Electronic Colloquium on Computational Complexity, Report No. 80 (2017)

The Quest for Strong Inapproximability Results with Perfect Completeness

Joshua Brakensiek* Venkatesan Guruswami†

Abstract

The Unique Games Conjecture (UGC) has pinned down the approximability of all constraint satisfaction problems (CSPs), showing that a natural semidefinite programming relaxation offers the optimal worst-case approximation ratio for any CSP. This elegant picture, however, does not apply for CSP instances that are perfectly satisfiable, due to the imperfect completeness inherent in the UGC. For the important case when the input CSP instance admits a satisfying assignment, in general it remains wide open to understand how well it can be approximated.

This work is motivated by the pursuit of a better understanding of the inapproximability of perfectly satisfiable instances of CSPs. Our main conceptual contribution is the formulation of a (hypergraph) version of Label Cover which we call "V label cover." Assuming a conjecture concerning the inapproximability of V label cover on perfectly satisfiable instances, we prove the following implications:

- There is an absolute constant c_0 such that for $k \geq 3$, given a satisfiable instance of Boolean k-CSP, it is hard to find an assignment satisfying more than $c_0 k^2/2^k$ fraction of the constraints.
- Given a k-uniform hypergraph, $k \geq 2$, for all $\epsilon > 0$, it is hard to tell if it is q-strongly colorable or has no independent set with an ϵ fraction of vertices, where $q = \lceil k + \sqrt{k} - \frac{1}{2} \rceil$.
- Given a k-uniform hypergraph, $k \geq 3$, for all $\epsilon > 0$, it is hard to tell if it is $(k-1)$-rainbow colorable or has no independent set with an ϵ fraction of vertices.

We further supplement the above results with a proof that an "almost Unique" version of Label Cover can be approximated within a constant factor on satisfiable instances.

*Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Email: jbrakens@andrew.cmu.edu. Research supported in part by an REU supplement to NSF CCF-1526092.

†Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213. Email: guruswami@cmu.edu. Research supported in part by NSF grants CCF-1422045 and CCF-1526092.

ISSN 1433-8092

Contents

1 Introduction **1**

1.1 Our contributions 2

1.2 Proof overview 4

1.3 Organization 5

2 Preliminaries **5**

2.1 Probability distributions 5

2.2 Influences 6

2.3 Invariance principles 7

3 V label cover **9**

3.1 Definition 9

3.2 Compatibility 10

3.3 Reduction from V label cover to P-CSP 11

4 Perfect-completeness approximation resistance and Max-k-CSP$_q$ **12**

5 Reduction to strong/rainbow hypergraph coloring **17**

5.1 Unpredictable predicates 18

5.2 Hardness results 21

A (L, s)-nearly 1–to–1 label cover **23**

B Proof of Theorem 2.7 **28**

C Existence of large 3-wise linearly independent subsets **32**

1 Introduction

The sustained progress on approximation algorithms and inapproximability results for optimization problems since the early 1990s has been nothing short of extraordinary. This has led to a sharp understanding of the approximability threshold of many fundamental problems, alongside the development of a rich body of techniques on the algorithmic, hardness, and mathematical programming aspects of approximate optimization. Yet there also remain many problems which have resisted resolution and for some there are in fact large gaps between the known algorithmic and hardness results. Examples include vertex cover, graph coloring, max-cut, feedback vertex set, undirected multicut, densest subgraph, and so on.

The Unique Games Conjecture of Khot [Kho02] postulates a strong inapproximability result for a particular class of arity two constraint satisfaction problems. This single assumption has a remarkable array of consequences, and implies *tight* inapproximability results for numerous problems including Vertex Cover [KR08], max-cut and indeed all constraint satisfaction problems (CSPs) [KKMO07, MOO10, Rag08], maximum acyclic subgraph and all ordering CSPs [GHM+11], scheduling problems [BK10, BK09], graph pricing [Lee15], and cut problems like directed multicut [Lee16], to name a few. Furthermore, for CSPs, the UGC implies that a standard semidefinite programming relaxation gives the best approximation ratio [Rag08, Rag09, BR15].

While the UGC has identified a common barrier against progress on a host of approximation problems, there are still several situations it does not apply to. Crucially, imperfect completeness, where Yes instances are only almost satisfiable, is inherent in the UGC, and this feature is inherited by the problems it reduces to. In particular, the UGC does not say *anything* about problems with *perfect completeness*, where Yes instances have a perfect solution obeying all the constraints. Important classes of such problems include satisfiable instances of CSPs (which have a perfect satisfying assignment and the goal is maximize the number of satisfied constraints) and coloring graphs/hypergraphs with approximately optimal number of colors.

Our understanding of approximating satisfiable instances of CSPs still has many gaps. Håstad's tight hardness result for approximating Max 3-SAT on satisfiable instances was much harder to prove than the analogous result for near-satisfiable instances, and was an early sign of the subtleties of ensuring perfect completeness. The approximability of satisfiable CSPs corresponds via a direct translation to the power of probabilistically checkable proof (PCP) systems with perfect completeness — the best soundness error one can achieve with a k query (non-adaptive) PCP is equal to the best inapproximability factor one can prove for a satisfiable arity k CSP. For $k = 3$, the best soundness is $5/8 + \epsilon$ for any $\epsilon > 0$, and this was established only recently via an intricate proof of the approximation resistance of satisfiable NTW (the arity 3 No-Two predicate which stipulates the number of true literals must be either $0, 1$ or 3) [Hå14]. As a basic open question that still remains wide open, we do not know the approximability of satisfiable Max NAE-3-SAT (not-all-equal 3-SAT) under any plausible (or even not so plausible!) conjecture.

The above-mentioned Unique Games hardness results consist of two components: (i) a dictatorship test that gives a way to test if a function is a dictator or is far from a dictator (e.g., has no influential coordinates), using constraints corresponding to the problem at hand (for NAE-3-SAT this would be checking if certain triples of function values are not all equal), and (ii) a reduction from Unique Games via the dictatorship test that establishes inapproximability under the UGC. The second step is standard, and it gives a "free pass" from the world of combinatorics/analysis of Boolean functions to the complexity world. When we require perfect completeness, no such conjectured off-the-shelf compiler from dictatorship tests to hardness is known (and such a passage even appears unlikely).

For instance, dictatorship tests with perfect completeness and optimal soundness are known or Max k-CSP [TY10] and Max NAE-3-SAT (folklore, and this has connections to robust forms of Arrow's theorem from social choice theory, as established using Fourier analysis [Kal02] and [O'D08, Sec. 4]). However, in both cases we do not have matching inapproximability results under any plausible conjecture.

The closest to a UGC surrogate in the literature is the d-to-1 conjecture also made in [Kho02]. The Unique Games problem is an arity 2 CSP whose constraints are bijections; the d-to-1 Label Cover is an arity 2 CSP whose constraints are d-to-1 functions. When $d \geq 2$, deciding satisfiability of a d-to-1 Label Cover instance is NP-complete, unlike Unique Games whose satisfiability is trivial to ascertain. Khot's d-to-1 conjecture states that d-to-1 Label Cover is also hard to approximate within any constant factor, even on satisfiable instances. Note that the UGC and d-to-1 conjecture are incomparable in strength; the UGC has simpler bijective constraints but the d-to-1 conjecture asserts perfect completeness which the UGC cannot.

The d-to-1 conjecture has been used to show some strong inapproximability results with perfect completeness. Such applications are, however, sporadic and also typically do not yield tight results. Some of these results are conditioned specifically on the 2-to-1 conjecture, such as a $\sqrt{2} - \epsilon$ inapproximability for vertex cover (mentioned in [Kho02] and explicit in [KMS16]), Max k-coloring with perfect completeness [GS13], and coloring 4-colorable graphs [DMR09]. The d-to-1 conjecture, for any fixed d, has been used to show the approximation resistance of NTW [OW09] and a similar result for larger arity [Hua12],[1] and finding independent sets in 2-colorable 3-uniform hypergraphs [KS14b]. Yet, the implications of the d-to-1 conjecture are limited, and it has become apparent that it is not a versatile starting point for hardness results with perfect completeness.

1.1 Our contributions

Given the above context, our work is motivated by the quest for a better starting point than 2-to-1 Label Cover for inapproximability results with perfect completeness, and which might be able to give striking consequences similar to the UGC.

Aggressive Unique Games variant. One version of Label Cover that is most similar to Unique Games, which we call (L, s)-nearly unique Label Cover, has constraint relations in[2] $[L] \times [L]$ consisting of a matching and s additional edges, for a small s that is a constant independent of L. For this version, it is NP-hard to decide satisfiability, and in fact one can give strong reductions matching the performance of dictatorship tests from it. However, this nearly unique form of Label Cover has a constant factor approximation algorithm with ratio depending only on s. We prove this result in Appendix A.

V label cover. Our main conceptual contribution is the formulation of a (hypergraph) version of Label Cover which we call "V label cover." This is an extension of 2-to-1 Label Cover, where the constraint predicates are 2-to-1 maps from $[2L]$ to $[L]$, whose relation graph can be visualized as L disjoint "V's." In V label cover of arity k, we have "longer V's" where the two branches involve k variables which coincide in single variable.[3] This is best illustrated by Figure 1 in Section 3. We put forth the *V label cover conjecture*, which asserts a strong inapproximability result for this problem. For completeness, we want an assignment where for every constraint, the k variables involved get values

[1]These were later improved to NP-hardness in [Hå14] and [Wen13].

[2]We denote $[L] = \{1, \ldots, L\}$.

[3]We should mention that our path to the formulation of V label cover was more circuitous, and has its origins in attempts to define hypergraph versions of the "α Label Cover" problem of [DMR09].

in a single "V-branch." For soundness, we insist that no assignment even *weakly satisfies* more than a tiny fraction of constraints, where a constraint is weakly satisfied if two of its k variables get values in some V-branch. [4] For this to make sense, the "junction" of the V's cannot all be on the same variable (as in 2-to-1 Label Cover), as in that case we will have a Unique Label Cover constraint between the other $(k-1)$ variables, which we can perfectly satisfy. Therefore, in our V label cover constraints, we have V's with junctions at all the k variables involved in the constraint. At a high level, this is similar to the correlation-breaking constraints of Chan [Cha16].

Near-optimal inapproximability for Max k-CSP with perfect completeness. Assuming the V label cover conjecture, we prove a near-tight inapproximability result for approximating satisfiable Max k-CSP over any fixed domain.

Theorem 1.1. *Assume the V label cover conjecture. There is an absolute constant c_0 such that for $k \geq 3$, given a satisfiable instance of Boolean k-CSP, it is hard to find an assignment satisfying more than $c_0 k^2/2^k$ fraction of the constraints. For CSP over domain size $q \geq 3$, where q is a prime power, it is hard to satisfy more than $c_0 k^3 q^3/q^k$ of the constraints.*

The approximability of Max k-CSP has been the subject of many papers in the past two decades since the advent of Håstad's optimal inapproximability results [Hås01]; a partial list includes [ST98, ST00, Eng04, HK05, EH08, ST09, GR08, AM09, Cha16, Hua14] on the hardness side, and [Tre98, Tre00, Has05, CMM09, GR08, MM14] on the algorithmic side.

The best known approximation guarantee for Max k-CSP over domain size q is $\Omega(kq/q^k)$ (for $k \geq \Omega(\log q)$, and $0.62k/2^k$ for the Boolean case [MM14]. This tight up to constant factors, due to Chan's inapproximability factor of $O(kq/q^k)$ [Cha16]. However, this hardness does not apply for satisfiable instances. For satisfiable instances, the best hardness factor is $2^{O(k^{1/3})}/2^k$ for Boolean Max k-CSP [Hua14], and $q^{O(\sqrt{k})}/q^k$ for Max k-CSP over domain size a prime q [HK05]. Note that our improved hardness factors (conditioned on the V label cover conjecture) from Theorem 1.1 are the first to get $\mathrm{poly}(k,q)/q^k$ type hardness for satisfiable instances (albeit only for prime powers) and are close to optimal. We note that satisfiable instances can be easier to approximate — Trevisan gave an elegant linear-algebra based factor $(k+1)/2^k$ approximation algorithm for satisfiable Boolean Max k-CSP [Tre00] long before Hast's $\Omega(k/2^k)$ algorithm for the general case [Has05].

Inapproximability for strong and rainbow colorable hypergraphs. Our other application of the V label cover conjecture is to hypergraph coloring, another fundamental problem where perfect completeness is crucial. We say a hypergraph is c-colorable if there is a coloring of its vertices with c colors so that no hyperedge is monochromatic. Given a 2-colorable k-uniform hypergraph for $k \geq 3$, strong inapproximability results that show the NP-hardness of coloring with any fixed ℓ number of colors are known [GHS02, DRS05], and recent developments show hardness (for $k \geq 8$) even for $\ell = \exp((\log n)^{\Omega(1)})$ where n is the number of vertices [KS14a, Var14, Hua15]. However, these results do not apply when the hypergraph has some form of balanced coloring that is stronger than just being 2-colorable. Specifically, we consider the notions of strong and rainbow colorability in this work. A hypergraph is q-strongly colorable, $q \geq k$ (resp. q-rainbow colorable, $q \leq k$) if it can be colored with q colors so that in every hyperedge, all vertices get distinct colors (resp. all q colors are represented). We refer the reader to the recent work [GL15, BG16, BGL15] for further context on these notions. When $k = q$, so that there is a perfectly balanced k-coloring where each hyperedge has exactly one vertex of each of the k colors, one can in polynomial time find a 2-coloring without any

[4] This stronger requirement in soundness is common in hypergraph versions of Label Cover. For general Label Cover the stronger soundness guarantee can be ensured with a minor loss in parameters, but for V label cover we do not know such a reduction.

monochromatic hyperedge [McD93]. Here we prove a strong hardness result for coloring hypergraphs (in fact for finding sizable independent sets), when this perfect balance condition is relaxed even slightly (specifically, $q = k - 1$ for rainbow coloring, and $q = k + o(k)$ for strong coloring).

A q-strong coloring of a hypergraph is also a legal q-coloring of the graph obtained by converting each of its hyperedges into a clique. For this reason, our hardness result for strongly colorable hypergraphs also implies hardness results in the more elementary setting of *approximate graph coloring*. There are several "pure" NP-hardness results known for graph coloring (e.g., the best known results in different regimes are [KLS00, GK04, Hua15, BG16]), but there is a gigantic gap between these results and the known algorithms. [DMR09] establishes much improved results, assuming variants of both the 2–to–1 conjecture as well as a new variant known as *alpha label cover*. Their main result is that for all $\epsilon > 0$, given a 3–colorable graph G, under these assumptions, it is NP–hard to locate an independent set with $|G|\epsilon$ vertices. In this work, assuming the V label cover–conjecture, we give a substantial generalization of this hardness.

Theorem 1.2. *Assume the V label cover conjecture.* [5]

- *Given a k-uniform hypergraph, $k \geq 2$, for all $\epsilon > 0$, it is hard to tell if it is q-strongly colorable or has no independent set with an ϵ fraction of vertices, where $q = \lceil k + \sqrt{k} - \frac{1}{2} \rceil$.*
- *Given a k-uniform hypergraph, $k \geq 3$, for all $\epsilon > 0$, it is hard to tell if it is $(k-1)$-rainbow colorable or has no independent set with an ϵ fraction of vertices.*

The authors of [GL15] showed that for any $\epsilon > 0$, it is NP-hard to distinguish if a k-uniform hypergraph (k even) is a $k/2$-rainbow colorable or does not have a independent set with ϵ fraction of the vertices. The results of [BG16] give results for strong coloring, but they only apply when $k = 2$ or when the weak coloring has only two colors. Thus, modulo the V label cover–conjecture, our results improve on those in the literature.

A path to unconditional results? In several cases, the UGC conditioned hardness results were later replaced by NP-hardness results. Examples include some geometric inapproximability results [GRSW16], hardness of Unique Coverage [GL16], inapproximability results for agnostic learning [FGRW12], tight hardness results for scheduling [SS13], Chan's breakthrough showing an asymptotically tight inapproximability result for (near-satisfiable) Max k-CSP [Cha16], etc. We hope that establishing a similar body of conditional results for perfect completeness, based on the V label cover conjecture or related variants, will point to strong inapproximability results and spur unconditional results in this domain.

