these improvements are negligible for small k , they are significant when k is reasonably large with respect to n .

4.2. Main Min-Coloring Result

Combining results from Section 4.1.2 with our degree reduction approximation schemes from the forthcoming sections, we obtain the following.

Theorem 4.1. *Consider any k -uniform hypergraph $H = (V, E)$ with n vertices. In $n^{c+O(1)}$ time, one can color H with*

$$\min \left\{ \left(\frac{n}{c \log n} \right)^\alpha, n^{\frac{1}{k}(1-\frac{3\beta}{2})}, \left(\frac{m}{n} \right)^{\frac{1}{k^2}} \right\} \log n \text{ colors,} \quad \text{if } H \text{ is } (k+1)\text{-strongly colorable.}$$

$$\min \left\{ \left(\frac{n}{c} \right)^\alpha, n^{\frac{1}{k}(1-\frac{5\beta}{4})}, \left(\frac{m}{n} \right)^{\frac{1}{k^2}} \right\} \log n \text{ colors,} \quad \text{if } H \text{ is } (k-1)\text{-rainbow colorable.}$$

$$\min \left\{ \left(\frac{n}{c} \right)^\alpha, \left(\frac{m}{n} \right)^{\frac{1}{k^2}} \right\} \log n \text{ colors,} \quad \text{if } H \text{ has discrepancy 1.}$$

$$\text{where, } \alpha = \frac{1}{k+2-o(1)}, \quad \beta = \frac{\log k}{\log n}$$

Remark. In all three promise cases the general polytime min-coloring guarantee parameterized by ℓ , is roughly $n^{\ell^2/k}$. Thus, the threshold value of ℓ , for which standard min-coloring techniques improve with k , is $o(\sqrt{k})$.

Degree Reduction Schemes under Promise. Wigderson [Wig83] and Alon et al. [AKMH96] studied degree reduction in the cases of 3-colorable graphs and 2-colorable hypergraphs, respectively. Assuming our proposed structures, we are able to combine some simple combinatorial ideas with counterparts of the observations made by Wigderson and Alon et al., to obtain degree reduction approximation schemes. Such approximation schemes are likely not possible assuming only 2-colorability.

4.3. Degree Reduction under strong colorability

Let $H = \left(V, E \subseteq \binom{V}{k} \right)$ be a k -uniform $(k+1)$ -strongly colorable hypergraph with n vertices and m edges. In this section, we give an algorithm that in $n^{c+O(1)}$ time, partially colors H with $3n(k+1) \log k / (t^{1/(k-1)} c \log n)$ colors, such that no edge in the colored subgraph is monochromatic, and furthermore, the subgraph induced by the uncolored vertices has max-degree t .

The following observations motivate the structure of our algorithm.

Observation 4.2. *For any $(k+1)$ -strong coloring $f : V \mapsto [k+1]$, of a k -uniform hypergraph H , and any subset of vertices \bar{V} satisfying, $\forall u, v \in \bar{V}, f(u) = f(v) = j$ (all of the same color), the subgraph F of H , induced by $N(\bar{V})$, is k -uniform and k -strongly colorable. This is because f is a strong coloring of F , and moreover, $\forall v \in N(\bar{V}), f(v) \neq j$, since v has a neighbor in \bar{V} with color j . Thus we can 2-color such a subgraph F in polynomial time.*

Observation 4.3. *By Observation 4.2, in order to $3(k+1)$ -color the subgraph induced by $\bar{V} \cup N(\bar{V})$ for an arbitrary subset \bar{V} of vertices, we need only search through all possible $(k+1)$ -colorings of \bar{V} , and then attempt to 2-color the neighborhood of each color class with two new colors. This process will always terminate with some proper coloring of $\bar{V} \cup N(\bar{V})$.*

We are now prepared to state the algorithm.

