hexagon:2611.00619v1

Circle configurations and the hardness of promise graph and hypergraph coloring

Yuwen Zhou
University of British Columbia Okanagan
zhouyw07@student.ubc.ca

November 2026

Abstract

In 2017 Brakensiek and Guruswami conjectured that promise graph homomorphism is NP-hard for every non-bipartite source and loopless target admitting a homomorphism from the source. At CCC 2016 they conjectured hardness of properly coloring a uniform hypergraph promised a strong coloring with one more color than its uniformity. We prove both conjectures by a configuration theorem on the circle. For any fixed finite configuration of distinct rational phases, the theorem gives a reduction whose satisfiable instances place every edge near a translate of the configuration and whose unsatisfiable instances have arbitrarily small independence ratio. An antipodal pair gives hardness for every odd cycle against every fixed clique of size at least three. A regular h-gon gives a strong $(h+1)$-coloring and a rainbow $(h-1)$-coloring of the same instance. A circle coloring built from single-coordinate changes gives hardness of $\mathrm{PCSP}(R, \mathrm{NAE}_h^{(d)})$ for every nonempty symmetric reconfigurable relation R with $R \to \mathrm{NAE}_h^{(d)}$, and the monochromatic-clique dichotomy of Banakh, Barsukov, and Nakajima without the Rich 2-to-1 Conjecture. For triples we obtain hardness of linearly ordered coloring for all fixed palettes $3 \le c \le d$. The proof subtracts an average of translated distance functions to detect an independent set, then uses conditional mean zero and a finite Ramsey alignment argument to decode a Label Cover assignment.

1 Introduction

For finite relational structures $A \to B$, the decision problem $\mathrm{PCSP}(A, B)$ accepts inputs X with a homomorphism $X \to A$ and rejects inputs X with $X \not\to B$. A homomorphism preserves every relation. For graphs it sends each edge to an edge. Taking A and B to be cliques gives approximate graph coloring.

For a hypergraph H, let $\alpha(H)$ be the maximum size of a set containing no entire edge. A proper coloring partitions the vertices into independent sets. A strong coloring gives distinct colors to all vertices of each edge. A rainbow coloring requires every edge to see every color of the specified palette. We use finite simple hypergraphs. An edge of uniformity h is a set of h distinct vertices.

Write $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ and give it the distance $d_{\mathbb{T}}(a, b) = \min_{n \in \mathbb{Z}} |a - b - n|$, computed with real lifts. The phase map in the following theorem is one map on the entire vertex set. Only the translation and the ordering vary from edge to edge.

Theorem 1.1 (Configuration theorem)**.** *Fix $h \ge 2$, distinct rational offsets $u_0 = 0, u_1, \dots, u_{h-1}$ in $\mathbb{T}$, a rational $\rho > 0$ such that*

$$2\rho < \min_{i \ne j} d_{\mathbb{T}}(u_i, u_j),$$

and a real $\delta > 0$. There is a deterministic polynomial-time reduction from 3SAT *to nonempty h-uniform hypergraphs H with these properties.*

(i) *If the input is satisfiable, there is a map $\varphi : V(H) \to \mathbb{T}$. Every edge has an ordering $x_0, \dots, x_{h-1}$ and a $t \in \mathbb{T}$ for which*

$$d_{\mathbb{T}}(\varphi(x_j), t + u_j) \le \rho \qquad (0 \le j < h).$$

(ii) *If the input is unsatisfiable, then $\alpha(H) < \delta |V(H)|$.*

The reduction outputs an explicit unweighted hypergraph. Its polynomial degree and constants may depend on the fixed data.

1.1 Coloring consequences

The first two applications use equal half-open intervals of the circle. For an odd cycle, their cyclic order matters. For a regular polygon, two partitions of different widths read different colorings from the same phase map.