1.2 Proof overview

We now briefly describe the steps needed to prove Theorem 1.1 and Theorem 1.2.

In each case, we reduce from a V label cover instance to a constraint satisfaction problem (with weighted constraints). In Section 3.3, we detail this reduction. The structure of the reduction has the same standard form as many other inapproximability results. Each vertex of the V label cover instance is replaced by a constellation of variables, known as a long code. Each hyperedge of the V label cover instance is replaced by a probability distribution of constraints between the variables in the correspond long codes. This is done carefully as to ensure that perfectly strongly satisfiable V label cover instances map to perfectly satisfiable CSPs.

[5] Technically, we need an "induced" version of the V label cover conjecture for this result.

For each problem type (Max-k-CSP, strong coloring, rainbow coloring), we craft a probability distribution which exploits its underlying structure. The probability distributions need to have a special correlation structure in order to be compatible with the V label cover constraints. We abstract a general notion termed *V label cover–compatibility* (Definition 3.1) which captures the properties common to these distributions. For example, we dictate that each vertex of each long code is sampled uniformly at random. Then, for each application, we outline the additional properties of our probability distributions in order for the reductions to have the proper soundness (Definitions 4.1 and 5.4).

For the soundness analysis, given a good approximation to the resulting CSP, we seek to find an approximate weak labeling of the original V label cover instance. To do that, we attempt to decode each long code by finding one (or many) low-degree influential coordinates; these coordinates can be viewed as candidate labels for the associated vertex. We then argue that for a sizable fraction of constraints, two of the decoded labels will belong to the a single V-branch in the constraint. We can then label our V label cover instance by assigning each vertex a label selected at random from among its decoded labels, which in expectation finds a good approximate weak labeling.

In order to guarantee these influential coordinates, we invoke a couple of invariance principles. For Max-k-CSP, we directly invoke a result due to Mossel (Theorem 2.5) on pairwise independent probability distributions. This version guarantees a common influential coordinate between *three* functions that belongs to a common "V." A pigeonhole principle then implies that two of these labels must be in the same branch. For the hypergraph coloring problems, where we do not have pairwise independence of the distributions, we generalize the invariance principles of Mossel (see Theorem 2.6) and [DMR09, Thm. 3.11] to yield a common influential coordinate for two functions that further lie on the same V-branch. This result, Theorem 2.7, is a key technical component of our reduction, which we hope will find other uses in the future.

1.3 Organization

In Section 2, we outline the necessary background on CSPs and probability spaces. In Section 3, we motivate and detail the V label cover–conjecture. In Section 4, we apply V label cover to the Max-k-CSP problem. In Section 5, we apply V label cover to the strong and rainbow hypergraph coloring problems.

In Appendix A, we show that (L, s)-nearly unique Label Cover has a polynomial-time approximation algorithm. In Appendix B, we prove Theorem 2.7. In Appendix C, we prove the existence of large subsets of a finite dimensional vector such that any three-element subset is linearly independent.

2 Preliminaries

2.1 Probability distributions

As is now commonplace in hardness of approximation reductions (e.g., [Cha16, DMR09, AM09, Mos10]), we utilize the following results on correlated probability spaces.

Definition 2.1 ([Hir35, Geb41, Ré59][6])**.** Let $X \times Y$ be a finite joint probability space with a probability

[6] See [AGKN13] for a history of this definition.

measure μ. The *correlation* between X and Y, denoted $\rho(X, Y)$ is defined to be

$$\rho(X, Y) = \sup_{\substack{f:X\to\mathbb{R}, g:Y\to\mathbb{R} \\ \mathbb{E}[f]=\mathbb{E}[g]=0, \quad \mathrm{Var}[f]=\mathrm{Var}[g]=1}} \left[\mathop{\mathbb{E}}_{(x,y)\sim\mu} [f(x)g(y)] \right].$$

This is then easily extended to the correlation of $n \geq 3$ spaces.

Definition 2.2 (Definition 1.9 of [Mos10])**.** Let $X_1 \times X_2 \times \cdots \times X_n$ be a finite joint probability space. Let $Z_i = X_1 \times X_2 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_n$. Then we define the correlation of $X_1, \ldots, X_n$ to be

$$\rho(X_1, X_2, \ldots, X_n) = \max_{1 \leq i \leq n} \rho(X_i, Z_i).$$

When a probability space can be decomposed into the product of independent subspaces, then the correlation behaves elegantly.

Lemma 2.1 (Theorem 1 of [Wit75])**.** *For all $i \in [n]$, let $X_i \times Y_i$ be a probability space with measure μ_i. Assume that $\mu_1, \ldots, \mu_n$ are independent. Then,*

$$\rho(X_1 \times X_2 \times \cdots \times X_n, Y_1 \times Y_2 \times \cdots \times Y_n) = \max_{1 \leq i \leq n} \rho(X_i, Y_i).$$

Often it can be difficult to bound the correlation of a distribution away from 1. The following result is key in reducing these complex correlation problems into rather elementary graph connectivity problems.

Lemma 2.2 (Lemma 2.9 of [Mos10])**.** *Let $X \times Y$ be a finite joint probability space with measure μ. Let G be the bipartite graph on $X \cup Y$ such that $(x, y) \in X \times Y$ is an edge iff $\Pr[x, y] > 0$ with respect to μ. Assume that G is connected, and let δ be the minimum nonzero probability in the joint distribution. Then, we have that*

$$\rho(X, Y) \leq 1 - \delta^2/2.$$

2.2 Influences

Recall the influence of a function over a probability space.

Definition 2.3. Let $X_1, \ldots, X_n$ be finite independent probability spaces, and let $f : X_1 \times \cdots \times X_n \to \mathbb{R}$ be a function. Let $Y_i = X_1 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_n$. The influence is

$$\mathrm{Inf}_i(f) = \mathop{\mathbb{E}}_{x \in Y_i} [\mathrm{Var}_{z \in X_i} f(x_1, \ldots, x_{i-1}, z, x_{i+1}, \ldots, x_n)].$$

Likewise, we need the notion of low-degree influences. We use the multilinear-polynomial definition used many times previously (e.g., [MOO10, DMR09, Mos10]).

Definition 2.4 (e.g., Definition 3.4, 3.7 of [MOO10])**.** Let $X_1, \ldots, X_n$ be finite independent probability spaces, and let $f : X_1 \times \cdots \times X_n \to \mathbb{R}$ be a function. For each $i \in [n]$, let q_i be the cardinality of the support of X_i. Let $\alpha_1^{(i)}, \ldots, \alpha_{q_i}^{(i)} : X_i \to \mathbb{R}$ be an orthonormal basis of functions such that $\alpha_1^{(i)} \equiv 1$. Let $\Sigma = [q_1] \times \cdots [q_n]$. Now, f can be uniquely expressed as

$$f = \sum_{\sigma \in \Sigma} c_\sigma \prod_{i=1}^{n} \alpha_{\sigma_i}^{(i)}.$$

for $c_\sigma \in \mathbb{R}$, which we call the Fourier coefficients. For $\sigma \in Q$, let $|\sigma| = |\{i \in [n] \mid \sigma_i \neq 1\}|$. The *low-degree influence* for $d \in [n]$ is

$$\mathrm{Inf}_i^{\le d} f = \sum_{\sigma \in \Sigma, |\sigma| \le d, \sigma_i \neq 1} c_\sigma^2.$$

The following is a key elementary fact concerning influences.

Lemma 2.3 (e.g., Proposition 3.8 [MOO10]). *Consider $f : X_1 \times \cdots \times X_n \to \mathbb{R}$. For all integers $d \ge 1$,*

$$\sum_{i=1}^n \mathrm{Inf}_i^{\le d} f \le d \operatorname{Var} f.$$

In particular, for all $\tau > 0$, $|\{i \in [n] \mid \mathrm{Inf}_i^{\le d} f \ge \tau\}| \le \frac{d \operatorname{Var} f}{\tau}$.

Sometimes, we look at f from the perspective of different marginal distributions. Consider $f : X_1 \times \cdots X_n \to \mathbb{R}$ where the X_i's are independent. Furthermore, assume that each X_i can be written as $X_i = Y_{i,1} \times \cdots Y_{i,\ell_i}$, where these $Y_{i,j}$'s are independent. Then, we let $\mathrm{Inf}_{X_i}^{\le d} f$ denote the low-degree influence of f in the ith coordinate with respect to the X_i's. Likewise, we let $\mathrm{Inf}_{Y_{i,j}}^{\le d} f$ be the influence of the (i,j)th coordinate when viewed from the perspective of $f : Y_{1,1} \times \cdots \times Y_{n,\ell_n} \to \mathbb{R}$.

For each (i,j), let $\beta_1^{(i,j)}, \ldots, \beta_{q_{i,j}}^{(i,j)} : Y_{i,j} \to \mathbb{R}$ be an orthonormal basis of functions such that $\beta_1^{(i,j)} \equiv 1$. Note that $q_i = \prod_{j=1}^{\ell_i} q_{i,j}$. Let $\Sigma' = [q_{1,1}] \times \cdots [q_{n,\ell_n}]$. Then, we have that there exist c_σ's such that $f = \sum_{\sigma \in \Sigma'} c'_\sigma \prod_{i=1}^n \alpha_{\sigma_i}^{(i)}$. If $\ell_i \le D$ for all i, then we have the following result

Lemma 2.4 (e.g., Claim 2.7 [DMR09]). *If $\ell_i \le D$ for all $i \in [n]$, then we have for all $i, d \in [n]$ that*

$$\mathrm{Inf}_{X_i}^{\le d} f \le \sum_{k=1}^{\ell_i} \mathrm{Inf}_{Y_{i,k}}^{\le Dd} f.$$

Thus, there exists $k \in [\ell_i]$ such that

$$\frac{1}{D} \mathrm{Inf}_{X_i}^{\le d} f \le \mathrm{Inf}_{Y_{i,k}}^{\le Dd} f.$$

Proof. The proof is a straightforward adaptation of the proof of Claim 2.7 in [DMR09]. □

For our applications, we only need the case $D = 2$.

2.3 Invariance principles

Like [AM09], we use the following result on pairwise independent probability spaces.

Theorem 2.5 (Lemmas 6.6, 6.9 [Mos10]). *Fix $k \ge 3$. For $1 \le i \le n$, let $\Omega_i = X_i^{(1)} \times \cdots \times X_i^{(k)}$ be finite pairwise independent probability spaces with probability measure μ_i such that the probability measures corresponding to $\mu_1, \ldots, \mu_n$ are independent. Let δ be the minimum positive probability among all the μ_i. Let*

$$\rho = \max_{1 \le i \le n} \rho(X_i^{(1)}, \ldots, X_i^{(k)})$$

and assume that $\rho < 1$. For every $\epsilon > 0$, there exists $\tau(\delta, \epsilon, \rho), d(\delta, \epsilon, \rho) > 0$ such that for any functions $f_1, \ldots, f_k$ where $f_i : X_1^{(i)} \times \cdots \times X_n^{(i)} \to [0, 1]$ if

$$\forall \ell \in [n], |\{i \mid \mathrm{Inf}^{\leq d}_{X_\ell^{(i)}} f_i > \tau\}| \leq 2$$

then

$$\left| \prod_{i=1}^{k} \mathbb{E}[f_i] - \mathbb{E}\left[\prod_{i=1}^{k} f_i\right] \right| \leq \epsilon.$$

In other words, if the product of the expected values and the expected value of the product significantly differ, then there must exist three functions with a common high low-degree influence coordinate. Note that the number "three" is crucially used in our reduction in Section 4.

As we cannot always obtain pairwise independent probability distributions (such as with our reduction to hypergraph coloring), we also need the following result on correlated probability spaces.

Theorem 2.6 (Theorem 1.14 [Mos10])**.** *Fix $k \geq 2$. For $1 \leq i \leq n$, let $\Omega_i = X_i^{(1)} \times \cdots \times X_i^{(k)}$ be a finite probability spaces with measures μ_i such that $\mu_1, \ldots, \mu_n$ are independent. Let δ be the minimum positive probability among all the μ_i. Let*

$$\rho = \max \left\{ \max_{1 \leq i \leq n} \rho(X_i^{(1)}, \ldots, X_i^{(k)}),\ \max_{\substack{1 \leq i \leq n \\ 1 \leq j < n}} \rho\left(\prod_{\ell=1}^{j} X_\ell, \prod_{\ell=j+1}^{k} X_\ell \right) \right\}.$$

and assume that $\rho < 1$. For every $\epsilon > 0$, there exists $\epsilon'(\delta, \epsilon, \rho), \tau(\delta, \epsilon, \rho) > 0$ such that for any functions $f_1, \ldots, f_k$ where $f_i : X_1^{(i)} \times \cdots \times X_n^{(i)} \to [0, 1]$ and $\mathbb{E}[f_i] \geq \epsilon$ if

$$\forall \ell \in [n], \forall i \in [k], \mathrm{Inf}_{X_\ell^{(i)}} f_i < \tau$$

then

$$\mathbb{E}\left[\prod_{i=1}^{k} f_i\right] \geq \epsilon'.$$

We need a stronger version of this theorem for our applications. We prove Theorem 2.7 in Appendix B.

Theorem 2.7. *Fix $k \geq 2$. For $1 \leq \ell \leq n$, let $\Omega_\ell = X_\ell^{(1)} \times \cdots \times X_\ell^{(k)}$ be a finite probability space with distributions μ_ℓ such that the μ_ℓ's are independent. Also, assume that for each $\ell \in [n]$ and $i \in [k]$, $X_\ell^{(i)} = \prod_{s=1}^{s_\ell^{(i)}} Y_{\ell,s}^{(i)}$, where the product is of otherwise independent distributions and $s_\ell^{(i)} \leq 2$ for all $i \in [k]$ and $\ell \in [n]$. Assume we also have the following key property*

- *If for distinct $i_1, i_2 \in [k]$ we have that $s_\ell^{(i_1)} = s_\ell^{(i_2)} = 2$, then $Y_{\ell,1}^{(i_1)}$ is independent of $Y_{\ell,2}^{(i_2)}$ (and $Y_{\ell,2}^{(i_2)}$ is independent of $Y_{\ell,1}^{(i_1)}$ by symmetry).*

For convenience of notation, if $s_\ell^{(i)} = 1$, let $Y_{\ell,2}^{(i)} := Y_{\ell,1}^{(i)}$. Let δ be the minimum positive probability among all the μ_ℓ's, $\ell \in [n]$. Let

$$\rho = \max \left\{ \max_{1 \leq \ell \leq n} \rho(X_\ell^{(1)}, \ldots, X_\ell^{(k)}),\ \max_{\substack{1 \leq \ell \leq n \\ 1 \leq j < n}} \rho\left(\prod_{\ell=1}^{j} X_\ell, \prod_{\ell=j+1}^{k} X_\ell \right) \right\}.$$

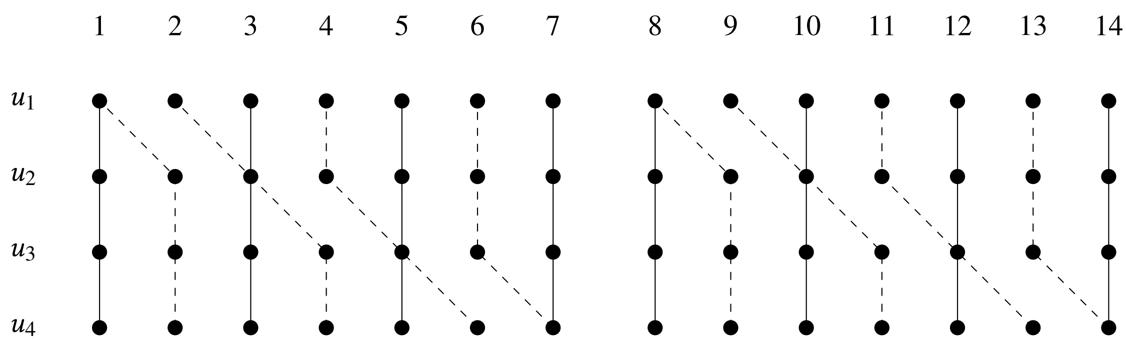

Figure 1: A schematic diagram of the branches for an edge $e = (u_1, u_2, u_3, u_4)$ of V label cover instance Ψ with parameters $k = 4$ and $L = 2$. The ith row represents $\pi_i^{(e)}$ and the jth column represents the input j. The dashed and dotted lines are to indicated the two different branches with the same values with respect to $\pi^{(e)}$. For example, we may deduce from this diagram that (10, 10, 10, 10) and (9, 10, 11, 11) are two branches of e. In particular, we have that $\pi_1^{(e)}(9) = \pi_2^{(e)}(10) = \pi_3^{(e)}(11) = \pi_4^{(e)}(11)$. Note that $\psi_i^{(e)}(j) = \perp$ exactly when the node of the ith row and jth column is at the intersection of two branches. Compare with Figure 1 of [DMR09].

and assume that $\rho < 1$. For every $\epsilon > 0$, there exists $\epsilon'(\delta, \epsilon, \rho), \tau(\delta, \epsilon, \rho), d(\delta, \epsilon, \rho) > 0$ such that for any functions $f_1, \ldots, f_k$ where $f_i : X_1^{(i)} \times \cdots \times X_n^{(i)} \to [0, 1]$ and $\mathbb{E}[f_i] \geq \epsilon$ if

$$\forall \ell \in [n], \forall s \in \{1, 2\}, |\{i \mid \mathrm{Inf}^{\leq d}_{Y^{(i)}_{\ell,s}} f_i \geq \tau\}| \leq 1$$

then

$$\mathbb{E}\left[\prod_{i=1}^{k} f_i\right] \geq \epsilon'.$$

3 V label cover

In this section, we propose a variant of hypergraph label cover which seems to plausibly have perfect completeness while also allowing for new hardness reductions. It can be thought of as a generalization of 2-to-1 label cover.