4.3.1. The Algorithm *SCDegreeReduce*

1. Let H_1 be a copy of H .
2. **While** H_1 contains a vertex of degree greater than t :
 - (A) Let H_2 be a copy of H_1 .
 - (B) Sequentially pick arbitrary vertices $\bar{V} = \{v_1, v_2 \dots v_s\}$ of degree at least t from H_2 , wherein we remove from H_2 the vertices $\{v_i\} \cup N(v_i)$ and all involved edges, after picking v_i and before picking v_{i+1} . We only stop when we have either picked $c \log n / \log k$ vertices, or H_2 has max-degree t .
 - (C) For every possible assignment of $k+1$ new colors $\{c_1, \dots, c_{k+1}\}$ to the vertices in \bar{V} :
 - (C1) Let $C_i = \{u \mid v \in \bar{V}, \text{color}(v) = c_i, u \in N_{H_1}(v)\}$. Then for each $i \in [k+1]$, 2-color the subgraph of H_1 induced by $N_{H_1}(C_i)$ using two new colors and the proper 2-coloring algorithm for r -uniform, r -strongly colorable graphs.
 - (C2) **If** no edge is monochromatic:
Stick with this $3(k+1)$ -coloring of $\bar{V} \cup N_{H_1}(\bar{V})$, remove $\bar{V} \cup N_{H_1}(\bar{V})$ and all edges containing any of these vertices, from H_1 , and stop iterating through assignments of \bar{V} .
 - (C3) **If** some edge is monochromatic:
Discard the coloring and continue iterating through assignments of \bar{V} .

End While

3. **Output** the partial coloring of H and the residual graph H_1 of max-degree t .

4.3.2. The Result

Theorem 4.4. Let $H = (V, E \subseteq \binom{V}{k})$ be a k -uniform $(k+1)$ -strongly colorable hypergraph with n vertices.

Algorithm 4.4.2 partially colors H in $n^{c+O(1)}$ time, with at most $\frac{3n(k+1)\log k}{t^{1/(k-1)}c\log n}$ colors, such that:

1. The subgraph of H induced by the colored vertices has no monochromatic vertices.
2. The subgraph of H induced by the uncolored vertices has maximum degree t .

Proof. Observation 4.2 combined with the fact that step (C1) uses two new colors for each C_i , establishes that step (C) of Algorithm 4.4.2 will always terminate with some proper coloring of $\bar{V} \cup N_{H_1}(\bar{V})$. Furthermore, any edge intersecting $\bar{V}_1 \cup N_{H_1}(\bar{V}_1)$ and $\bar{V}_2 \cup N_{H_1}(\bar{V}_2)$ for V_1 and V_2 taken from different iterations of Algorithm 4.4.2, cannot be monochromatic since we use new colors in each iteration. Thus the partial coloring is proper.

For the claim on number of colors, observe that a vertex of degree at least t , must have at least $(k-1)t^{1/(k-1)}$ distinct neighbors. Thus step (C) can be run at most $n/t^{1/(k-1)}$ times, using $3(k+1)$ new colors each time. Lastly for the runtime, note that for each run of step (C), there are at most $(k+1)^{c\log n/\log k} = n^{c+O(1)}$ assignments to try, and the rest of the work takes $n^{O(1)}$ time. ■

Remark. We contrast Theorem 4.4 with the results of Alon et al. [AKMH96], who give a polynomial time algorithm for degree reduction in 2-colorable k -uniform hypergraphs using $O(n/t^{1/(k-1)})$ colors. The strong coloring property, gives us additional power, namely, we obtain an approximation scheme, and furthermore, for constant c , Theorem 4.4 uses fewer colors than the result of Alon et al., by a factor of about $k \log n / \log k$. The arguments in this section and the next are readily generalizable - One can modify the degree reduction algorithm, such that the bound on colors used, would be a function of the strong colorability parameter of the hypergraph.

4.4. Degree Reduction under Low Discrepancy

For odd k , let $H = (V, E)$ be a k -uniform hypergraph with n vertices, that admits a discrepancy 1 coloring. In this section, we give an algorithm that in $n^{c+O(1)}$ time, partially colors H with $3n(k+1)/(t^{1/(k-1)}c\log n)$ colors, such that no edge in the induced colored subgraph is monochromatic, and furthermore, the subgraph induced by the the uncolored vertices has max-degree t .

First, we present a warmup algorithm that exposes the key ideas. The following observations motivate the structure of our algorithm.

Observation 4.5. For any discrepancy 1 coloring $f : V \mapsto \{-1, 1\}$, of a k -uniform hypergraph H , and any size $k-1$ subset of vertices S , we have:

- (A) If $N(S)$ is an independent set, we can properly 2-color the subgraph induced by $S \cup N(S)$.
- (B) If $N(S)$ contains an edge, then the set S has discrepancy 0 in the coloring f . This is because, an edge cannot be monochromatic in the coloring f , and by assumption, S must be have a neighbor with color -1 and a neighbor with color $+1$.