Corollary 1.2 (Odd cycles and graph homomorphisms)**.** *For every odd $g \ge 3$ and every $0 < \delta < 1/3$, there is a deterministic polynomial-time reduction from* 3SAT *to nonempty graphs such that satisfiable instances map to C_g and unsatisfiable instances have independence ratio less than δ. Consequently* $\mathrm{PCSP}(C_g, K_c)$ *is NP-hard for every $c \ge 3$. More generally,* $\mathrm{PCSP}(A, B)$ *is NP-hard for every pair of finite loopless undirected graphs $A \to B$ with A non-bipartite.*

The last statement is Conjecture 1.2 of Brakensiek and Guruswami in the April 2017 revision of ECCC TR16-183 [8]. The bound $\delta < 1/3$ keeps the two promises disjoint for every C_g. Indeed, write $g = 2k + 1$. A uniform rotation of a fixed independent k-set of C_g contains each vertex with probability k/g. Its preimage under $X \to C_g$ has expected size $k|V(X)|/g \ge |V(X)|/3$.

Corollary 1.3 (Strong and rainbow colorings)**.** *For every $h \ge 3$ and $\delta > 0$, a deterministic polynomial-time reduction from* 3SAT *produces an h-uniform hypergraph with a strong $(h+1)$-coloring and a rainbow $(h-1)$-coloring in the satisfiable case, and independence ratio below δ in the unsatisfiable case.*

For all fixed $k, c \ge 2$ except $(k, c) = (2, 2)$, it is NP-hard to find a proper c-coloring of a k-uniform hypergraph promised to be strongly $(k+1)$-colorable. For $k \ge 3$ this also holds under the rainbow $(k-1)$-coloring promise and under both promises together.

The strong-coloring assertion settles CCC 2016 Conjecture 1.2. The rainbow boundary was posed separately in the table on page 14:4 of the same paper [7].

For a palette of size d, let $\mathrm{NAE}_h^{(d)}$ be the relation of nonconstant h-tuples. A relation is *symmetric* if it is closed under coordinate permutations. It is *reconfigurable* if any two of its tuples are joined inside the relation by a path whose steps change one coordinate. We identify a relation with its one-relation structure.

Write $\mathcal{F}_k^{(c)}$ for the family consisting of a monochromatic k-clique in each of c colors. An arrow $\mathcal{F}_k^{(c)} \to \mathcal{F}_\ell^{(d)}$ means that a fixed map of palettes takes every coloring avoiding the first family to a coloring avoiding the second. The corresponding PMMSNP problem accepts graphs admitting the first coloring and rejects graphs admitting no second coloring [6].

Corollary 1.4 (Reconfigurable relations and monochromatic cliques)**.** *Let $R \subseteq A^h$ be a nonempty symmetric reconfigurable relation on a finite set, where $h \ge 2$. For every d with $R \to \mathrm{NAE}_h^{(d)}$, the problem* $\mathrm{PCSP}(R, \mathrm{NAE}_h^{(d)})$ *is NP-hard.*

Let c, d, k, ℓ be positive integers with $\mathcal{F}_k^{(c)} \to \mathcal{F}_\ell^{(d)}$. If $\ell \geq c(k-1)$, then $\mathrm{PMMSNP}(\mathcal{F}_k^{(c)}, \mathcal{F}_\ell^{(d)})$ *is solvable in polynomial time via affine integer programming. If $\ell < c(k-1)$, then it is NP-hard.*

A linearly ordered coloring gives each edge a unique vertex of maximum color. Let LO_c denote the relation of triples with a unique maximum on an ordered c-element palette.

Corollary 1.5 (Linearly ordered triples)**.** *For every fixed $3 \leq c \leq d$, the problem* $\mathrm{PCSP}(\mathrm{LO}_c, \mathrm{LO}_d)$ *on 3-uniform hypergraphs is NP-hard.*

1.2 Proof overview

A point in a phase grid assigns one circle value to each possible answer of a question tuple. An answer projection π lifts a phase vector y to $y \circ \pi$. We give this link length zero. Moving within a grid costs the maximum change in any coordinate. A satisfying assignment chooses a coordinate in every grid, compatibly with every projection. Evaluating that coordinate gives a circle map that cannot increase path length. Adding u to every coordinate adds u to the evaluation. We can therefore test a configuration by looking for one grid point w near whose prescribed translates the vertices of an edge all lie.