3.1 Definition

Let $k \geq 2$ and $L \geq 1$ be positive integers. An instance of k-uniform V-label cover is a k-uniform hypergraph on vertex set U. The constraints are on k-tuples $E \subseteq U^k$. Each edge $e = (u_1, \ldots, u_k)$ also has projection maps $\pi_1^{(e)}, \ldots, \pi_k^{(e)} : [(2k-1)L] \to [kL]$ with the following special property.

- The maps are surjective, in particular for all $i \in [k]$ and $j \in [kL]$,

$$|(\pi_i^{(e)})^{-1}(j)| = \begin{cases} 1 & i \equiv j \mod k \\ 2 & \text{otherwise} \end{cases}$$

In addition we would like to be able to distinguish between the two labels which map to a common value. To do this, we supplement the projection maps with *distinguishing functions* $\psi_1, \ldots, \psi_k : [(2k-1)L] \to \{0, 1, \perp\}$ such that for all $i \in [k]$, the map $x \mapsto (\pi_i^{(e)}(x), \psi_i(x))$ is injective. Furthermore, if $|(\pi_i^{(e)})^{-1}(\pi_i^{(e)}(x))| = 1$, then we define $\psi_i(x) = \perp$, and otherwise $\psi_i(x) \in \{0, 1\}$. We say that $(t_1, \ldots, t_k) \in [(2k-1)L]^k$ is a *branch* of e if there is $\ell \in [kL]$ and $b \in \{0, 1\}$ such that for all i, $(\pi_i^{(e)}(t_i), \psi_i^{(e)}(t_i))$ equals (ℓ, b) or $(\ell, \perp)$. Note that for each branch, there is exactly one $j \in [k]$ such that $\psi_j^{(e)}(t_j) = \perp$. In fact such such an index satisfies $j \equiv \pi_i^{(e)}(t_i) \mod k$ for all j. We say that i is the *junction* of the branch.

To better understand the setup, see Figure 1.

The goal of V-label cover is to produce a labeling of the vertices $\sigma : U \to [(2k-1)L]$. We say that a hyperedge $e = (u_1, \ldots, u_k)$ is *strongly satisfied* if $(\sigma(u_1), \ldots, \sigma(u_k))$ is a branch. In other words, for all $i, j \in [k]$, $\pi_i^{(e)}(\sigma(u_i)) = \pi_j^{(e)}(\sigma(u_j))$ *and* either $\psi_i^{(e)}(\sigma(u_i)) = \psi_j^{(e)}(\sigma(u_j)) \neq \perp$ or exactly one of $\psi_i^{(e)}(\sigma(u_i)), \psi_j^{(e)}(\sigma(u_j))$ is $\perp$. Another way to express this is that $(\pi_i^{(e)}(\sigma(u_i)), \psi_i^{(e)}(\sigma(u_i)))$ is uniform except for one i for which $\psi_i^{(e)}(\sigma(u_i)) = \perp$ (the meeting point in the 'V' of the two branches).

We say the hyperedge is *weakly satisfied* if for some distinct $i, j \in [k]$, $\pi_i^{(e)}(\sigma(u_i)) = \pi_j^{(e)}(\sigma(u_j))$ *and $\sigma(u_i)$ and $\sigma(u_j)$ are in the same branch.*

We now formally state our conjectured intractability of approximating V label cover. Below we state an "induced" version where in the soundness guarantee, for every labeling, most of the hyperedges within any subset of vertices of density ϵ fail to be weakly satisfied. The induced version is needed for our reduction to hypergraph coloring (this is similar to the α conjecture of [DMR09] which was also defined in an induced form). For our Max k-CSP result, it suffices to assume the soundness condition that at most ϵ fraction of edges are weakly satisfiable. For simplicity, we only state the stronger induced version below.

Conjecture 3.1 (V label cover–conjecture, induced version)**.** *For all $k \geq 2$ and $\epsilon > 0$, there exists an $L \geq 1$ such that for any k-uniform V label cover instance Ψ on label set L and vertex set U and hyperedge set E, it is NP-hard to distinguish between*

- *YES: There exists a labeling for which every hyperedge is strongly satisfied.*
- *NO: For every labeling and every subset $U' \subset U$ with $|U'| \geq |U|\epsilon$, less than ϵ fraction of the edges in $(U')^k \cap E$ are weakly satisfied by the labeling.*

3.2 Compatibility

Consider a domain size $q \geq 2$, an arity $k \geq 2$, and a predicate $P \subseteq [q]^k$. In order to understand the "V label cover–hardness" of this predicate P, for each edge $e = (u_1, \ldots, u_k)$ of our V label cover instance we seek to construct probability distributions on $[q]^{k \times (2k-1)L}$ such that the marginal distribution of each branch of e is supported by P. We define the notion of *V label cover–compatibility* in order to capture exactly what we need.

Definition 3.1. For a predicate $P \subseteq [q]^k$, consider $\mu_1, \ldots, \mu_k$ supported on P^2. For $i, j \in [k]$, let $X_{i,j} \sim [q]^2$ be the marginal distribution of μ_i on the jth coordinates. That is, for all $(a, b) \in [q]^2$,

$$\Pr_{(x', y') \sim X_{i,j}} [(x', y') = (a, b)] = \Pr_{(x, y) \sim \mu_i} [(x_j, y_j) = (a, b)].$$

We call the distributions $\mu_1, \dots, \mu_k$ a *V label cover–compatible family* if they satisfy the following properties.

1. For all $i \in [k]$, $X_{i,i}$ is uniform on $\{(a,a) \mid a \in [q]\}$.
2. For all $i, j \in [k]$ with $i \neq j$ and $X_{i,j}$ is uniform on $[q]^2$.
3. For all $i \in [k]$, $\rho(\mu_i) < 1$, which we define to be

$$\rho(\mu_i) := \rho(X_{i,1}, \dots, X_{i,k}).$$

We say that P is *V label cover–compatible* if a V label cover–compatible family $\mu_1, \dots, \mu_k$ exists.

The reason we have k different distributions is because the two connected branches can intersect in k different rows (see Figure 1).

Property (3) of Definition 3.1 precludes any algebraic structure in our predicate that would permit a polynomial-time algorithm. For example, the uniform distribution on the predicate $\{x \in \mathbb{Z}_2^n \mid x_1 + \cdots + x_n = 0\}$ has correlation 1 and allows for Gaussian-elimination to solve exactly.

3.3 Reduction from V label cover to P-CSP

Let $P \subseteq [q]^k$ be a predicate for $q, k \geq 2$ which is V label cover–compatible with distributions $\mu_1, \dots, \mu_k$. In this section, we show how to reduce an arbitrary instance of V label cover into an instance of P-CSP, the constraint satisfaction problem where all clauses are of the form $(x_1, \dots, x_n) \in P$. Furthermore, we assign weights to the clauses of this CSP, in which the weights are determined by these distributions μ_i. This reduction is the starting point for showing the conditional NP-hardness results in Sections 4 and 5.

Let $\Psi = (U, E, L, \{\pi_i^{(e)}\}_{e \in E, i \in [k]}, \{\psi_i^{(e)}\}_{e \in E, i \in [k]})$ be our instance of k-uniform V label cover. For each $u \in U$, we construct $q^{(2k-1)L}$ variables $x_s^{(u)}$, where $s \in [q]^{(2k-1)L}$. Now, for every edge $e = (u_1, \dots, u_k) \in E$ and every $s^{(1)}, \dots, s^{(k)} \in [q]^{(2k-1)L}$ with the following property

- For any $t_1, \dots, t_k \in [(2k-1)L]$ such that $(t_1, \dots, t_k)$ is a branch of e, we have $(s_{t_1}^{(1)}, \dots, s_{t_k}^{(k)}) \in P$,

we add the constraint $(x_{s^{(1)}}^{(u_1)}, \dots, x_{s^{(k)}}^{(u_k)}) \in P$. Looking back at Figure 1, we have that any assignment of values from $[q]$ to the nodes of the schematic such that each branch is an element of P corresponds to some choice $(s^{(1)}, \dots, s^{(k)})$.

Let Φ be the resulting instance. Although we have described the clauses, we have not yet determined the relative weights of the clauses.

Claim 3.2. *If Ψ has a labeling $\sigma : U \to [(2k-1)L]$ which strongly satisfies every hyperedge, then we have that Φ has a perfect satisfying assignment. In other words, this reduction has perfect completeness.*

Proof. For each $u \in U$, and $s \in [q]^{(2k-1)L}$, we let $x_s^{(u)} = s_{\sigma(u)}$. One can verify this assignment satisfies Φ. □

Now, fix $e = (u_1, \dots, u_k) \in E$. For each $\ell \in [kL]$, let $(a_1, \dots, a_k), (b_1, \dots, b_k)$ be the two branches of e such that $\pi_i^{(e)}(a_i) = \pi_i^{(e)}(b_i) = \ell$ for all i. Let $j \in [k]$ be the unique index for which $a_j = b_j$, (i.e., j is the junction). Let I be the index set $I := \{(i, a_i) \mid i \in [k]\} \cup \{(i, b_i) \mid i \in [k]\}$; note that $|I| = 2k - 1$. Let $\Omega_\ell^{(e)} \sim [q]^I$ be the probability distribution isomorphic to μ_j such that the marginals $x_1, \ldots, x_k, y_1, \ldots, y_k$ of μ_j correspond to the marginals indexed by $(1, a_1), \ldots, (k, a_k), (1, b_1), \ldots, (k, b_k)$ of $\Omega_\ell^{(e)}$.

Let

$$\nu^{(e)} := \prod_{\ell \in [kL]} \Omega_\ell^{(e)},$$

where the product is over independent distributions. Note that the support of $\nu^{(e)}$ can be identified with $[q]^{[k]\times[(2k-1)L]}$ since each $(i, a_i) \in [k] \times [(2k-1)L]$ is accounted for in some branch. We let $Y_j^{(e,i)}$ be the marginal distribution of coordinate $(i, j) \in [k] \times [(2k-1)L]$ of $\nu^{(e)}$. For any $i \in [k]$ and $\ell \in [kL]$, we let $X_{i,\ell}^{(e)}$ be the marginal distribution on the indices $\{(i, t) \mid \pi_i^{(e)}(t) = \ell\}$. In particular, if i is a junction, the meeting point of the branches, then $Y_t^{(e,i)} = X_{i,\ell}^{(e)}$. Otherwise, $X_{i,\ell}^{(e)}$ is the product of two Y's:

$$X_{i,\ell}^{(e)} = \prod_{t \in (\pi_i^{(e)})^{-1}(\ell)} Y_t^{(e,i)}.$$

This distribution $\nu^{(e)}$ specifies the probability distribution of the clauses corresponding to a particular edge of the label cover instance. These probabilities are the relative weights of the clauses in the instance.

4 Perfect-completeness approximation resistance and Max-k-CSP$_q$

A natural question to ask concerning V label cover is if it reduces to natural families of predicates which are hard to approximate, even when guaranteed perfect completeness. In the case of imperfect completeness, Austrin and Mossel [AM09] showed assuming the Unique Games Conjecture that if a predicate $P \subseteq [q]^k$, for some finite domain size q, supports a balanced pairwise independent distribution, then P is *approximation resistant*. That is, for all $\epsilon > 0$, it is NP-hard to distinguish between $1 - \epsilon$-satisfiable and $\frac{|P|}{q^k} + \epsilon$-satisfiable P-CSPs. Only a few years later, in a breakthrough by Chan [Cha13], unconditional approximation resistance was shown for any P which supports a balanced pairwise independent subgroup. We hope that establishing a similar conditional results for perfect completeness will spur unconditional results in this domain.

In order to reduce from V label cover, we need a more stringent criteria than merely supporting a balanced pairwise independent distribution. We call these more structured distributions *pairwise-independent V label cover-compatible*.

Definition 4.1. Let $q \geq 2, k \geq 3$ be parameters. Let $P \subseteq [q]^k$ be a predicate. We say that P is *pairwise–independent V label cover–compatible* if there exists a V label cover–compatible family $\mu_1, \ldots, \mu_k$ supported on P^2 (with marginals $X_{i,j}$, $i, j \in [k]$) with the additional property that

4. For all $i \in [k]$ and $j \neq j' \in [k]$, we have that $X_{i,j}$ and $X_{i,j'}$ are pairwise independent.

To motivate the definition, one way to view property (4), when combined with properties (1) and (2) of Definition 3.1, is that P does not just support a pairwise independent distribution, but that the distribution can preserve pairwise independence even when conditioning on the value of a coordinate.[7] Assuming the V label cover-conjecture, this property suffices to establish perfect-completeness approximation resistance if we allow what are known as *folded* predicates.[8] Assume that $[q]$ has a + operator (e.g., addition modulo q). We specify that we may use folded versions of our predicate P to be the predicates

$$a \in [q]^k,\ P^{(a)} := \{(x_1 + a_1, \ldots, x_k + a_k) \mid (x_1, \ldots, x_k) \in P\}.$$

Each $P^{(a)}$ has the same cardinality, so incorporating these extra predicates can only increase the severity of the hardness of approximation. Thus, more precisely we say that the *family* of predicates $\{P^{(a)} \mid a \in [q]^k\}$ is perfect-completeness approximation resistant. That is, for every $\epsilon > 0$, it is NP-hard to distinguish whether a CSP with predicates from $\{P^{(a)} \mid a \in [q]^k\}$ is perfectly satisfiable or is $\frac{|P|}{q^k} + \epsilon$ satisfiable.

Theorem 4.1. *Let $P \subseteq [q]^k$ be a predicate which supports a pairwise-independent V label cover–compatible distribution. Then, assuming the V label cover–conjecture, we have that the collection of predicates $\{P^{(a)} \mid a \in [q]^k\}$ is perfect-completeness approximation resistant.*

Proof. The high-level structure of our proof is analogous to that of Austrin and Mossel [AM09]. The proof proceeds in a couple of stages. First, we describe the reduction from a V label cover instance to an instance of P-CSP, and note that such a reduction preserves perfect completeness. Second, we analyze the soundness of our reduction using Theorem 2.5 to show that if our P-CSP can be well-approximated, then our original V label cover instance also admits an approximation.

Reduction. Let $\Psi = (U, E, L, \{\pi_i^{(e)}\}_{e \in E, i \in [k]}, \{\psi_i^{(e)}\}_{e \in E, i \in [k]})$ be our instance of k-uniform V label cover. Let Φ be the instance of P-CSP guaranteed by the construction in Section 3.3. Let $\nu^{(e)} \in [q]^{[k] \times [(2k-1)L]}$ be the weighting distributions on the clauses corresponding to the hyperedges. Let $\Omega_\ell^{(e)}, X_{i,j}^{(e)}, Y_j^{(e,i)}$ be the marginal distributions described in Section 3.3. By Claim 3.2, our reduction has perfect completeness.