Though Observation 4.5 and Observation 4.2 are functionally similar, the two-pronged nature of Observation 4.5 almost wholly accounts for the gap in power between the respective degree reduction algorithms. Intuitively, the primary weakness comes from the fact that $N(S)$ being an independent set tells us nothing about the discrepancy of S .

Nevertheless, we may still exploit some aspects of this observation.

Observation 4.6. *Consider any discrepancy 1 coloring $f : V \mapsto \{-1, 1\}$, of a k -uniform hypergraph H , and any set of subsets S_1, \dots, S_m each of size $(k-1)$ and discrepancy 0 in the coloring f . The $(k-1)$ -uniform hypergraph F with vertex set $\bigcup_i S_i$ and edge set $\{S_1, \dots, S_m\}$, has a discrepancy 0 coloring (f). Thus we can properly 2-color F in polynomial time.*

We are now ready to state the warmup algorithm, whose correctness is evident from Observation 4.5 and Observation 4.6

4.4.1. Warmup Algorithm

1. Let H_1 be a copy of H , and set $\text{MARKED} \leftarrow \phi$
2. **While** H_1 contains a size $(k-1)$ subset S such that $N_{H_1}(S) > t$:
 - (A) **If** $N_{H_1}(S)$ contains an edge:
Delete from H_1 all edges that completely contain S . Also, add S to MARKED .
 - (B) **If** $N_{H_1}(S)$ is an independent set:
Use 2 new colors, color S one color and $N_{H_1}(S)$ the other, remove $S \cup N_{H_1}(S)$ and all edges containing any of these vertices from H_1 .

End While

3. Let F be the $(k-1)$ -uniform hypergraph whose vertex set is the union of the sets in MARKED , and whose edge set is MARKED . Using 2 new colors, properly 2-color the vertices of F using the 2-coloring algorithm for discrepancy 0 hypergraphs. Remove these vertices and all involved edges, from H_1 .
4. **Output** the partial coloring of H and the residual graph H_1 of max-degree t .

4.4.2. The Algorithm *LDDegreeReduce*

1. Let H_1 and H_2 be copies of H , $\text{MARKED} \leftarrow \phi$ and $T \leftarrow \phi$.
2. **While** H_2 contains a size $(k-1)$ subset S of vertices, such that $|N_{H_2}(S)| > t$:
 - (A) Delete $N_{H_2}(S)$ and all edges involving these vertices, from H_2 .
 - (B) **If** $N_{H_1}(S)$ contains an edge:
Delete from H_1 all edges that completely contain S . Also, add S to MARKED .
 - (C) **If** $N_{H_1}(S)$ is an independent set:
Add S to T .
 - (D) For every size c subset $\bar{V} = \{S'_1, \dots, S'_c\}$ of T :

Fix two new colors c_1, c_2 .

For every possible assignment of c_1, c_2 to \bar{V} , such that each S'_i has discrepancy 2, (We define $bias(S'_i) = c_1$ (resp. c_2) for coloring bias towards c_1 (resp. c_2)):

(D1) For $i = 1, 2$, let $C_i = \{u \mid S' \in \bar{V}, bias(S') = c_i, u \in N_{H_1}(S')\}$. Then color $N_{H_1}(C_1)$ with just c_2 and $N_{H_1}(C_2)$ with just c_1 .

(D2) **If** no edge is monochromatic:

Stick with this proper 2-coloring of the vertices in $\bar{V}, N_{H_1}(\bar{V})$.

Remove \bar{V} from T , i.e. $T \leftarrow T \setminus \bar{V}$

Remove $\cup_i(S'_i \cup N_{H_1}(S'_i))$ and all edges containing any of these vertices, from H_1 and H_2 , and stop iterating through assignments of \bar{V} .

(D3) **If** some edge is monochromatic:

Discard the coloring and continue iterating through assignments of \bar{V} .

End While

3. For every subset B of T of size less than c :

(1) Let $A \leftarrow T \setminus B$.

(2) Using two new colors, run the proper 2-coloring algorithm for discrepancy zero hypergraphs on the $(k-1)$ -uniform hypergraph whose edge set is A .

(3) Using two new colors, iterate through all assignments of B , and attempt to 2-color $N_{H_1}(B)$ just as in Step (D1).

(4) **If** both colorings succeed:

Stick with this proper 2-coloring of the vertices in $T, N_{H_1}(B)$.