The layered Label Cover framework with phase grids, list decoding, and distributional alignment was introduced by OpenAI in item 106 for the independent-set gap on three-colorable graphs [26]. The accompanying scope document describes the Lean formalization of that graph theorem [31].

Consider an independent set and let b be a truncated distance function equal to one at its locations. At a selected location x, at least one of the translates of x must have bump value zero. Otherwise the nearby selected vertices would form an edge. Thus the difference between $b(x)$ and the average of its translated values is at least $1/h$. Extend the bump to the torus before taking this difference. The resulting function has integral zero on every diagonal circle, including circles of product measure zero.

The sampled vertex combines several layers. Each layer contributes independent rotation entries, multiplied by a coefficient in $\{0, 1/m, \ldots, 1\}$. These real rotation entries give a product space even when answer projections identify many coordinates. Austin's theorem approximates the function of the rotations by a function depending on boundedly many entries. Their indices are answer lists. If two lists from separated layer sets project to a common final answer, they give compatible candidates on the two sides of a single Label Cover test. Small Label Cover value makes such intersections rare.

The remaining issue is to assign a layer to every final answer. When layer i has coefficient zero, we use the list for the other layers. A final answer in that list must be assigned somewhere other than i. A finite Ramsey argument chooses a distribution on layer sets before the lists are known. With high probability all these assignments can be made at once. The proof uses the fact that a label occurring on several lists has a point common to the intervals spanned by their index sets. Two test points in a gap between indices control the probability of rounding this point to the deleted index.

Give each chosen layer an auxiliary uniform phase. Each answer reads the auxiliary phase of its assigned layer. Simultaneously translating these phases translates every input coordinate, so the function of the auxiliary phases has mean zero for each fixed choice of the rotations and coefficients. A large mean square forces an influential layer. A coefficient equal to one can absorb all the auxiliary phases. This makes a zero-coefficient layer nearly uninfluential by the assignment property. At a positive coefficient, decreasing the coefficient by one level changes the function very little. The geometric law on levels and a bound on the number of influential layers make this second possibility

rare as well. The mean square bound passes to the grid and bounds the density of the independent set.

1.3 Related work

Hell and Nešetřil classified the complexity of homomorphisms to a fixed undirected graph. A loopless target gives a polynomial-time problem when it is bipartite and an NP-complete problem otherwise [25]. The approximate coloring problem goes back to Garey and Johnson [20]. Krokhin, Opršal, Wrochna, and Živný observed that the graph promise conjecture is equivalent to hardness for all odd cycles against all fixed cliques of size at least three [29].

Khanna, Linial, and Safra proved hardness of four-coloring a three-colorable graph [27]. Guruswami and Khanna gave another proof and obtained bounded-degree hardness [21]. For a-colorable inputs, $a \geq 3$, Brakensiek and Guruswami reached $2a - 2$ colors [7]. Bulín, Krokhin, and Opršal reached $2a - 1$, including five colors at $a = 3$ [11, 5]. Krokhin and Opršal proved hardness against K_3 for every non-bipartite three-colorable source [28]. Wrochna and Živný reached $\binom{a}{\lfloor a/2 \rfloor} - 1$ colors for $a \geq 4$. Their topological method also gives hardness for square-free targets and circular cliques of ratio strictly between two and four [37, 29]. Krokhin, Opršal, Wrochna, and Živný's arc-graph reduction shows that hardness against all constant palettes for one fixed clique source implies the same result for K_3 [29, Theorem 1.8]. Avvakumov, Filakovský, Opršal, Tasinato, and Wagner proved hardness against K_4 for every non-bipartite four-colorable source, settling the conjecture for four-colorable targets [4].