We now modify the CSP Φ into a new CSP Φ' which incorporates folding. For each constraint $(x^{(u_1)}_{s^{(1)}}, \ldots, x^{(u_k)}_{s^{(k)}}) \in P$ and for each $i \in [k]$, let $(s^{(i)})' = s^{(i)} - s^{(i)}_1$ (i.e., subtract $s^{(i)}_1$ from every coordinate). Then, we specify that

$$(x^{(u_1)}_{(s^{(1)})'}, \ldots, x^{(u_k)}_{(s^{(k)})'}) \in P^{(s^{(1)}_1, \ldots, s^{(k)}_1)}.$$

One may check that this modification preserves perfect completeness.

Soundness. We view an assignment to Φ' as a collection of functions $\mathcal{F} = \{f_u : [q]^{(2k-1)L} \to [q] \mid u \in U\}$, where $f_u(s)$ is the assigned value for x^u_s. Because of our modification to the CSP, we only specify constraints for $f_u(s)$ when $s_1 = q$. Thus, we may assume that each f_u is *folded*. That is, $f_u(s) + a \equiv f_u(s + (a, \ldots, a)) \mod q$ for all $a \in [q]$. One may check that the f_u's satisfy a clause in Φ' if and only if they satisfy the corresponding clause in Φ. Thus, it is equivalent to focus on the f_u's satisfaction of Φ.

For $a \in [q]$, we let

$$f_u^{(a)}(x) = \begin{cases} 1 & f_u(x) = a \\ 0 & \text{otherwise} \end{cases}$$

[7]The definition permits a slightly broader class of P (i.e., the distribution can change depending on which coordinate is conditioned on), but our applications will construct P of the type specified in the motivation.

[8]This is a standard assumption in the CSP literature, e.g., [AM09].

We define the influences and low-degree influences (Definitions 2.3 and 2.4) of the $f_u^{(a)}$'s to be with respect to the uniform distribution.

Let $\Phi(\mathcal{F})$ be the fraction of constraints of Φ satisfied by $\mathcal{F}$, using the weights specified by the $\nu^{(e)}$ distributions. We seek to show for any $\epsilon > 0$ if there exists a $\mathcal{F}$ such that $\Phi(\mathcal{F}) > \frac{|P|}{q^k} + \epsilon$, then there exists $\delta > 0$ and $\sigma : U \to [(2k-1)L]$ such that σ weakly satisfies δ fraction of the constraints of Ψ.

It is evident from the construction, that a group of constraints are associated with each $e \in E$. Let $e(\mathcal{F})$ be the fraction of constraints corresponding to ϕ satisfied by $\mathcal{F}$ (that is the measure with respect to $\nu^{(e)}$ of the clauses satisfied by $\mathcal{F}$). We have that

$$\Phi(\mathcal{F}) = \frac{1}{|E|}\sum_{e\in E} e(\mathcal{F}).$$

Thus, if $\Phi(\mathcal{F}) > \frac{|P|}{q^k} + \epsilon$, there exists a subset $E' \subseteq E$ such that $|E'| > (\epsilon/2)|E|$ and $e(\mathcal{F}) \geq \frac{|P|}{q^k} + \epsilon/2$ for all $e \in E'$; as otherwise, $\Phi(\mathcal{F}) \leq \epsilon/2 \cdot 1 + (1-\epsilon/2)\cdot\left(\frac{|P|}{q^k} + \epsilon/2\right) < \frac{|P|}{q^k} + \epsilon$.

Fix, $e = (u_1, \ldots, u_k) \in E'$. Note that

$$\begin{aligned} e(\mathcal{F}) &= \mathop{\mathbb{E}}_{(s_1,\ldots,s_k)\sim\nu^{(e)}}[(f_{(u_1)}(s_1), \ldots, f_{(u_k)}(s_k)) \in P] \\ &= \sum_{r\in P} \mathop{\mathbb{E}}_{(s_1,\ldots,s_k)\sim\nu^{(e)}}[f_{u_1}^{(r_1)}(s_1)\cdots f_{u_k}^{r_k}(s_k)]. \end{aligned}$$

Thus, for some $r \in P$, we have that

$$\mathop{\mathbb{E}}_{(s_1,\ldots,s_k)\sim\nu^{(e)}}[f_{u_1}^{(r_1)}(s_1)\cdots f_{u_k}^{r_k}(s_k)] > \frac{1}{q^k} + \frac{\epsilon}{2|P|}.$$

Let $\epsilon' = \epsilon/(2|P|) > 0$. Also, for all $i \in [k]$, let $\Pi_i^{(e)} = \prod_{\ell=1}^{kL} X_{i,\ell}^{(e)}$. Since each $\Pi_i^{(e)}$ is uniform and f_{u_i} is folded, we have that

$$\mathop{\mathbb{E}}_{s_i\sim\Pi_i^{(e)}}[f_{u_i}^{(r_i)}(s_i)] = \frac{1}{q}.$$

In particular, this implies that

$$\left|\prod_{i=1}^{\ell}\mathbb{E}[f_{u_i}^{(r_i)}(s_i)] - \mathbb{E}\left[\prod_{i=1}^{\ell} f_{u_i}^{(r_i)}(s_i)\right]\right| > \epsilon'.$$

Note that $\nu^{(e)} = \Omega_1^{(e)} \times \cdots \times \Omega_{kL}^{(e)}$ meets the requirements of Theorem 2.5. Thus, there exists $\tau, d > 0$, which are functions of only ϵ' and parameters of $|P|$, such that

$$\exists \ell \in [kL], |\{i : \mathrm{Inf}_{X_{i,\ell}^{(e)}}^{\leq d} f_{u_i}^{(r_i)} > \tau\}| \geq 3.$$

Let $i_1, i_2, i_3 \in [k]$ be three of these coordinates and let $\ell \in [kL]$ be the guaranteed value of ℓ. Observe that we can also write $\Pi_{i_a}^{(e)}$ as

$$\Pi_{i_a}^{(e)} = \prod_{t\in[(2k-1)L]} Y_t^{(e,i_a)}.$$

Note that each $X^{(e)}_{i_a,\ell}$ can be written as the product distribution of at most 2 $Y^{(e,i_a)}_t$'s, where $\pi^{(e)}_{i_a}(t) = \ell$. By invoking Lemma 2.4 with $D = 2$, we have that there exists t_1, t_2, t_3 such that $\pi^{(e)}_{i_a}(t_a) = \ell$ for all $a \in \{1, 2, 3\}$ and

$$\mathrm{Inf}^{\leq 2d}_{Y^{(e,i_a)}_{t_a}} f^{(r_i)}_{u_{i_a}} = \mathrm{Inf}^{\leq 2d}_{t_a} > \frac{\tau}{2},$$

where the equality is due to the fact that the $Y^{(e,i_a)}_{t_a}$ distributions are uniform distributions on $[q]$.

Note that since each 'component' of (e) has two branches, by the Pigeonhole principle, some two of $\{t_1, t_2, t_3\}$ are in the same branch. Thus, any assignment σ for which $\sigma(u_{i_a}) = t_a$ for all $a \in \{1, 2, 3\}$ weakly satisfies e.

For each $u \in U$. Let $S_u \subseteq [(2k-1)L]$ be the set of labels j for which $\mathrm{Inf}^{\leq 2d}_j f^{(a)}_u > \tau/2$ for some $a \in [q]$. Since $\mathrm{Var}\, f^{(a)}_u \leq \max(f^{(a)}_u)^2 = 1$, we have by Lemma 2.3 that $|S_u| \leq 4dq/\tau$, which is independent of L. Construct a random labeling $\sigma : U \to [(2k-1)L]$ by sampling each $\sigma(u)$ from S_u independently and uniformly at random (if S_u is empty, let $\sigma(u) = 1$). For each $e \in E'$, we established that there exists $i, i' \in [k]$ and $\ell \in S_{u_i}$ and $\ell' \in S_{u_{i'}}$ such that setting $\sigma(u_i) = \ell$ and $\sigma(u_{i'}) = \ell'$ weakly satisfies e. Thus, in expectation at least

$$\frac{|E'|}{|E|} \cdot \frac{1}{(\max |S_u|)^2} = \frac{\tau^2 \epsilon}{16 d^2 q^2} > 0$$

of the edges are weakly satisfied. Note that this expression is independent of L and the size of Ψ, as desired. □

We use this theorem to obtain hardness of approximation results for Max-k-CSP$_q$ when $q \geq 2$ is a prime power.

Lemma 4.2. *For all $q \geq 2$ a prime power and $k \geq 2$, there exists $P \subseteq [q]^k$ which is pairwise–independent V label cover-compatible with $|P| = 2k^3q^3$.*

Proof. We use a modification of the constructions of [AM09] and [TY10]. Let $\ell \geq 3$ be the least odd integer such that $q^{(\ell-1)/2} \geq k$. Thus, $q^\ell \leq k^2 q^3$. View $\mathbb{F}_q^\ell$ as a vector space over $\mathbb{F}_q$. By Lemma C.1 there exists $S \subset \mathbb{F}_q^\ell$ with $|S| \geq q^{(\ell-1)/2} \geq k$ such that S is 3-wise linearly independent (i.e., every 3-element subset is linearly independent). Let $v^{(1)}, \ldots, v^{(k)} \in S$ be k distinct elements from this set. Define $\langle \cdot, \cdot \rangle$ to be the canonical bilinear form on $\mathbb{F}_q^\ell$. That is, $\langle x, y \rangle = \sum_{i=1}^{\ell} x_i y_i$.

We give an initial attempt to construct our predicate. Let[9]

$$P_0 = \{(\langle v^{(1)}, X \rangle, \ldots, \langle v^{(k)}, X \rangle) : X \in \mathbb{F}_q^\ell\}.$$

We have that $|P_0| \leq q^\ell \leq k^2 q^3$. We show that P_0 satisfies properties (1), (2), and (4). Note that the definition of P_0 defined a natural probability distribution μ. It is clear that μ has uniform marginal distributions (since each $v^{(i)}$ is nonzero and X is uniform). Furthermore, the marginal distributions are 3-wise independent (and thus 3-wise uniform) since the $v^{(i)}$'s are 3-wise linearly independent. (We omit the proof, a similar result for pairwise independence is Lemma 4.2 of [AM09].)

Now, fix $i \in [k]$, define μ_i to be

$$\mu_i := \{x, y \sim \mu \text{ independent} : x_i = y_i\}.$$

[9]Note that we identify $[q]$ with $\mathbb{F}_q$ in some canonical way.

Let $X_{i,j}$ with $j \in [k]$ be the marginals of μ_i. We seek to show μ_i satisfies properties (1), (2), and (4). Property (1) follows immediately from the uniform marginals of μ. Now, fix $j \neq i$, since (x_i, x_j) and $(x_i = y_i, y_j)$ are uniform distributions and x_j and y_j are conditionally independent given x_i, we have that

$$\Pr[x_i \wedge x_j \wedge y_j] = \Pr[x_j \wedge y_j | x_i] \Pr[x_i] = \Pr[x_j | x_i] \Pr[y_j | x_i] \Pr[x_i] = \Pr[x_j] \Pr[y_j] \Pr[x_i].$$

Therefore, (x_i, x_j, y_j) is uniform on $\mathbb{F}_q^{\ell}$. Thus, property (2) and the case $j' = i$ of property (4) follow.

To finish establishing property (4), consider $j \neq j' \in [k] \setminus \{i\}$. We seek to show that $(x_j, x_{j'}, y_j, y_{j'})$ is uniform for which it suffices to show that $(x_i, x_j, x_{j'}, y_j, y_{j'})$ is uniform. Like before,

$$\begin{aligned} \Pr[x_i \wedge x_j \wedge x_{j'} \wedge y_j \wedge y_{j'}] &= \Pr[x_j \wedge x_{j'} | x_i] \Pr[y_j \wedge y_{j'} | x_i] \Pr[x_i] \\ &= \Pr[x_j] \Pr[x_{j'}] \Pr[y_j] \Pr[y_{j'}] \Pr[x_i] \text{ (3-wise independence of } \mu). \end{aligned}$$

Thus, the μ_i's satisfy properties (1), (2), and (4). Sadly, due to the nice algebraic structure of P_0, we have that $\rho(\mu_i) = 1$ for all i. To rectify this, we create a 'noisy' version of P_0. For $x \in \mathbb{F}_q^k$, let $|x|$ be the number of nonzero coordinates of x. Then, we define P to be

$$P := \{x \in \mathbb{F}_q^k \mid \exists y \in P_0, |x - y| \leq 1\}.$$

Note that $|P| \leq (k+1)|P_0| \leq 2k^3 q^3$. Now, modify the μ_i's to get μ_i''s by the following procedure.

1. Sample $(x, y) \in \mu_i$.
2. Sample $j \in [k]$ and $a, b \in \mathbb{F}_q$ uniformly.
3. If $i = j$, set $x_i = y_j = a$. Otherwise, set $x_i = a$ and $y_j = b$.

Clearly the support of μ_i' is P^2. Also μ_i' preserves properties (1), (2), and (4) of being V label cover-compatible since re-randomizing coordinates can only assist in maintaining pairwise independent distributions.

It remains to show that μ_i' satisfies property (3). The proof of this is similar to that of Lemma 4.6 of [TY10]. Let

$$Z_{i,j} := \prod_{j=1, j \neq i}^{k} X_{i,j}.$$

It suffices to show that $\rho(X_{i,j}, Z_{i,j}) < 1$. To do that, it suffices to show by Lemma 2.2 that the bipartite graph whose edges are the support of $X_{i,j} \times Z_{i,j}$ is connected. For any $(\alpha, \beta) \in X_{i,j} \times Z_{i,j}$, since with nonzero probability the jth coordinate is rerandomized, we have that $(\alpha', \beta) \in X_{i,j} \times Z_{i,j}$ for all α' in the support of $X_{i,j}$. From this connectivity immediately follows.

Therefore, P has the desired properties. □

Using the same proof techniques, we have the following corollary.

Corollary 4.3. *For $q = 2$ and all $k \geq 2$, there exists $P \subseteq [2]^k$ which is pairwise–independent V label cover–compatible and $|P| = O(k^2)$.*

Proof. Repeat the proof of Lemma 4.2, but note that $S = \{x \in \mathbb{F}_2^{\ell} : \sum_{i=1}^{\ell} x_i = 1\}$ is a 3-wise-independent subset of size $2^{\ell-1}$. □

Now we may obtain Theorem 1.1.

Proof of Theorem 1.1. The case $q = 2$ follows immediately from Corollary 4.3 and Theorem 4.1. Similarly, if $q \geq 3$ is a prime power, then the result follows from Lemma 4.2 and Theorem 4.1. $\square$

Remark. If q is not a prime power, we cannot invoke the monotonicity result of [AM09, Cor. B.1], since they crucially assume a lack of perfect completeness. In fact, their reduction does not even produce instances which are near-perfectly satisfiable. If for a general q, we can find a distribution $\mu \sim [q]^k$ whose support is of size $\text{poly}(q, k)$, has uniform marginals, and has 3-wise independence, then by Theorem 4.1 we can extend our result to Max-k-CSP_q.

This is the first conditional NP-hardness reduction which obtains a soundness of $\frac{\text{poly}(q,k)}{q^k}$ for even one fixed q. Previously, a long code test due to Tamaki and Yoshida [TY10] obtained $\frac{O(k)}{2^k}$ for when $q = 2$. The currently best known unconditional result for Max-k-CSP_2 is $\frac{2^{O(k^{1/3})}}{2^k}$ due to Huang [Hua14]. For $q \geq 3$, the best known result is [HK05] [MM14].

Remark. Using a modification of the predicate of [TY10], we speculate that it is possible to improve the hardness factor for Boolean Max-k-CSP to $O(k/2^k)$.

5 Reduction to strong/rainbow hypergraph coloring

Recall the notions of strong and rainbow graph coloring [GL15, BG16, BGL15].

Definition 5.1. Let $H = (V, E)$ be a hypergraph of uniformity $k \geq 2$. Let $q \geq k$ be a positive integer. A function $\chi : V \to [q]$ is a (k, q)-*strong coloring* of H if for all $e \in E$, $\chi \restriction e$ is an injection. In other words, no two vertices in the same hyperedge receive the same color.