Remove from H_1 the vertices $\cup_{S' \in A} S'$ and $\cup_{S' \in B}(S' \cup N_{H_1}(S'))$

and all edges involving any of these vertices, and stop iterating through subsets of T .

(5) **If** either coloring fails:

Discard the coloring and continue iterating through subsets of T .

4. **Output** the proper partial coloring of H and the residual graph H_1 of max-degree $\binom{n-1}{k-2}t$.

4.4.3. The Result.

Theorem 4.7. For odd k , let $H = (V, E \subseteq \binom{V}{k})$ be a k -uniform discrepancy 1 hypergraph with n vertices.

Algorithm 4.4.2 partially colors H in $n^{c+O(1)}$ time, with at most $2n/ct$ colors, such that:

1. The subgraph of H induced by the colored vertices has no monochromatic vertices.

2. The subgraph of H induced by the uncolored vertices has maximum degree $\binom{n-1}{k-2}t$.

Proof. The proof goes very similarly to that of Theorem 4.4, thus we just state the key observations required to complete the proof.

(A) In any discrepancy 1 coloring of H , any size $k-1$ set S' either has discrepancy 2, or discrepancy 0.

(B) Consider any discrepancy 1 coloring of H . If a size $k-1$ set S' has discrepancy 2, then $N(S')$ is monochromatic.

- (C) At the end of any iteration of Step 2., there is no size c subset of T such that every set in the subset has discrepancy 2 in any discrepancy 1 coloring of H .
- (D) When we reach Step 3., at least $|T| - c$ sets in T , all have discrepancy 0 in EVERY discrepancy 1 coloring of H . ■

4.5. Degree Reduction under Rainbow Colorability

Now, the equivalent algorithm in the case of rainbow colorability is virtually identical to that of Section 4.4. Thus we merely state the result.

Theorem 4.8. *Let $H = (V, E \subseteq \binom{V}{k})$ be a k -uniform $(k-1)$ -rainbow colorable hypergraph with n vertices. Algorithm 4.4.2 partially colors H in $n^{c+O(1)}$ time, with at most $(k-1)n/ct$ colors, such that:*

1. *The subgraph of H induced by the colored vertices has no monochromatic vertices.*
2. *The subgraph of H induced by the uncolored vertices has maximum degree $\binom{n-1}{k-2}t$.*

References

- [A⁺77] Kazuhiko Aomoto et al. Analytic structure of schläfli function. *Nagoya Math. J*, 68:1–16, 1977.
- [AGH14] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2 + \epsilon)$ -SAT is NP-hard. 2014. To appear in FOCS '14.
- [AH05] Geir Agnarsson and Magnús M Halldórsson. Strong colorings of hypergraphs. In *Approximation and Online Algorithms*, pages 253–266. 2005.
- [AH13] Per Austrin and Johan Håstad. On the usefulness of predicates. *ACM Trans. Comput. Theory*, 5(1):1:1–1:24, May 2013.
- [AKMH96] Noga Alon, Pierre Kelsen, Sanjeev Mahajan, and Ramesh Hariharan. Approximate hypergraph coloring. *Nordic Journal of Computing*, 3(4):425–439, 1996.
- [Alo14] Noga Alon. Personal communication. 2014.
- [AM09] Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *computational complexity*, 18(2):249–271, 2009.
- [Ban10] Nikhil Bansal. Constructive algorithms for discrepancy minimization. In *Proceedings of the 51st annual IEEE symposium on Foundations of Computer Science, FOCS '10*, pages 3–10. IEEE, 2010.
- [BJH99] Károly Böröczky Jr and Martin Henk. Random projections of regular polytopes. *Archiv der Mathematik*, 73(6):465–473, 1999.
- [BK10] Nikhil Bansal and Subhash Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming, ICALP '10*, pages 250–261, 2010.