Dinur, Mossel, and Regev obtained three-colorable completeness and small independent sets from their fish-shaped Label Cover conjecture [14]. Under ordinary perfect-completeness d-to-1 hardness, Guruswami and Sandeep obtained the density gap with $2d$-colorable completeness and constant-palette hardness with three-colorable completeness [23, 29]. Fei, Minzer, and Wang proved perfect-completeness 4-to-1 hardness. Their Corollaries 1.7 and 1.9 give, respectively, constant-palette hardness for three-colorable graphs and the density gap for eight-colorable graphs [18]. OpenAI's ordinary perfect-completeness 2-to-1 theorem also implies the constant-palette conclusion by the preceding reductions [19, 23, 29]. In the odd-cycle versus clique formulation, these results leave odd $g \geq 5$ against $c \geq 5$.

Braverman, Khot, Lifshitz, and Minzer derived the three-colorable density gap from the Rich 2-to-1 Games Conjecture [10]. Banakh, Barsukov, and Nakajima observed that their conditional Theorem 8 implies the full graph conjecture [6]. Richness concerns the distribution at a fixed left question. Braverman et al. require a uniform partition into pairs. Banakh et al. require uniformity over the labeled two-to-one maps. Ordinary two-to-one hardness specifies the fiber sizes [10, Definition 1.16][6, Definition 27].

For strong colorings, Brakensiek and Guruswami proved two-color hardness with $\lceil 3k/2 \rceil$ strong colors for every $k \geq 3$, and with $k+1$ strong colors for $k = 3, 4$ [7, Theorem 1.3]. Under their induced V Label Cover conjecture, they obtained separate density gaps for $\lceil k + \sqrt{k} - 1/2 \rceil$ strong colors when $k \geq 2$, and $k - 1$ rainbow colors when $k \geq 3$ [9, Theorem 1.2 and footnote 6]. Guruswami and Lee obtained small independent sets under an almost balanced rainbow promise [22]. Austrin, Bhangale, and Potukuchi proved two-color hardness with $k - 2\lfloor \sqrt{k} \rfloor$ rainbow colors in the range where this is a rainbow promise, and settled the boundary $k - 1$ for $k = 4$ [1]. Guruswami and Sandeep extended boundary two-color hardness to $k = 5, 6$ [24].

For triples, rainbow two-colorability is ordinary two-colorability. Dinur, Regev, and Smyth proved hardness against every constant proper palette under this promise [15]. Also, pairing the colors of a strong four-coloring gives a rainbow two-coloring. Thus the $k = 3$ strong-coloring result of Brakensiek and Guruswami already supplied both promises for two-color hardness. The original

ECCC abstract of Fei, Minzer, and Wang, posted September 14, 2026, announced the density gap for ordinary two-colorable triples. It is Corollary 1.10 of their full paper [18].

Banakh, Barsukov, and Nakajima proved the two statements of Corollary 1.4 under the Rich 2-to-1 Conjecture. We use their LICS numbering. Theorems 5, 8, and 9 and Corollary 10 correspond to Theorems 1.5, 1.8, and 1.13 and Corollary 1.14 in the second arXiv version [6]. Their Theorem 8 also gives a conditional predecessor for simultaneous strong and rainbow completeness. Take the product of the injective h-tuples on $h+1$ colors and the surjective h-tuples on $h-1$ colors. The first relation can change an injection using a spare color. In the second, mark one position of each color. To move a mark to the vacant position, write its color there and release its old position. Swaps with the vacancy permute the marks arbitrarily. After placing them as prescribed, set the unmarked entry to its target value. Each move preserves surjectivity. The product reconfigures one factor at a time. Projecting onto the second factor and grouping its colors into two nonempty parts gives a map to $\mathrm{NAE}_h^{(2)}$. A coloring by the product gives exactly the two promised colorings. The proof of their Claim 34 decodes from a single edge-free set of measure at least $1/d$ for the fixed target palette d [6].