Definition 5.2. Let $H = (V, E)$ be a hypergraph of uniformity $k \geq 2$. Let $q \leq k$ be a positive integer. A function $\chi : V \to [q]$ is a (k, q)-*rainbow coloring* of H if for all $e \in E$, $\chi \restriction e$ is a surjection. That is, for all $e \in E$ and $c \in [q]$, there is $v \in e$ such that $\chi(v) = c$.

Note that the notions of strong and rainbow coloring coincide when $k = q$. In these hypergraphs, we would like to know if we can tractably identify large *weak independent sets*.

Definition 5.3. Let $H = (V, E)$ be a hypergraph. A subset $W \subseteq V$ is an *weak independent set,* if for all $e \in E$, $e \cap W \neq e$.

Theorem 5.1. *Assume the induced version of the V label cover-conjecture (Conjecture 3.1). For all $k \geq 2$, $q > k + \sqrt{k} - \frac{1}{2}$ and $\epsilon > 0$, given a k-uniform hypergraph $H = (V, E)$, it is NP-hard to distinguish between the following two settings.*

- *YES: H admits a (k, q)-strong coloring.*
- *NO: H does not have a weak independent set of density ϵ ($|V|\epsilon$ vertices).*

Theorem 5.2. *Assume the induced version of the V label cover-conjecture (Conjecture 3.1). For all $k > q \geq 2$ and $\epsilon > 0$, given a k-uniform hypergraph $H = (V, E)$, it is NP-hard to distinguish between the following two settings.*

- *YES: H admits a (k, q)-rainbow coloring.*

- *NO: H does not have a weak independent set of density ϵ ($|V|\epsilon$ vertices).*

We can view strong and rainbow hypergraph coloring as CSPs. In particular, let

$$S_{k,q} = \{(c_1, \ldots, c_k) \in [q]^k \mid \forall i, j \in [k], \text{ if } i \neq j \text{ then } c_i \neq c_j\}$$

be the strong coloring predicate, and let

$$R_{k,q} = \{(c_1, \ldots, c_k) \in [q]^k \mid \forall c \in [q], \exists i \in [k], c = c_i\}$$

be the rainbow coloring predicate.

These predicates have structure which we call *unpredictable*.

5.1 Unpredictable predicates

In this section, we supplement Definition 3.1 to give our distributions additional properties that we need for our hardness reduction.

Definition 5.4. Let $q, k \geq 2$ be parameters. Let $P \subseteq [q]^k$ be a predicate. We say that P is *unpredictably V label cover–compatible* if there exists a V label cover–compatible family $\mu_1, \ldots, \mu_k$ supported on P^2 (with marginals $X_{i,j}$, $i, j \in [k]$) with the additional properties that:

4. For all $i \in [k]$ and $1 \leq j < k$, we have that

$$\rho\left(\prod_{\ell=1}^{j} X_{i,\ell}, \prod_{\ell=j+1}^{k} X_{i,\ell}\right) < 1.$$

5. Each $i \in [k]$ and $j_1, j_2 \in [k] \setminus \{i\}$ with $j_1 \neq j_2$, we have that the marginal distribution of (x_{j_1}, y_{j_2}) in μ_i (recall that x_{j_1} and y_{j_2} are in separate 'branches' of μ_i) is uniform over $[q]^2$.

As the properties are rather technical, the following definition helps to streamline our understanding.

Definition 5.5 (c.f., Section 1.4 [Mos10])**.** Let $\Omega = X^{(1)} \times \cdots \times X^{(k)}$ be a probability space. We say that Ω is *connected* if for all atoms (elements with nonzero probability) $x, y \in \Omega$, there exists a sequence $z_0, \ldots, z_n \in \Omega$ of atoms such that $x = z_0, y = z_n$, and z_i and z_{i-1} differ in exactly one of the k coordinates for all $i \in [n]$.

The following lemma demonstrates the utility of connected predicates.

Lemma 5.3. *If P admits a family $\mu_1, \ldots, \mu_k$ (with marginals $X_{i,j}$, $i, j \in [k]$) of probability distributions such that they are connected. Then P satisfies property (3) of Definition 3.1 and property (4) of Definition 5.4.*

Proof. First we verify property (4) of Definition 5.4. Fix $i \in [k]$. It suffices to check for all $1 \leq j < k$ that

$$\rho\left(\prod_{\ell=1}^{j} X_{i,\ell}, \prod_{\ell=j+1}^{k} X_{i,\ell}\right) < 1.$$

By Lemma 2.2, it suffices to check that the bipartite graph $G_{i,j} := \prod_{\ell=1}^{j} X_{i,\ell} \times \prod_{\ell=j+1}^{k} X_{i,\ell}$ corresponding to nonzero probability events is connected. Consider any atom $x \in \prod_{\ell=1}^{j} X_{i,\ell}$ and $y \in \prod_{\ell=1}^{j} X_{i,\ell}$. Since x and y are marginals with nonzero probability, there exist atoms $x', y' \in \mu_i$ such that x is a prefix of x' and y is a suffix of y'. Since μ_i is connected, there exists $z_0, \ldots, z_n$ such that $z_0 = x', z_n = y'$ and z_i and z_{i-1} differ in exactly one coordinate for all $i \in [n]$. In particular, this implies that each z_i corresponds to an edge of $G_{i,j}$ and consecutive edges share a vertex. Thus, x and y are connected; therefore $G_{i,j}$ is connected. Hence, the μ_i's satisfy property (4) of Definition 5.4.

By essentially the same argument, we can see that the μ_i's satisfy property (3) of Definition 3.1. □

We can apply this lemma to obtain results about the CSPs corresponding to strong and rainbow hypergraph coloring.

Lemma 5.4. *For all $k \geq 2$ and $q \geq k + \sqrt{k} + \frac{1}{2}$, $S_{k,q}$ is unpredictably V label cover-compatible.*

Proof. Since $S_{k,q}$ is a symmetric predicate, it suffices without loss of generality to construct the distribution μ_1. The distribution μ_1 corresponds to the following algorithm

1. Pick $m \in \{2k - q - 1, \ldots, k - 1\}$ according to a distribution Ω to be specified.
2. Pick uniformly at random a partial matching $(a_1, b_1), \ldots, (a_m, b_m) \in \{2, \ldots, k\}^2$. Such that $a_i \neq a_j$ and $b_i \neq b_j$ for all distinct $i, j \in [m]$.
3. Define $S' \subset S_{k,q}$ to be

$$S' := \{((x_1, \ldots, x_k), (y_1, \ldots, y_k)) \in S_{k,q} \mid x_1 = y_1, \forall a, b \in \{2, \ldots, k\},$$
$$y_a = x_b \text{ iff } \exists i \in [m], (a, b) = (a_i, b_i)\}.$$

   Pick $((x_1, \ldots, x_k), (y_1, \ldots, y_k)) \sim S'$ uniformly at random.

Our sample from μ_1 is then $((x_1, \ldots, x_k), (y_1, \ldots, y_k))$. In order for this to be sensible, we need to verify the following claim.

Claim 5.5. *For any choice of $m \in \{2k - q - 1, \ldots, k - 1\}$ and $(a_1, b_1), \ldots, (a_m, b_m)$, we have that S' is nonempty.*

Proof. By symmetry, we may assume without loss of generality that $a_i = b_i = i + 1$ for all $i \in [m]$. Let $x_i = i$ for all $i \in [k]$. Let $y_i = i$ for all $i \in [m + 1]$. For $i \in \{m + 1, \ldots, k\}$, let $y_i = k + i - (m + 1) \leq k + k - (2k - q) = q$. Thus, $((x_1, \ldots, x_n), (y_1, \ldots, y_n)) \in S'$, as desired. □

Now, we pick our distribution Ω to satisfy the property guaranteed by the following claim

Claim 5.6. *There is a distribution Ω supported on $\{2k - q + 1, \ldots, k - 1\}$ such that each element of the set has nonzero probability and*

$$\mathbb{E}[\Omega] = \frac{(k-1)^2}{q}.$$

Proof. Since $q > k + \sqrt{k} - \frac{1}{2}$, we have that

$$2k - q + 1 < \frac{(k-1)^2}{q} < k - 1.$$

By an application of the intermediate value theorem, there must be some distribution on Ω with the desired mean which gives every $m \in \{2k - q + 1, \dots, k - 1\}$ nonzero probability. □

Since the algorithm is symmetric with respect to the colors, we have that x_1 (and thus also y_1) is chosen uniformly at random. Therefore, μ_1 has property (1) of Definition 3.1. Fix $i, j \in \{2, \dots, k\}$ (not necessarily distinct). Since our algorithm is symmetric with respect to these pairs (i, j), Claim 5.6 guarantees that $x_j = y_j$ with probability $1/q$. Thus, $x_j \neq y_j$ with probability $(q-1)/q$. These probabilities are consistent with the uniform distribution on $[q]^2$. By the symmetry of the algorithm, we have that once we decide whether $x_j = y_j$ or $x_j \neq y_j$, the coloring is chosen uniformly from the valid options. Thus, (x_j, y_j) is a uniform distribution on $[q]^2$, so μ_1 has property (2) of Definition 3.1 and property (5) of Definition 5.4.

The last thing to verify is that μ_1 is a connected distribution. Let $S''_{k,q} := \{(x, y) \in S^2_{k,q} \mid x_1 = y_1\}$ note that each element $(x, y) \in S''_{k,q}$ has nonzero probability in μ_1, since there is a nonzero probability that m is chosen and $\{(a_i, b_i) \mid i \in [m]\}$ are drawn in order to equal to $\{(i, j) \in \{2, \dots, k\}^2 \mid x_i = y_j\}$. Then, since $(x, y) \in S'$, there is a nonzero probability (x, y) is drawn. Thus, to show that μ_1 is connected, it suffices to show that each $(x, y) \in S''_{k,q}$ can reach $((1, 2, \dots, k), (1, 2, \dots, k)) \in S'_{k,q}$ by changing pairs (x_i, y_i) while staying in $S''_{k,q}$. This can be done by Algorithm 1.

for $c \in [k]$ **do**
    for $j \in [k] \setminus \{1\}$ ***with*** $x_j = c$ **do**
        Set x_j to some color in $[q] \setminus \{x_1, \dots, x_k\}$
    end
    for $j \in [k] \setminus \{1\}$ ***with*** $y_j = c$ **do**
        Set y_j to some color in $[q] \setminus \{y_1, \dots, y_k\}$
    end
    Set $x_c = y_c = c$
end

Algorithm 1: Algorithm demonstrating connectivity of $S''_{k,q}$.

In the two internal for loops, the modification is always legal, as we purposely select a color not among those used by the other x_k's. The last line is also legal for $c = 1$ since every other variable has value other than 1. The last line is also legal for $c > 1$ since $x_j, y_j \neq c$ for all $j \in \{2, \dots, n\}$ and $x_1 = y_1 = 1 \neq c$. Thus, we have that μ_1 is connected.

Thus, by Lemma 5.3, we have that μ_1 is unpredictably V label cover-compatible. □

Lemma 5.7. *For all $k \geq 3$, $R_{k,k-1}$ is unpredictably V label cover-compatible*

Proof. Again, it suffices to construct μ_1 only. Consider the following distribution. Note that the support of this distribution is a strict subset of $R_{k,q}$, where $q = k - 1$.

1. Let $(x_2, \dots, x_k)$ and $(y_2, \dots, y_k)$ be independently chosen uniformly random permutations of $(1, \dots, q)$.

2. Pick $b \in \{0, 1\}$ and $\ell \in \{2, \ldots, k\}$ uniformly at random.

3. If $b = 0$, set $x_1 = y_1 = x_\ell$ and then recolor x_ℓ uniformly at random (possibly the same color). Otherwise, if $b = 1$, set $x_1 = y_1 = y_\ell$ and recolor y_ℓ uniformly at random.

Like usual, $((x_1, \ldots, x_k), (y_1, \ldots, y_k))$ is our sample from μ_1. It is straightforward to verify that this distribution μ_1 satisfies properties (1) and (2) of Definition 3.1 and property (5) of Definition 4.1. To verify the other properties, by Lemma 5.3, if suffices to show that the support of μ_1 is connected. We do this by demonstrating that everything connects to $\{(1, 1, 2, \ldots, q), (1, 1, 2, \ldots, q)\}$.

First, note that for any $(x, y) \in \mu_1$, we have that (x, y) is connected to $(x', y') \in \mu_1$ such that $(x'_2, \ldots, x'_k)$ and $(y'_2, \ldots, y'_k)$ are permutations of $(1, \ldots, q)$, because by Step (3) we can change the color of either x_ℓ or y_ℓ to make the permutations.

Second, observe that if $(x, y) \in \mu_1$ has the property that $(x_2, \ldots, x_k)$ and $(y_2, \ldots, y_k)$ are permutations of $(1, \ldots, q)$, then the modification (x', y') with $x'_1 = y'_1 = 1$, but otherwise equal to (x, y), is also in the support of μ_1.

Next, we show that if $(x, y) \in \mu_1$ has $x_1 = y_1 = 1$ and $(x_2, \ldots, x_k)$ and $(y_2, \ldots, y_k)$ are permutations of $(1, \ldots, q)$, then for any distinct $i, j \in \{2, \ldots, k\}$, $(x', y') \in \mu_1$ with $x'_j = x_i$ and $x'_i = j$, but otherwise equal to (x, y), is connected to (x, y). We do this as follows.

1. Set $x_1 = y_1 = x_i$.

2. Set $x_i = x_j$.

3. Set $x_j = x_1$.

4. Set $x_1 = 1$.

It is clear a similar result holds for transposing the elements of y instead of the elements of x.

Now, by applying a standard sorting algorithm, we can see that all $(x, y) \in \mu_1$ are connected to $((1, 1, 2, \ldots, q), (1, 1, 2, \ldots, q))$, as desired. Thus, $R_{k,k-1}$ is unpredictably V label cover–compatible. □

5.2 Hardness results

Now that we know are predicates are unpredictably V label cover-compatible, we may proceed with establishing Theorems 5.1 and 5.2.

If Φ is a P-CSP, in which $P \subseteq [q]^k$, define the *underlying k-uniform hypergraph of* Φ to be the k-uniform hypergraph who vertices are the variables of Φ and those hyperedges are the clauses of Φ.

Theorem 5.8. *Let $P \subseteq [q]^k$ $(q, k \geq 2)$, be a predicate which supports a unpredictably V label cover-compatible distribution. Then, assuming the induced version of the V label cover-conjecture (Conjecture 3.1), for all $\epsilon > 0$, it is NP-hard to distinguish the following for a P-CSP Φ.*

- *YES: Φ is perfectly satisfiable.*
- *NO: The underlying k-uniform hypergraph of Φ does not have an ϵ-density weak independent set.*

Proof. The proof mirrors the structure of Theorem 4.1 and also incorporates some ideas from [DMR09]. First, we describe the reduction from a V label cover instance to an instance of P-CSP, and note that such a reduction preserves perfect completeness. Second, we analyze the soundness of our reduction using Theorem 2.6 to show that if the underlying hypergraph of the P-CSP has a large weak-independent set, then our original V label cover instance also admits an approximation.

Reduction. The reduction is exactly that specified in Section 3.3. This time, we make no modifications for folding. In particular by Claim 3.2, the reduction has perfect completeness.

Soundness. Assume the the underlying hypergraph $H_\Phi = (V_\Phi, E_\Phi)$ has a large weak independent set $I \subset V_\Phi$ with $|I| \geq \epsilon |V_\Phi|$. We view I as a collection of functions $\mathcal{F} = \{f_u : [q]^{(2k-1)L} \to \{0,1\} : u \in U\}$, where $f_u(s) = 1$ if and only if $x_s^{(u)} \in I$. From this, it is clear that

$$\frac{1}{|U|} \sum_{u \in U} \mathbb{E}[f_u] \geq \epsilon,$$

where the expectation is taken over the uniform distribution on $[q]^{(2k-1)L}$. We also define the influences and the low-degree influences of the f_u's with respect to the uniform distribution of $[q]^{(2k-1)L}$. Thus, there exists a subset $U' \subseteq U$ of size $|U'| > (\epsilon/2)|U|$ for which $\mathbb{E}[f_u] > \epsilon/2$ for all $u \in U'$. As otherwise,

$$\frac{1}{|U|} \sum_{u \in U} \mathbb{E}[f_u] \geq \epsilon < \frac{\epsilon}{2}(1) + \left(1 - \frac{\epsilon}{2}\right) < \epsilon.$$

For each $e = (u_1, \ldots, u_k) \in E \cap (U')^k$, since I is a weak independent set of H_Φ.

$$\begin{aligned} 0 &= \mathop{\mathbb{E}}_{(s_1,\ldots,s_k)\sim\nu^{(e)}}[x_{s_1}^{(u_1)} \in I \wedge \cdots \wedge x_{s_k}^{(u_k)} \in I] \\ &= \mathop{\mathbb{E}}_{(s_1,\ldots,s_k)\sim\nu^{(e)}}[f_{u_1}(s_1) \cdots f_{u_k}(s_k)]. \end{aligned}$$

For all $i \in [k]$, let $\Pi_i^{(e)} = \prod_{\ell=1}^{kL} X_{i,\ell}^{(e)}$. Since each $\Pi_i^{(e)}$ is uniform, $\mathop{\mathbb{E}}_{s_i \sim \Pi_i^{(e)}}[f_{u_i}] > (\epsilon/2)$.