- [BPRS13] Béla Bollobás, David Pritchard, Thomas Rothvoß, and Alex Scott. Cover-decomposition and polychromatic numbers. *SIAM Journal on Discrete Mathematics*, 27(1):240–256, 2013.
- [CF96] Hui Chen and Alan M Frieze. Coloring bipartite hypergraphs. In *Proceedings of the 5th international conference on Integer Programming and Combinatorial Optimization, IPCO '96*, pages 345–358, 1996.
- [DG13] Irit Dinur and Venkatesan Guruswami. PCPs via low-degree long code and hardness for constrained hypergraph coloring. In *Proceedings of the 54th annual symposium on Foundations of Computer Science, FOCS 13*, pages 340–349, 2013.
- [DRS05] Irit Dinur, Oded Regev, and Clifford D. Smyth. The hardness of 3-Uniform hypergraph coloring. *Combinatorica*, 25(1):519–535, 2005.
- [EL75] Paul Erdos and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. *Infinite and finite sets*, 10(2):609–627, 1975.
- [GHH⁺14] Venkatesan Guruswami, Johan Håstad, Prahladh Harsha, Srikanth Srinivasan, and Girish Varma. Super-polylogarithmic hypergraph coloring hardness via low-degree long codes. In *Proceedings of the 46th annual ACM Symposium on Theory of Computing, STOC '14*, 2014.
- [GHS02] Venkatesan Guruswami, Johan Håstad, and Madhu Sudan. Hardness of approximate hypergraph coloring. *SIAM Journal on Computing*, 31(6):1663–1686, 2002.
- [GL15a] Venkatesan Guruswami and Euiwoong Lee. Strong inapproximability results on balanced rainbow-colorable hypergraphs. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '15*, pages 822–836, 2015.
- [GL15b] Venkatesan Guruswami and Euiwoong Lee. Towards a characterization of approximation resistance for symmetric CSPs. *Manuscript*, 2015.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, July 2001.
- [Hua15] Sangxia Huang. $2^{(\log N)^{1/4-o(1)}}$ hardness for hypergraph coloring. *Electronic Colloquium on Computational Complexity (ECCC)*, 15-062, 2015.
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for Max-Cut and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [KMS98] David Karger, Rajeev Motwani, and Madhu Sudan. Approximate graph coloring by semidefinite programming. *Journal of the ACM (JACM)*, 45(2):246–265, 1998.
- [Kne36] Hellmuth Kneser. Der Simplexinhalt in der nichteuklidischen Geometrie. *Deutsche Math*, 1:337–340, 1936.
- [KS14] Subhash Khot and Rishi Saket. Hardness of coloring 2-colorable 12-uniform hypergraphs with $\exp(\log^{\Omega(1)} n)$ colors. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 206–215. IEEE, 2014.
- [LM12] Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. In *Proceedings of the 53rd annual IEEE symposium on Foundations of Computer Science, FOCS 12*, pages 61–67, 2012.

- [Lov79] László Lovász. On the shannon capacity of a graph. *Information Theory, IEEE Transactions on*, 25(1):1–7, 1979.
- [McD93] Colin McDiarmid. A random recolouring method for graphs and hypergraphs. *Combinatorics, Probability and Computing*, 2(03):363–365, 1993.
- [MY05] Jun Murakami and Masakazu Yano. On the volume of a hyperbolic and spherical tetrahedron. *Communications in analysis and geometry*, 13(2):379, 2005.
- [Rog61] CA Rogers. An asymptotic expansion for certain schläfli functions. *Journal of the London Mathematical Society*, 1(1):78–80, 1961.
- [Rog64] Claude Ambrose Rogers. *Packing and covering*. Number 54. Cambridge University Press, 1964.
- [Rot14] Thomas Rothvoß. Constructive discrepancy minimization for convex sets. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 140–145. IEEE, 2014.
- [Sch58] Ludwig Schläfli. On the multiple integral $\int \dots \int dx dy \dots dz$, whose limits are $p_1 = a_1x + b_1y + \dots + h_1z > 0, p_2 > 0, \dots, p_n > 0$, and $x^2 + y^2 + \dots + z^2 < 1$. *Quart. J. Math*, 2(1858):269–300, 1858.
- [Spe85] Joel Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289(2):679–706, 1985.
- [SS13] S. Sachdeva and R. Saket. Optimal inapproximability for scheduling problems via structural hardness for hypergraph vertex cover. In *Proceedings of the 28th annual IEEE Conference on Computational Complexity, CCC '13*, pages 219–229, 2013.
- [Wen14] Cenny Wenner. Parity is positively useless. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: The 17th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 433–448. Schloss Dagstuhl, 2014.
- [Wig83] Avi Wigderson. Improving the performance guarantee for approximate graph coloring. *Journal of the ACM (JACM)*, 30(4):729–735, 1983.
- [Zwi98] Uri Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *SODA*, volume 98, pages 201–210, 1998.