For linearly ordered triples, Filakovský, Nakajima, Opršal, Tasinato, and Wagner proved hardness from three colors to four [17]. Krokhin and Vagnozzi proved hardness from two to three [30]. Nakajima and Živný's palette reduction raises both palettes by one [34, Theorem 5.6]. Banakh et al. state the resulting adjacent-palette hardness for all $c \geq 2$ as Theorem 9, citing the Krokhin–Vagnozzi preprint for the starting case. Their Corollary 10 gives all $3 \leq c \leq d$ conditionally [6]. For uniformity at least four, Nakajima, Verwimp, Wrochna, and Živný proved hardness for all $3 \leq c \leq d$ [33].

2 Label Cover and coordinate approximation

Write $[n] = \{1, \dots, n\}$. All product metrics are sup metrics. Uniform measure on a real cube means product Lebesgue probability. Uniform measure on a product of circles means product Haar probability. Conditional expectations and L^2 norms below use these laws.

2.1 The source problem

A Label Cover instance has finite question sets U, V, nonempty answer sets Σ_L, Σ_R, and a nonempty explicit list $\mathcal{C}$ of test occurrences. Each occurrence e specifies a left question p_e, a right question q_e, and a total map $P_e : \Sigma_L \to \Sigma_R$. Its value is

$$\mathrm{val}(\mathcal{C}) = \max_{\ell: U \to \Sigma_L,\, r': V \to \Sigma_R} \Pr_{e \in \mathcal{C}} \{P_e(\ell(p_e)) = r'(q_e)\}.$$

The law on occurrences is uniform. Occurrences with the same pair of questions may have different projections.

Theorem 2.1 (Perfect-completeness Label Cover)**.** *For every fixed rational $0 < \sigma < 1$, there are fixed finite answer sets and a deterministic polynomial-time reduction from* 3SAT *to explicit Label Cover instances as above. Satisfiable formulas have value one. Unsatisfiable formulas have value at most σ.*

This is the projection-game form of the PCP theorem with parallel repetition [2, 36]. A constant gap is amplified with a fixed number of repetitions chosen for σ. A product of uniform occurrence lists is again an explicit uniform occurrence list, and product projections are total. Perfect completeness is preserved at each repetition. See also the quantitative treatment of Dinur and Steurer and the total-projection formulation of Dinur, Khot, Kindler, Minzer, and Safra [16, 13]. The questions unused by all occurrences can be discarded. No regularity or bound on projection fibers will enter the proof.

2.2 Two bounds on coordinate influences

We use the following consequence of Austin's junta theorem. For every $L < \infty$ with $L \geq 0$ and every $\tau > 0$, there is an integer $J(L, \tau) \geq 1$ with the following property. For any n and any L-Lipschitz $G : [0,1]^n \to [-1,1]$, there is a coordinate set S, $|S| \leq J(L,\tau)$, such that

$$\|G - \mathbb{E}[G \mid x_S]\|_2 < \tau.$$

Indeed, apply Austin's Theorem 1.1 in arXiv version 4 to the uniform interval with modulus $t \mapsto Lt$. Take the L^1 error below $\tau^2/2$. The pointwise difference has absolute value at most two, giving the stated L^2 bound [3]. The quotient map from the cube to the product of circles preserves uniform probability and is nonexpanding. The same approximation therefore holds on a product of circles. A function of the selected coordinates will be called a junta.

For $H \in L^2(\mathbb{T}^n)$, let E_i average coordinate i while retaining all others. Set

$$Q_i = 1 - E_i, \qquad D_i(H) = \|Q_i H\|_2.$$

Both E_i and Q_i are orthogonal projections. In particular, changing a function by η in L^2 changes D_i by at most η.