Because P is unpredictably V label cover–compatible, $\nu^{(e)} = \Omega_1^{(e)} \times \cdots \times \Omega_{kL}^{(e)}$ meets the requirements of Theorem 2.7. Thus, there exists $\epsilon', \tau, d > 0$, which are functions of only $\epsilon/2$ and parameters of $|P|$, such that there are $i_1, i_2 \in [k]$ and $t_1, t_2 \in [(2k-1)L]$ such that (i_1, t_1) and (i_2, t_2) are in the same branch and

$$\mathrm{Inf}^{\leq d}_{Y_{t_a}^{(e,i_a)}} f_{u_i} = \mathrm{Inf}^{\leq d}_{t_a} f_{u_i} > \tau.$$

For each $u \in U'$. Let $S_u \subseteq [(2k-1)L]$ be the set of labels j for which $\mathrm{Inf}^{\leq d}_j f_u > \tau$. Since $\mathrm{Var}\, f_u \leq \max(f_u)^2 = 1$, we have by Lemma 2.3 that $|S_u| \leq d/\tau$, which is independent of L. Construct a random partial labeling $\sigma : U' \to [(2k-1)L]$ by sampling each $\sigma(u)$ from S_u independently and uniformly at random (if S_u is empty, let $\sigma(u) = 1$). For each $e \in E \cap (U')^k$, we established that there exists $i, i' \in [k]$ and $\ell \in S_{u_i}$ and $\ell' \in S_{u_{i'}}$ such that setting $\sigma(u_i) = \ell$ and $\sigma(u_{i'}) = \ell'$ weakly satisfies e. Thus, inside U' expectation at least

$$\frac{1}{(\max |S_u|)^2} = \frac{\tau^2}{d^2} > 0$$

of the edges are weakly satisfied. Note that this expression is independent of L and the size of Ψ, as desired. □

Note that Theorem 5.1 follows as a corollary of Theorem 5.8 combined with Lemma 5.4.

Proof of Theorem 5.2. Theorem 5.8 and Lemma 5.7 imply the case $q = k - 1$. For $q < k - 1$, one can see that a $(k, k-1)$-rainbow colorable hypergraph is also a (k, q)-rainbow colorable hypergraph since we can 'merge' colors together while preserving the rainbow property. Therefore, since the V label cover–conjecture implies for $\epsilon > 0$, it is NP-hard to distinguish $(k, k-1)$-rainbow colorable hypergraphs from graphs without an ϵ-density independent set, then for any $q \leq k - 1$ it must be NP-hard to distinguish (k, q)-rainbow colorable hypergraphs from graphs without an ϵ-density independent set. □

Theorems 5.1 and 5.2 together imply Theorem 1.2.

Acknowledgments

We would like to thank Elchanan Mossel for a useful discussion on a generalization of Theorem 2.6. We would also like to thank anonymous reviewers for helpful comments.

A (L, s)-nearly 1–to–1 label cover

Consider the following variant of the classic Label Cover problem.

Definition A.1. Let L be a positive integer and $s \in \{0, \ldots, L^2\}$. An instance of (L, s)-*nearly 1-to-1 Label Cover* consists of $\Psi = (V, E, \{S_e\}_{e\in E}, \{\pi_{e,u}\}_{e\in E, u\in e})$, where (V, E) is a regular graph,[10], the $S_e \subseteq [L] \times [L]$ have size s,[11] and the maps $\pi_{e,u} : [L] \to [L]$ are permutations. A labeling is a function $\sigma : V \to [L]$. An edge $e \in E$ is satisfied if $(\pi_{e,u}(\sigma(u)), \pi_{e,v}(\sigma(v))) \in \{(\ell, \ell) : \ell \in [L]\} \cup S_e$.

Assume $s \geq 1$ (as the case $s = 0$ is unique games with perfect completeness). We show that when s is a constant relative to L, the (L, s)-nearly 1-to-1 Label Cover problem is efficiently approximable.

Theorem A.1. *There exists a function $\eta : \mathbb{N} \to (0, 1]$ (presumably decreasing) such that there is a randomized polynomial time algorithm which with high probability distinguishes the following two types of instances $\Psi = (V, E, \{\pi_{e,u}\}_{e\in E, u\in E})$ of (L, s)-nearly 1-to-1 Label Cover.*

- *Accept: Ψ is perfectly satisfiable.*
- *Reject: every labeling of Ψ satisfies strictly less than $\eta(s)$ fraction of the edges.*

In fact, one may take $\eta(s) = \frac{1}{1024s^2}$.

For each $e \in E$, let $T_e = \{x : (x, y) \in S_e\} \cup \{y : (x, y) \in S_e\}$. Note that $|T_e| \leq 2s$.

Assume that a perfect labeling exists for Ψ and let $\Sigma : V \to [L]$ be such a labeling. We show that we can efficiently construct a labeling $\sigma : V \to [L]$ which satisfies at least $\eta(s)$ fraction of the edges. Such an algorithm will suffice to distinguish the two cases specified in the theorem statement.

[10] We assume that (V, E) is a regular graph for simplicity of presentation. The authors believe the same result should hold for general graphs.

[11] If S_e is not symmetric, then the edge e is technically directed, but it is fine to assume that (V, E) is undirected for most of our analysis.

For each $e \in E$, say that $e = (u, v)$ is *type-1-satisfied* by Σ if $\pi_{e,u}(\Sigma(u)) \notin T_e$ and $\pi_{e,v}(\Sigma(v)) \notin T_e$. Otherwise, say that e is *type-2-satisfied* by Σ. Let $E_1 \subseteq E$ be the type-1-satisfied edges, and let E_2 be the type-2 satisfied. Let D be the degree of each vertex of (V, E). Let $d_i(v)$ be the number of vertex incident with vertex v which are type-i satisfied by σ.

First, we use a standard DFS algorithm to construct partial, but perfect labelings of Ψ.

Lemma A.2. *Given $v_0 \in V$ and $\ell \in L$ there is a polynomial time algorithm which outputs a subset $W \subseteq V$ and a partial labeling $\sigma : W \to [L]$ with the following properties.*

- *$v_0 \in W$ and $\sigma(v_0) = \ell$.*
- *Every $e \in E \cap W \times W$ is satisfied by σ.*
- *For every $\sigma' : V \to [L]$ which extends σ (i.e., $\sigma'(v) = \sigma(v)$ for all $v \in W$) which perfectly satisfied Ψ, every edge in the cut $E \cap W \times (V \setminus W)$ must be type-2-satisfied.*

If there is no satisfying assignment $\sigma : V \to [L]$ to Ψ with $f(v) = \ell$, then the algorithm returns $\perp$.

Informally, the last condition means that the partial labeling cannot be extended any further by type-1 satisfying edges.

Proof. Consider the DFS/BFS-like Algorithm 2.

```
Function Partial-Type-1-Labeling(Ψ, W, σ) do
    Data: (L, S)-nearly 1-to-1 Label Cover instance Ψ = (V, E, {S_e}_{e∈E}, {π_{e,u}}_{e∈E,u∈e}), W ⊆ V,
          σ : W → [L]
    Result: Either ⊥ or a pair (W′, σ′) where W′ ⊆ V and σ′ : V → [L].
    for v ∈ W do
        for e ∈ E where v ∈ e do
            Set u to other vertex of e
            if u ∈ W then
                if σ does not satisfy e then return ⊥
            end
            else if π_{e,v}(σ(v)) ∉ T_e then
                Set W′ = W ∪ {u}
                Set σ′ ↾ W = σ
                Set σ′(u) = (π_{e,u}^{-1} ∘ π_{e,v})(σ(v))
                return Partial-Type-1-Labeling(Ψ, W′, σ′)
            end
        end
    end
    return (W, σ)
end
```

Algorithm 2: Finding a partial solution using type-1-satisfied edges.

We claim that calling `Partial-Type-1-Labeling(`Ψ, $\{v_0\}$, $v_0 \mapsto \ell$`)`[12] is the correct procedure. To prove efficiency, it is easy to see that during each recursive call, W will grow by at least one

[12] $v_0 \mapsto \ell$ is shorthand for the function $\sigma : \{v_0\} \to [L]$ such that $\sigma(v_0) = \ell$.

element or the procedure will terminate. Hence there can be at most $|V|$ recursive calls (including the initial call). Furthermore, within one recursive call only a polynomial amount of work is done. Thus, the procedure runs in polynomial time.

To prove correctness, note that the final recursive call will verify that every edge inside the vertices of W is correctly labeled and every edge between W and $V \setminus W$ must be type-2 satisfied. Thus, if the algorithm outputs (W, σ), we know that W and σ will have the required properties. Furthermore, observe that the algorithm adds a new vertex to W only when the label of that vertex is forced. Thus, any contradiction found is proof that there is no fully satisfiable way to extend the initial choice that $\sigma(v_0) = \ell$.

Thus, the algorithm is correct and efficient. □

Note that the above algorithm will do quite well when Σ type-1-satisfies most of the edges. The following algorithm deals with the case in which most of the edges are type-2-satisfied. Let $\delta = |E_2|/|E|$.

Lemma A.3. *Assume Ψ is satisfiable and $\delta \geq 1/2$. Then there is a randomized polynomial-time algorithm which finds a labeling $\sigma : V \to [L]$ which satisfies at least $f(s) = \frac{1}{1024s^2}$ of the constraints of Ψ with probability $1 - \frac{1}{2^{poly(|V|)}}$.*

Remark. We set $\eta(s) = f(s)$.

Function `Approx-Type-2-Labeling(Ψ)` **do**
  Data: (L, S)-nearly 1-to-1 Label Cover instance $\Psi = (V, E, \{S_e\}_{e\in E}, \{\pi_{e,u}\}_{e\in E, u\in e})$
  Result: An approximately satisfying labeling $\sigma : V \to [L]$.
  for $v \in V$ **do**
    Pick $e \in E$ uniformly at random such that $v \in e$.
    Pick $\ell \in T_e$ uniformly at random.
    Set $\sigma(v) = \pi_{e,v}^{-1}(\ell)$
  end
  return σ
end

Algorithm 3: Finding a good approximate solution when there are many type-2 edges.

Proof. Consider Algorithm 3. Clearly the algorithm runs in polynomial time. It suffices to show that the above algorithm succeeds in finding an $\eta(s)$ approximation with constant probability, as one may repeat the subroutine polynomially many times and take the best solution. The first step of our analysis is the following simple claim.

Claim A.4. *For each $v \in V$, with probability at least $\frac{d_2(v)}{2sD}$, $\sigma(v) = \Sigma(v)$.*

Proof. With probability $\frac{d_2(v)}{D}$ we pick an edge e which is type-2 satisfied by σ. With probability $\frac{1}{|T_e|} \geq \frac{1}{2s}$ we then subsequently pick $\sigma(v)$ since $\pi_{e,v}(\sigma(v)) \in T_e$. □

Define a vertex $v \in V$ to be *good* if $d_2(v) \geq |D|/4$. Let $V' \subseteq V$ be the set of good vertices. By Markov's inequality and the fact that V is regular, $|V'| \geq |V|/4$. Define an edge $(u, v) \in E$ to be good if $u, v \in V'$. Let $E' \subseteq E$ be the set of good edges.

Claim A.5. *At least $1/8$ fraction of the edges are good.*

Proof. Let $\psi = |E'|/|E_2|$. Pick a uniformly random edge $e \in E_2$ and pick a uniformly random vertex u of e. The probability that u is not good is at least $(1-\psi)/2$. Note that the probability that any particular u is picked is $\frac{d_2(u)}{2|E_2|}$. Thus,

$$\begin{aligned}
\frac{1-\psi}{2} &\le \sum_{v\in V\setminus V'} \frac{d_2(u)}{2|E_2|} \\
&\le |V \setminus V'| \frac{D}{8|E_2|} \\
&\le \frac{3|V|D}{32|E_2|} \\
&= \frac{3|E|}{16|E_2|} \\
&\le \frac{3}{8}.
\end{aligned}$$

Thus, $\psi \ge \frac{1}{4}$, so $|E'|/|E| \ge \psi\delta \ge \frac{1}{8}$. ☐

The expected fraction of edges satisfied is then

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{|E|}\sum_{e\in E} \mathbf{1}[\sigma \text{ satisfies } e]\right] &\ge \frac{1}{|E|}\sum_{(u,v)\in E} \mathbb{E}[\mathbf{1}[(\sigma(u),\sigma(v)) = (\Sigma(u),\Sigma(v))]] \\
&\ge \frac{1}{|E|}\sum_{(u,v)\in E} \frac{d_2(u)d_2(v)}{(2s)^2D^2} \text{ (Claim A.4 and independence).} \\
&\ge \frac{1}{(2s)^2|E|}\sum_{(u,v)\in E'} \frac{d_2(u)d_2(v)}{D^2}. \\
&\ge \frac{|E'|}{64s^2|E|} \\
&\ge \frac{1}{512s^2} \text{ (Claim A.5).}
\end{aligned}$$

Thus, by Markov's inequality, the above algorithm will find a solution satisfying at least $\frac{1}{1024s^2}$ fraction of the edges with probability at least $\frac{1}{1024s^2}$. ☐

Consider the case $\delta \le 1/2$, Thus, most of the edges of E are type-1 satisfied by σ. Assume that (V, E_1) has k connected components $V_1 \cup \cdots \cup V_k = V$. We would like to show that one of these sets $V_1, \ldots, V_k$ has small *edge expansion*. First, recall the definition of edge expansion.

Definition A.2. The *edge expansion* of a subset $V' \subseteq V$ of an undirected D-regular graph (V, E) is

$$\Phi(V') = \frac{|E \cap (V' \times (V \setminus V'))|}{D|V'|}.$$

Intuitively, if we can perfectly label the induced edges of a connected component V_i with poor edge expansion, we have made good progress toward a labeling satisfying a constant fraction of the edges, as we can recursively apply our algorithm to find an approximate labeling of $V \setminus V_i$ and union it with our labeling of V_i. The following lemma shows that such a V_i always exists.

Lemma A.6. *Let (V, E) be an undirected D-regular graph, and let $V_1, \ldots, V_k$ be a partition of the vertices. Assume that at most δ fraction of the edges of E are between different V_i. Then there exists an $i \in [k]$ such that $\Phi(V_i) \leq \delta$.*

Proof. Let E' be the set of edges between the V_i. Note that

$$\begin{aligned} 2|E'| &= \sum_{i \in [k]} |E \cap (V_i \times (V \setminus V_i))| \text{ (each edge of } E' \text{ is between two of the } V_i\text{'s)} \\ &= \sum_{i \in [k]} D|V_i|\Phi(V_i) \\ &\geq D|V| \min_i \Phi(V_i). \end{aligned}$$

Since $2|E'| = \delta(2|E|) = D|V|\delta$, we have that $\delta \geq \min_i \Phi(V_i)$. □

With this lemma proven, we may now state the final algorithm (4).

Function `Approximate-Labeling(Ψ)` **do**
  Data: (L, S)-nearly 1-to-1 Label Cover instance $\Psi = (V, E, \{\pi_{e,u}\}_{e \in E, u \in e})$
  Result: Either $\perp$ or an approximately-satisfying labeling $\sigma : V \to [L]$.
  Set σ_1 = `Approx-Type-2-Labeling` (Ψ)
  if σ_1 *satisfies* $\eta(s)$ *fraction of the edges of* Ψ **then**
    return σ_1
  end
  for $v \in V$, $\ell \in [L]$ **do**
    Set $\sigma_2(v) = \ell$.
    Set τ = `Partial-Type-1-Labeling`(Ψ, v, σ_2)
    if $\tau = (W, \sigma_2')$ ***and*** $\Phi(W) \leq 1/2$ **then**
      Set $V' = V \setminus W$
      Set $E' = E \cap (V')^2$
      Set $\Psi' = (V', E', \{\pi_{e,u}\}_{e \in E', u \in e})$
      Set σ_3 = `Approximate-Labeling`(Ψ')
      if $\sigma_3 \neq \perp$ **then return** $\sigma_2' \cup \sigma_3$
    end
  end
  return $\perp$
end

Algorithm 4: The full algorithm.

Proof of Theorem A.1. We prove the algorithm works by strong induction on $|V|$.

Assume Ψ is perfectly satisfiable. If $\delta = |E_2'|/|E| \geq 1/2$, then Lemma A.3 guarantees that we will find an $\eta(s)$ approximation with high probability. Otherwise, if $\delta \leq 1/2$, we know by Lemma A.6 there exists $W \subseteq V$ and a perfect partial labeling $\sigma_2' : W \to [L]$ such that W is connected by edges type-1-satisfied by σ and $\Phi(W) \leq 1/2$. By Lemma A.2, the above for loop will succeed in finding some (W, σ_2') with these properties in polynomial time. By the strong induction hypothesis, we can with high probability find a $\eta(s)$-approximate labeling σ_3 to the instance Ψ' induced by $V \setminus W$. Thus, the labeling $\sigma_2' \cup \sigma_3$ satisfies at least $1/2$ of the edges incident with at least one vertex of W (since $\Phi(W) \leq 1/2$ and the edges inside of W are perfectly satisfied) and at least $\eta(s)$ of the edges not incident with W. Thus, we have efficiently found a $\min(1/2, \eta(s)) = \eta(s)$ approximation for Ψ. If the algorithm does not succeed, then Ψ is not perfectly satisfiable. $\square$

B Proof of Theorem 2.7

Recall the statement of Theorem 2.7.

Theorem B.1 (Theorem 2.7)**.** *Fix $k \geq 2$. For $1 \leq \ell \leq n$, let $\Omega_\ell = X_\ell^{(1)} \times \cdots \times X_\ell^{(k)}$ be a finite probability space with distributions μ_ℓ such that the μ_ℓ's are independent. Also, assume that for each $\ell \in [n]$ and $i \in [k]$, $X_\ell^{(i)} = \prod_{s=1}^{s_\ell^{(i)}} Y_{\ell,s}^{(i)}$, where the product is of otherwise independent distributions and $s_\ell^{(i)} \leq 2$ for all $i \in [k]$ and $\ell \in [n]$. Assume we also have the following key property*

- *If for distinct $i_1, i_2 \in [k]$ we have that $s_\ell^{(i_1)} = s_\ell^{(i_2)} = 2$, then $Y_{\ell,1}^{(i_1)}$ is independent of $Y_{\ell,2}^{(i_2)}$ (and $Y_{\ell,2}^{(i_2)}$ is independent of $Y_{\ell,1}^{(i_1)}$ by symmetry).*

For convenience of notation, if $s_\ell^{(i)} = 1$, let $Y_{\ell,2}^{(i)} := Y_{\ell,1}^{(i)}$. Let δ be the minimum positive probability among all the μ_ℓ's, $\ell \in [n]$. Let

$$\rho = \max\left\{\max_{1\leq \ell \leq n} \rho(X_\ell^{(1)}, \ldots, X_\ell^{(k)}),\ \max_{\substack{1\leq \ell\leq n\\ 1\leq j<n}} \rho\left(\prod_{\ell=1}^{j} X_\ell,\ \prod_{\ell=j+1}^{k} X_\ell\right)\right\}.$$

and assume that $\rho < 1$. For every $\epsilon > 0$, there exists $\epsilon'(\delta,\epsilon,\rho), \tau(\delta,\epsilon,\rho), d(\delta,\epsilon,\rho) > 0$ such that for any functions $f_1, \ldots, f_k$ where $f_i : X_1^{(i)} \times \cdots \times X_n^{(i)} \to [0,1]$ and $\mathbb{E}[f_i] \geq \epsilon$ if

$$\forall \ell \in [n], \forall s \in \{1,2\}, |\{i \mid \mathrm{Inf}^{\leq d}_{Y_{\ell,s}^{(i)}} f_i \geq \tau\}| \leq 1$$

then

$$\mathbb{E}\left[\prod_{i=1}^{k} f_i\right] \geq \epsilon'.$$

Proof. The proof of this theorem follows a similar structure to the proof of Theorem 3.11 of [DMR09]. Let ϵ_0', τ_0 be the values guaranteed by Theorem 2.6 for parameters ρ, ϵ, δ. For each $i \in [k]$, we define

$$\Sigma^{(i)} = [|X_1^{(i)}|] \times \cdots \times [|X_n^{(i)}|]$$
$$(\Sigma^{(i)})' = [|Y_{1,1}^{(i)}|] \times [|Y_{1,2}^{(i)}|] \times \cdots \times [|Y_{n,2}^{(i)}|]$$
$$\forall \ell \in [n], \alpha_1^{(i,\ell)}, \ldots, \alpha_{|X_\ell^{(i)}|}^{(i,\ell)} : X_\ell^{(i)} \to \mathbb{R} \text{ orthonormal basis with } \alpha_1^{(i,\ell)} \equiv 1$$
$$\forall (\ell, s) \in [n] \times [2], \beta_1^{(i,\ell,s)}, \ldots, \beta_{|Y_{\ell,s}^{(i)}|}^{(i,\ell,s)} : Y_{\ell,s}^{(i)} \to \mathbb{R} \text{ orthonormal basis with } \beta_1^{(i,\ell,s)} \equiv 1$$

We also require that α's and β's are consistent in the following sense. Since $X_\ell^{(i)} = Y_{\ell,1}^{(i)} \times \cdots \times Y_{\ell,s_\ell^{(i)}}^{(i)}$, we have that

$$\beta_{j_1}^{(i,\ell,1)} \cdots \beta_{j_k}^{(i,\ell,s_\ell^{(i)})}$$

is an orthonormal basis of the functions from $X_\ell^{(i)}$ to $\mathbb{R}$, where $(j_1, \ldots, j_{s_\ell^{(i)}}) \in [|Y_{\ell,1}^{(i)}|] \times \cdots \times [|Y_{\ell, s_\ell^{(i)}}^{(i)}|]$. Since $\beta_1^{(i,\ell,1)} \cdots \beta_1^{(i,\ell,s_\ell^{(i)})} = 1$, we may assume that the $\alpha_j^{(i,\ell)}$'s are some enumeration of this basis.

We define the Fourier coefficients of the f_i's (see Definition 2.4 for notation) to be

$$\begin{aligned} f_i &:= \sum_{\sigma \in \Sigma^{(i)}} c_\sigma^{(i)} \prod_{\ell=1}^n \alpha_{\sigma_\ell}^{(i,\ell)} && (X_\ell^{(i)} \text{ marginals}) \\ &= \sum_{\sigma' \in (\Sigma^{(i)})'} c_{\sigma'}^{(i)} \prod_{\ell=1}^n \prod_{s=1}^{s_\ell^{(i)}} \beta_{\sigma'_{\ell,s}}^{(i,\ell,s)} && (Y_{\ell,s}^{(i)} \text{ marginals}), \end{aligned}$$

where $c_{\sigma'}^{(i)} := c_\sigma^{(i)}$ if $\prod_{\ell=1}^n \alpha_{\sigma_\ell}^{(i,\ell)} = \prod_{\ell=1}^n \prod_{s=1}^{s_\ell^{(i)}} \beta_{\sigma'_{\ell,s}}^{(i,\ell,s)}$.

Denote $|\sigma| = \{\ell \in [n] \mid \sigma_\ell \neq 1\}$ for $\sigma \in \Sigma^{(i)}$. For $\sigma' \in (\Sigma^{(i)})'$, we denote $|\sigma'| = \{\ell \in [n] \mid s \in [s_\ell^i], \sigma_{i,s} \neq 1\}$. A key property is that if $\prod_{\ell=1}^n \alpha_{\sigma_\ell}^{(i,\ell)} = \prod_{\ell=1}^n \prod_{s=1}^{s_\ell^{(i)}} \beta_{\sigma'_{\ell,s}}^{(i,\ell,s)}$, then $|\sigma| = |\sigma'|$.

To not be concerned with low-degree influences, we first replace each f_i, with a noised version $T_\eta^X f_i$,[13] which is defined in terms of Fourier coefficients to be

$$T_\eta^X f_i := \sum_{\sigma \in \Sigma^{(i)}} \eta^{|\sigma|} c_\sigma^{(i)} \prod_{\ell=1}^n \alpha_{\sigma_\ell}^{(i,\ell)}.$$

Note that this noise operator is applied to the $X_\ell^{(i)}$ marginals. Rewriting this in terms of the $Y_{\ell,s}^{(i)}$ basis,

$$T_\eta^X f_i = \sum_{\sigma' \in (\Sigma^{(i)})'} \eta^{|\sigma'|} c_{\sigma'}^{(i)} \prod_{\ell=1}^n \prod_{s=1}^{s_\ell^{(i)}} \beta_{\sigma'_{\ell,s}}^{(i,\ell,s)}.$$

Since the range of each f_i is a subset of $[0, 1]$, it well-known that $T_\eta^X f_i$'s range is also a subset of $[0, 1]$ (e.g., Definition 8.28 of [O'D14]).

Let $\epsilon_1 = \epsilon_0'/(4k) > 0$. Since ρ, our correlation, is bounded away from 1, by Lemma 6.2 of [Mos10], there exists $\eta < 1$ such that

$$\begin{aligned} \left| \mathbb{E}\left[\prod_{i=1}^k f_i\right] - \mathbb{E}\left[\prod_{i=1}^k (T_\eta^X f_i)\right] \right| &\leq \epsilon_1 \sum_{i=1}^k \sqrt{\mathrm{Var}[f_i]} \sqrt{\mathrm{Var}\left[\prod_{j<i} T_\eta^X f_j \prod_{j>i} f_j\right]} \\ &\leq \epsilon_1 k = \frac{\epsilon_0'}{4}. \end{aligned} \tag{1}$$

The second inequality follows from the fact that range for both f_i and $\left(\prod_{j<i} T_\eta^X f_j \prod_{j>i} f_j\right)$ are inside $[0, 1]$, so their variances are bounded by 1. Let $g_i := T_\eta^X f_i$ for all i. Note that $\mathbb{E}[g_i] = \mathbb{E}[f_i] \geq \epsilon$ and $\mathrm{Var}[g_i] \leq \mathrm{Var}[f_i] \leq 1$. From (1), it suffices to give a lower bound on $\mathbb{E}\left[\prod_{i=1}^k g_i\right]$.

Similar to [DMR09], we $d \in \mathbb{N}$ such that $2^8(d+1)\eta^d < (\epsilon_0')^2 \tau_0$. Also fix

$$\tau = \frac{\tau_0 (\epsilon_0')^2}{2^8 d k^3} < \frac{\tau_0}{4}.$$

Thus, $\eta^d < \tau$. We need the following quantitative bound

[13]We add a superscript X to signify that the noise operator is with respect to the $X_1^{(i)} \cdots \times \cdots X_n^{(i)}$ basis.

Claim B.2. *For all $i \in [k]$ and $(\ell, s) \in [n] \times [2]$, we have that*

$$\operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i \le \operatorname{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i + \tau.$$

Proof. We have that

$$\begin{aligned}
\operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i &= \sum_{\substack{\sigma' \in (\Sigma^{(i)})' \\ \sigma'_{\ell,s} \neq 1}} \eta^{2|\sigma'|} (c^{(i)}_{\sigma'})^2 \\
&\le \sum_{\substack{\sigma' \in (\Sigma^{(i)})' \\ \sigma'_{\ell,s} \neq 1}} \eta^{|\{(\ell',s') \mid \sigma'_{\ell',s'} \neq 1\}|} (c^{(i)}_{\sigma'})^2 \qquad (= \langle f_i, T^Y_\eta f_i \rangle) \\
&\le \operatorname{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i + \eta^d \operatorname{Var}[f_i] \\
&\le \operatorname{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i + \tau,
\end{aligned}$$

as desired. □

Now consider,

$$B := \{\ell \in [n] \mid \exists s \in \{1,2\}, \exists i \in [k], \operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i \ge \tau_0/2\}.$$

(This is analogous to the "B" in the proof of Theorem 3.11 of [DMR09].) Then note that for all $\ell \in B$, there exists $s \in [2]$ and $i \in [k]$ such that

$$\frac{\tau_0}{2} \le \operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i \le \operatorname{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i + \tau,$$

where the inequality follows from Claim B.2. Since $\tau < \tau_0/4$, we have that $\operatorname{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i > \tau_0/4$. Therefore, by Lemma 2.3, we have that each $i \in [k]$ has at most $4d/\tau_0$ marginals $Y^{(i)}_{\ell,s}$ with $\operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i \ge \tau_0/2$. Thus, $|B| \le \frac{4kd}{\tau_0}$.

Now, we show that we can 'smooth out' these high-influence coordinates without substantially changing the product of the g_i's. For each $i \in [k]$, define

$$\begin{aligned}
B^{\mathsf{lo}}_i &:= \{(\ell, s) \in B \times [2] \mid \operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i \le 2\tau\} \\
B^{\mathsf{hi}}_i &:= \{(\ell, s) \in B \times [2] \mid \operatorname{Inf}_{Y^{(i)}_{\ell,s}} g_i > 2\tau\}.
\end{aligned}$$

For each $i \in [k]$, we define $g'_i, g''_i : X^{(i)}_1 \times \cdots \times X^{(i)}_n$ to be

$$\begin{aligned}
g'_i &:= \mathop{\mathbb{E}}_{\prod_{(\ell,s) \in B^{\mathsf{lo}}_i} Y^{(i)}_{\ell,s}} [g_i] \\
g''_i &:= \mathop{\mathbb{E}}_{\prod_{(\ell,s) \in B^{\mathsf{hi}}_i} Y^{(i)}_{\ell,s}} [g'_i].
\end{aligned}$$

(c.f., the "averaging operator" in Section 2.1 of [DMR09]). In other words, in g'_i, we average out all of the low-influence marginals in the blocks $\Omega_\ell = X^{(1)}_\ell \times \cdots \times X^{(k)}_\ell$ which contain a high-influence marginal. In g_i'', we then average out the remaining marginals. For each $i \in [k]$, we average out at most $2|B| \le \frac{4kd}{\tau_0}$ coordinates, so we have that

$$\mathrm{Var}[g_i - g_i'] \le \sum_{(\ell,s)\in B_i^{lo}} \mathrm{Inf}_{Y^{(i)}_{\ell,s}}\, g_i \le 2|B|(2\tau) = \frac{16kd\tau}{\tau_0} \le \frac{(\epsilon_0')^2}{16k^2}.$$

Thus, by Cauchy-Schwartz,

$$\begin{aligned}\left|\mathbb{E}\left[\prod_{i=1}^k g_i - \prod_{i=1}^k g_i'\right]\right| &\le \sum_{i=1}^k \sqrt{\mathrm{Var}[g_i - g_i']}\sqrt{\mathrm{Var}\left[\prod_{j<i} g_j' \prod_{j>i} g_j\right]} \\ &\le \sum_{i=1}^k \frac{\epsilon_0'}{4k} = \frac{\epsilon_0'}{4}. \end{aligned} \tag{2}$$

Now, we claim the (rather remarkable) fact that

$$\mathbb{E}\left[\prod_{i=1}^k g_i'\right] = \mathbb{E}\left[\prod_{i=1}^k g_i''\right]. \tag{3}$$

First recall the assumption that for all $(\ell, s) \in [n] \times [2]$, we have that

$$|\{i \mid \mathrm{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i \ge \tau\}| \le 1. \tag{4}$$

By Claim B.2, if $(\ell, s) \in B_i^{hi}$ for some $i \in [k]$, then $\mathrm{Inf}^{\le d}_{Y^{(i)}_{\ell,s}} f_i \ge \tau$. Thus, each $(\ell, s) \in [n] \times [2]$ is in at most one B_i^{hi}. Consider the set of marginal distributions

$$S := \{Y^{(i)}_{\ell,s} \mid i \in [k], (\ell, s) \in [n] \times [2] \text{ s.t. } (\ell, s) \in B_i^{\mathsf{hi}}\}.$$

We claim that this set of marginal distributions is independent. Since $\Omega_1, \dots, \Omega_n$ are independent, it suffices to check that each subset

$$S \supset S_\ell := \{Y^{(i)}_{\ell,s} \mid i \in [k], s \in [2] \text{ s.t. } (\ell, s) \in B_i^{\mathsf{hi}}\}$$

for $\ell \in [n]$. Note that $|S_\ell| \le 2$ for all ℓ because of (4). If $|S_\ell| = 1$ we are immediately done. If $Y^{(i)}_{\ell,s} \in S_\ell$ in which $s^{(i)}_\ell = 1$, then recall that $Y^{(i)}_{\ell,1} = Y^{(i)}_{\ell,2}$, so $(\ell, 1), (\ell, 2) \in B_i^{hi}$. Thus, there cannot be any more elements of S_ℓ besides $Y^{(i)}_{\ell,1}$. Thus, $|S_\ell| = 1$, so S_ℓ is vacuously a set of independent random variables. In the last case, we have $|S_\ell| = 2$ and $Y^{(i_1)}_{\ell,1}, Y^{(i_2)}_{\ell,2} \in S_\ell$ have the property that $s^{(i_1)}_\ell = s^{(i_2)}_\ell = 2$. Then, by our assumption on the marginal distributions, we have that $Y^{(i_1)}_{\ell,1}$ and $Y^{(i_2)}_{\ell,2}$ are independent. Thus, S_ℓ is independent for all $\ell \in [n]$. Therefore, S is independent.