Lemma 2.2 (Influence detection and count)**.** *For every $L < \infty$, $L \geq 0$, and $\varepsilon > 0$, there are $\beta > 0$ and $K \geq 1$ independent of n such that every L-Lipschitz $H : \mathbb{T}^n \to [-1, 1]$ satisfies*

$$\begin{aligned} &\mathbb{E}H = 0, \quad \mathbb{E}H^2 > \varepsilon \quad \Longrightarrow \quad D_i(H) \geq \beta \text{ for some } i, \\ &\#\{i : D_i(H) \geq \beta/2\} \leq K. \end{aligned}$$

Proof. Set $J_0 = J(L, \sqrt{\varepsilon}/2)$. Choose a set S of at most J_0 coordinates so that $G = \mathbb{E}[H \mid z_S]$ satisfies $\|H - G\|_2 < \sqrt{\varepsilon}/2$. If H has mean zero and $\|H\|_2 > \sqrt{\varepsilon}$, then G has mean zero and $\|G\|_2 > \sqrt{\varepsilon}/2$.

To bound this last norm, expand H using the commuting projections E_i, Q_i. For $A \subseteq [n]$, put

$$H_A = \left(\prod_{i \in A} Q_i\right)\left(\prod_{i \notin A} E_i\right) H.$$

Their sum is H, since $Q_i + E_i = 1$. Two distinct summands are orthogonal by averaging at a coordinate in their symmetric difference. Conditioning on z_S retains exactly the summands with $A \subseteq S$. The constant summand vanishes. Consequently

$$\frac{\varepsilon}{4} < \|G\|_2^2 = \sum_{\varnothing \neq A \subseteq S} \|H_A\|_2^2 \leq \sum_{i \in S} \sum_{A \ni i} \|H_A\|_2^2 = \sum_{i \in S} D_i(H)^2.$$

Take $\beta = \sqrt{\varepsilon/(4J_0)}$.

For the count, approximate H to error less than $\beta/2$ by a junta G' on at most $K = J(L, \beta/2)$ coordinates. Outside its coordinate set $Q_i G' = 0$, so $D_i(H) \leq \|H - G'\|_2 < \beta/2$. This part holds for every mean. □

3 The reduction

Fix the configuration and density in Theorem 1.1. The construction uses integers $r \geq s \geq 2$, $m \geq 1$, and $P \geq 1$. Take P divisible by the denominators of all the offsets. Other finite data are a rational number $0 < \lambda < 1$ and a rational probability law μ on $\binom{[r]}{s}$. The choices in Section 8 fix all these data before the input formula is read. Put

$$D = mP, \qquad \Gamma = (D^{-1}\mathbb{Z})/\mathbb{Z}.$$

We describe the reduction on the output of Theorem 2.1.

3.1 A space of phase assignments

At layer $i \in [r]$, form all question tuples and the corresponding answer set

$$\mathbf{q} \in V^{i-1} \times U^{r-i}, \qquad M_i = \Sigma_R^{i-1} \times \Sigma_L^{r-i}.$$

Both products have $r-1$ positions. Occurrences of the same question at different positions remain separate positions. At adjacent layers $i, i+1$, take all pairs of question tuples that agree off position i. For each occurrence between their questions at position i, impose the map $M_i \to M_{i+1}$ applying its projection at that position and the identity elsewhere.

Place a separate copy of Γ^{M_i} at each question tuple of layer i. A grid point x is an assignment of a phase x_κ to every answer κ of that tuple. Let $\mathscr{X}$ be the disjoint union of the grids. Define an undirected weighted link graph on $\mathscr{X}$. Within a grid, link each pair with length its sup circle distance. For an imposed map $\pi : M_i \to M_{i+1}$, link each point y of its target grid to $y \circ \pi$ of its source grid with length zero. The direction of this pullback is opposite to the answer map.

Let dist be shortest-path distance in the link graph, with value infinity between components. It is an extended pseudometric. Different grid points can have distance zero. For $u \in \Gamma$, let $S_u x = x + u\mathbf{1}$ in each grid. Within-grid lengths are translation invariant. On a projection link,

$$(S_u y) \circ \pi = S_u(y \circ \pi).$$

Thus S_u and its inverse take links to links of the same length. Each S_u is an isometry of dist. We use the same notation for arbitrary diagonal translations on a continuous product torus.

3.2 Configuration tests

Choose an integer $q \geq h$ with $(h-1)/q < \delta$. Denote by $K_q^{(h)}$ the hypergraph of all h-subsets of $[q]$. First test whether there exist $w \in \mathscr{X}$ and $i < j$ with

$$\operatorname{dist}(S_{u_i} w, S_{u_j} w) \leq 2\rho. \tag{1}$$

If so, output $K_q^{(h)}$ and stop. We call this the fallback. On the other branch, the closed balls of radius ρ about the h translates of any one center are pairwise disjoint.

The vertices on this branch will have locations $x_v \in \mathscr{X}$. Declare a set of h distinct vertices to be an edge if it has an ordering $v_0, \ldots, v_{h-1}$ and a witness $w \in \mathscr{X}$ such that

$$\operatorname{dist}(x_{v_j}, S_{u_j} w) \leq \rho \qquad (0 \leq j < h). \tag{2}$$

Both tests range over the entire grid union. In particular, w need not be the location of an output vertex. The order in (2) is an existential witness for an unordered edge.

3.3 The law of a vertex

Sample $r-1$ independent uniform occurrences $e_1, \dots, e_{r-1}$ from $\mathcal{C}$. Write their questions as (p_i, q_i). They determine the occurrence chain

$$\mathbf{q}_i = (q_1, \dots, q_{i-1}, p_i, \dots, p_{r-1}) \qquad (i \in [r]).$$

For $i \le j$, let $\pi_{ij} : M_i \to M_j$ be the composite of the imposed maps along this chain. It acts at positions i through $j-1$. We use $\pi_{ii} = \mathrm{id}$ and $\pi_{jk}\pi_{ij} = \pi_{ik}$.

Independently sample $B \sim \mu$. In each layer $j \in B$, draw a coefficient t_j independently from the law

$$\mathcal{T}_m = \{0, 1/m, \dots, 1\}, \qquad \Pr(t_j = k/m) = p_k = \frac{\lambda^k}{\sum_{v=0}^{m} \lambda^v}. \tag{3}$$

Draw independent rotations $\widehat{\theta}_{j,\ell}$ uniformly from $\{0, 1/P, \dots, (P-1)/P\}$, for $j \in B, \ell \in M_j$. These are real representatives in $[0,1)$. We multiply by the coefficient before reducing modulo one.

It will be useful to use the same formula on other layer sets and on real rotations. For nonempty $I \subseteq [r]$, $c = \min I$, define the phase vector

$$X_I(t,\theta)_\kappa = \sum_{j \in I} t_j \theta_{j,\pi_{cj}(\kappa)} \pmod 1, \qquad \kappa \in M_c. \tag{4}$$

The location of the sampled outcome is $X_B(t, \widehat{\theta})$ in the grid at $\mathbf{q}_{\min B}$. Each summand has denominator dividing mP, so this is a point of the stated grid.

Proposition 3.1 (Deterministic encoding)**.** *For fixed parameters and alphabets the construction outputs a nonempty finite simple h-uniform hypergraph in deterministic polynomial time. In the sampling branch a uniform vertex has exactly the law just specified, followed by a uniform choice among the copies of its outcome.*

Proof. Let $N = |\mathcal{C}| \ge 1$. Each of the N^{r-1} occurrence chains has the same probability. Conditional on a chain, the probabilities of all remaining elementary outcomes are rational. Their denominators have a common multiple C_0 independent of the instance. Indeed, μ, the coefficient table, and the possible numbers of rotation entries are fixed. For each conditional outcome of probability w emit $C_0 w$ distinct vertices with its location. The count is C_0 per chain and $C_0 N^{r-1}$ overall. Uniform sampling now has the required law, including when distinct outcomes have the same location.