Let $Z = \prod_{\ell \in [n]\setminus B} \Omega_\ell$. It is easy to see that $S \cup Z$ is independent. Note that $\prod_{i=1}^k g_i'$ is a function of only $Z \times \prod_{Y\in S} Y$ and that $\prod_{i=1}^k g_i''$ is a function of only Z. Thus,

$$\begin{aligned}\mathop{\mathbb{E}}_{Z\times\prod_{Y\in S} Y}\left[\prod_{i=1}^k g_i'\right] &= \mathop{\mathbb{E}}_{Z}\left[\mathop{\mathbb{E}}_{\prod_{Y\in S} Y}\left[\prod_{i=1}^k g_i'\right]\right] \text{ (independence)} \\ &= \mathop{\mathbb{E}}_{Z}\left[\left[\prod_{i=1}^k \mathop{\mathbb{E}}_{\prod_{Y^{(i)}\in S} Y^{(i)}} g_i'\right]\right] \text{ (magic: independence)}\end{aligned}$$

$$= \mathop{\mathbb{E}}_{Z}\left[\prod_{i=1}^{k} g_i''\right].$$

The second equality follows from the fact that each $Y_{\ell,s}^{(i)} \in S$ only affects the value of g_i'. Thus, (3) holds. Now, note that for all $i \in [k]$ and $\ell \in [n]$, we have that

$$\mathrm{Inf}_{X_\ell^{(i)}} g_i'' < \tau_0,$$

as otherwise by Lemma 2.4 (where $d = n$), we would have that there exists $s \in [2]$ for which $\mathrm{Inf}_{Y_{\ell,s}^{(i)}} g_i'' \geq \tau_0/2$, but such coordinates were averaged out from g_i from construction of g_i''. Thus, we may invoke Theorem 2.6 to obtain that

$$\mathbb{E}\left[\prod_{i=1}^{k} g_i''\right] \geq \epsilon_0'. \tag{5}$$

Thus, by combining (1, 2, 3, 5), we have that

$$\mathbb{E}\left[\prod_{i=1}^{k} f_i\right] \geq \frac{\epsilon_0'}{2}.$$

Therefore, we may set $\epsilon' = \epsilon_0'/2 > 0$. □

Remark. Unlike [DMR09], we do not require that the distributions are symmetric.

C Existence of large 3-wise linearly independent subsets

Lemma C.1. *Let $q \geq 2$ be a prime power, and let $\ell \geq 1$ be odd. There exists $S \subset \mathbb{F}_q^\ell$ with $|S| = q^{(\ell-1)/2}$ such that S is 3-wise linearly independent over $\mathbb{F}_q$. That is, each three-element subset of S is linearly independent.*

Remark. Because of the recent breakthrough that subsets of $\mathbb{Z}_q^n$ which do not have an arithmetic progress of length three have size at most q^{cn} for some $c < 1$, [CLP16, EG16], it is impossible to improve that factor of $1/2$ in the exponent to 1 when $q \geq 3$. In particular, Lemma 4.2 can at best be improved to $O_q(k^{2+\gamma})$ for some $\gamma > 0$ (where the O_q notation hides the dependence of q).

Proof. This construction is inspired by user2566092's post on Math StackExchange [uh15]. Let $\ell' = (\ell - 1)/2$. Consider

$$S = \{(1, x_1, x_1^2, x_2, x_2^2, \ldots, x_{\ell'}, x_{\ell'}^2) : x \in \mathbb{F}_q^{\ell'}\}.$$

Consider distinct $u, v, w \in S$ and $a, b, c \in \mathbb{F}_q$ such that $au + bv + cw = 0$. We seek to show that $a = b = c = 0$. First, note that $0 = au_1 + bv_1 + cw_1 = a + b + c$. Since $u \neq v$, by nature of S, there exists $i \in [\ell]$ even such that $u_i \neq v_i$. Thus, we have that

$$\begin{pmatrix} 1 & 1 & 1 \\ u_i & v_i & w_i \\ u_i^2 & v_i^2 & w_i^2 \end{pmatrix}\begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

Doing row elimination, we get that

$$\begin{pmatrix} 1 & 1 & 1 \\ 0 & v_i - u_i & w_i - u_i \\ 0 & 0 & (w_i - u_i)(w_i - v_i) \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

Thus, if $w_i \neq u_i, v_i$ then the 3×3 matrix is nonsingular, implying that $a = b = c = 0$. Otherwise, if $w_i = u_i$ then since $v_i - u_i \neq 0$, we may deduce that $b = 0$. This implies that $a = -c$ so $au = aw$. Since u and w differ in at least one coordinate, we must have that $a = b = c = 0$. Finally if $w_i = v_i$, we deduce that $(v_i - u_i)(b + c) = 0$, so $b + c = 0$ and so $a = 0$. By the same reasoning as before, we have that $a = b = c = 0$.

Thus, S is 3-wise linearly independent. □

References

[AGKN13] Venkat Anantharam, Amin Gohari, Sudeep Kamath, and Chandra Nair. On Maximal Correlation, Hypercontractivity, and the Data Processing Inequality studied by Erkip and Cover. *arXiv:1304.6133 [cs, math]*, April 2013. arXiv: 1304.6133.

[AM09] Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *computational complexity*, 18(2):249–271, 2009.

[BG16] Joshua Brakensiek and Venkatesan Guruswami. New Hardness Results for Graph and Hypergraph Colorings. In Ran Raz, editor, *31st Conference on Computational Complexity (CCC 2016)*, volume 50 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:27, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[BGL15] Vijay V. S. P. Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Approximate hypergraph coloring under low-discrepancy and related promises. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, pages 152–174, 2015.

[BK09] N. Bansal and S. Khot. Optimal Long Code Test with One Free Bit. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 453–462, October 2009.

[BK10] Nikhil Bansal and Subhash Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. In *37th International Colloquium on Automata, Languages and Programming*, pages 250–261, 2010.

[BR15] Jonah Brown-Cohen and Prasad Raghavendra. Combinatorial optimization algorithms via polymorphisms. *CoRR*, abs/1501.01598, 2015.

[Cha13] Siu On Chan. Approximation resistance from pairwise independent subgroups. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 447–456, New York, NY, USA, 2013. ACM.

[Cha16] Siu On Chan. Approximation resistance from pairwise-independent subgroups. *J. ACM*, 63(3):27, 2016.

[CLP16] Ernie Croot, Vsevolod Lev, and Peter Pach. Progression-free sets in Z_4^n are exponentially small. *arXiv:1605.01506 [math]*, May 2016. arXiv: 1605.01506.

[CMM09] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. *ACM Trans. Algorithms*, 5(3), 2009.

[DMR09] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM Journal on Computing*, 39(3):843–873, 2009.

[DRS05] Irit Dinur, Oded Regev, and Clifford D. Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica*, 25(5):519–535, 2005.

[EG16] Jordan S. Ellenberg and Dion Gijswijt. On large subsets of F_q^n with no three-term arithmetic progression. *arXiv:1605.09223 [math]*, May 2016. arXiv: 1605.09223.

[EH08] Lars Engebretsen and Jonas Holmerin. More efficient queries in pcps for NP and improved approximation hardness of maximum CSP. *Random Struct. Algorithms*, 33(4):497–514, 2008.

[Eng04] Lars Engebretsen. The nonapproximability of non-boolean predicates. *SIAM J. Discrete Math.*, 18(1):114–129, 2004.

[FGRW12] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM J. Comput.*, 41(6):1558–1590, 2012.

[Geb41] Hans Gebelein. Das statistische Problem der Korrelation als Variations- und Eigenwertproblem und sein Zusammenhang mit der Ausgleichsrechnung. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 21(6):364–379, January 1941.

[GHM+11] Venkatesan Guruswami, Johan Håstad, Rajsekar Manokaran, Prasad Raghavendra, and Moses Charikar. Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM J. Comput.*, 40(3):878–914, 2011.

[GHS02] Venkatesan Guruswami, Johan Håstad, and Madhu Sudan. Hardness of approximate hypergraph coloring. *SIAM Journal on Computing*, 31(6):1663–1686, 2002.

[GK04] Venkatesan Guruswami and Sanjeev Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.*, 18(1):30–40, 2004.

[GL15] Venkatesan Guruswami and Euiwoong Lee. Strong inapproximability results on balanced rainbow-colorable hypergraphs. *Combinatorica*, 2015. Accepted.

[GL16] Venkatesan Guruswami and Euiwoong Lee. Nearly optimal NP-hardness of unique coverage. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1724–1730, 2016.

[GR08] Venkatesan Guruswami and Prasad Raghavendra. Constraint satisfaction over a non-boolean domain: Approximation algorithms and Unique-Games hardness. In *Proceedings of the 11th International Workshop on Approximation, Randomization and Combinatorial Optimization (APPROX)*, pages 77–90, 2008.

[GRSW16] Venkatesan Guruswami, Prasad Raghavendra, Rishi Saket, and Yi Wu. Bypassing UGC from some optimal geometric inapproximability results. *ACM Trans. Algorithms*, 12(1):6, 2016.

[GS13] Venkatesan Guruswami and Ali Kemal Sinop. Improved inapproximability results for maximum k-colorable subgraph. *Theory of Computing*, 9:413–435, 2013.

[Hå14] Johan Håstad. On the np-hardness of max-not-2. *SIAM Journal on Computing*, 43(1):179–193, 2014.

[Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.

[Has05] Gustav Hast. Approximating max k-CSP - outperforming a random assignment with almost a linear factor. In *32nd International Colloquium on Automata, Languages and Programming*, pages 956–968, 2005.

[Hir35] H. O. Hirschfeld. A Connection between Correlation and Contingency. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):520–524, October 1935.

[HK05] Johan Håstad and Subhash Khot. Query Efficient PCPs with Perfect Completeness. *Theory of Computing*, 1:119–148, September 2005.

[Hua12] Sangxia Huang. Approximation resistance on satisfiable instances for predicates strictly dominating parity. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:40, 2012.

[Hua14] Sangxia Huang. Approximation resistance on satisfiable instances for sparse predicates. *Theory of Computing*, 10:359–388, 2014.

[Hua15] Sangxia Huang. $2^{(\log N)^{1/4-o(1)}}$ hardness for hypergraph coloring. Technical report, Electronic Colloquium on Computational Complexity (ECCC), 2015.

[Kal02] Gil Kalai. A Fourier-theoretic perspective on the Condorcet paradox and arrow's theorem. *Advances in Applied Mathematics*, 29(3):412 – 426, 2002.

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing*, 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 Journal on Computing*, 37(1):319–357, 2007.

[KLS00] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the chromatic number. *Combinatorica*, 20(3):393–415, 2000.

[KMS16] Subhash Khot, Dor Minzer, and Muli Safra. On Independent Sets, 2-to-2 Games and Grassmann Graphs. Technical Report 124, Electrontic Colloquium on Computational Complexity (ECCC), August 2016.

[KR08] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-epsilon. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.

[KS14a] Subhash Khot and Rishi Saket. Hardness of coloring 2-colorable 12-uniform hypergraphs with ith $2^{(\log n)^{\Omega(1)}}$ colors. In *55th IEEE Annual Symposium on Foundations of Computer Science*, pages 206–215, 2014.

[KS14b] Subhash Khot and Rishi Saket. Hardness of finding independent sets in 2-colorable and almost 2-colorable hypergraphs. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1607–1625, 2014.

[Lee15] Euiwoong Lee. Hardness of graph pricing through generalized max-dicut. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 391–399, 2015.

[Lee16] Euiwoong Lee. Improved hardness for cut, interdiction, and firefighter problems. *CoRR*, abs/1607.05133, 2016.

[McD93] Colin McDiarmid. A random recolouring method for graphs and hypergraphs. *Combinatorics, Probability and Computing*, 2:363–365, 9 1993.

[MM14] Konstantin Makarychev and Yury Makarychev. Approximation Algorithm for Non-Boolean Max-k-CSP. *Theory of Computing*, 10:341–358, October 2014.

[MOO10] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *Ann. of Math. (2)*, 171(1):295–341, 2010.

[Mos10] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010.

[O'D08] Ryan O'Donnell. Some topics in analysis of boolean functions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 569–578, 2008.

[O'D14] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[OW09] Ryan O'Donnell and Yi Wu. Conditional hardness for satisfiable 3-CSPs. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 493–502, New York, NY, USA, 2009. ACM.

[Ré59] A. Rényi. On measures of dependence. *Acta Mathematica Academiae Scientiarum Hungarica*, 10(3-4):441–451, September 1959.

[Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 245–254, 2008.

[Rag09] Prasad Raghavendra. *Approximating NP-hard problems: Efficient algorithms and their limits*. PhD thesis, University of Washington, 2009.

[SS13] S. Sachdeva and R. Saket. Optimal Inapproximability for Scheduling Problems via Structural Hardness for Hypergraph Vertex Cover. In *2013 IEEE Conference on Computational Complexity*, pages 219–229, June 2013.

[ST98] Madhu Sudan and Luca Trevisan. Probabilistically checkable proofs with low amortized query complexity. In *39th Annual Symposium on Foundations of Computer Science*, pages 18–27, 1998.

[ST00] Alex Samorodnitsky and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 191–199, 2000.

[ST09] A. Samorodnitsky and L. Trevisan. Gowers Uniformity, Influence of Variables, and PCPs. *SIAM Journal on Computing*, 39(1):323–360, January 2009.

[Tre98] Luca Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, 1998.

[Tre00] Luca Trevisan. Approximating satisfiable satisfiability problems. *Algorithmica*, 28(1):145–172, 2000.

[TY10] Suguru Tamaki and Yuichi Yoshida. *A Query Efficient Non-adaptive Long Code Test with Perfect Completeness*, pages 738–751. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[uh15] user2566092 (http://math.stackexchange.com/users/87313/user2566092). Finding subsets of all the 3-dimensional vectors over finite field, such that every triple is independent. Mathematics Stack Exchange, 2015. URL:http://math.stackexchange.com/q/1306844 (version: 2015-05-31).

[Var14] Girish Varma. Reducing uniformity in Khot-Saket hypergraph coloring hardness reductions. *arXiv:1408.0262 [cs]*, August 2014. arXiv: 1408.0262.

[Wen13] Cenny Wenner. Circumventing d-to-1 for approximation resistance of satisfiable predicates strictly containing parity of width at least four. *Theory of Computing*, 9(23):703–757, 2013.

[Wit75] H. Witsenhausen. On Sequences of Pairs of Dependent Random Variables. *SIAM Journal on Applied Mathematics*, 28(1):100–113, January 1975.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il