There are polynomially many question tuples and a constant number of grid points per tuple. Enumerate the links and multiply their lengths by D. The weights are nonnegative integers at most D. A minimum finite distance is realized on a simple path, so its integer representation has polynomial bit length. Exact all-pairs shortest paths give the comparisons in both tests. More explicitly, if $\rho = p/n$ and a is a finite scaled distance, the tests are $na \le 2pD$ and $na \le pD$, respectively. An infinite distance fails either comparison.

For the second test enumerate centers and ordered h-tuples of distinct vertex indices. Store each successful edge once under its sorted indices. This costs $O(h|\mathscr{X}|\,|V(H)|^h)$ comparisons after the distances are available. All loops here have polynomial total cost for fixed r, h and alphabets.

There is even an edge in the sampling branch. Fix a chain and a B of positive probability. Give its first layer coefficient one and the remaining layers coefficient zero. In the first block one may set all rotations equal to u_j, for each j, since their denominators divide P. These h positive-probability outcomes have distinct constant locations and witness an edge about the zero vector of that grid. The fallback also contains an edge. □

4 Completeness and circle colorings

Lemma 4.1 (Evaluation)**.** *A labeling satisfying every occurrence rules out the fallback and gives the phase map in Theorem 1.1.*

Proof. For a question tuple $\mathbf{q}$, let $\kappa_{\mathbf{q}}$ be the product of its assigned answers. Define $\varphi(x) = x_{\kappa_{\mathbf{q}}}$ on its grid. Evaluation is nonexpanding on within-grid links. If π is an imposed map from $\mathbf{q}$ to $\mathbf{q}'$, satisfaction gives $\pi(\kappa_{\mathbf{q}}) = \kappa_{\mathbf{q}'}$. Evaluations therefore agree across the corresponding zero links. Along any path, the circle distance between its endpoint evaluations is bounded by the sum of its link lengths. Hence

$$d_{\mathbb{T}}(\varphi(x), \varphi(y)) \le \operatorname{dist}(x, y), \qquad \varphi(S_u x) = \varphi(x) + u.$$

At any w, the distance between translates i, j is at least $d_{\mathbb{T}}(u_i, u_j) > 2\rho$. This excludes (1). Give vertex v the phase $\varphi(x_v)$. For an edge with witness w, nonexpansion in (2) gives the required inequalities with $t = \varphi(w)$. □

4.1 An antipodal pair

Use C_g on $\mathbb{Z}/g\mathbb{Z}$ with edges of difference ± 1. For two offsets $0, 1/2$, the two errors relative to a common translate give an error at most 2ρ in the displacement.

Lemma 4.2 (Odd-cycle rounding)**.** *Let $g = 2k + 1 \ge 3$. Suppose a graph has a phase map under which the displacement on every edge has circle distance at most $1/(4g)$ from $1/2$. Then the graph maps to C_g.*

Proof. At a vertex with real phase lift u, give the color

$$k^{-1}\lfloor gu \rfloor \pmod g.$$

Here $k^{-1} = -2 \bmod g$. Changing the lift adds a multiple of g to the bin index. On an edge, choose the other lift v so that

$$g(v - u) \in [k + 1/4, k + 3/4].$$

Writing $\{gu\}$ for the fractional part gives

$$\lfloor gv \rfloor - \lfloor gu \rfloor = \lfloor \{gu\} + g(v - u) \rfloor \in \{k, k + 1\}.$$

The interval includes its endpoints, and the floor assigns a boundary point to the bin starting there. Since $k + 1 \equiv -k \bmod g$, the colors differ by 1 or -1. □

Remark 4.3. For $g = 3$, Open Mathematics Courses, Exercise 6.5, observes that the three-bin completeness argument works for any antipodal displacement cutoff strictly below $1/6$ [12]. For the configuration theorem the displacement cutoff is twice the per-vertex radius. We use $\rho = 1/(8g)$.

4.2 A regular polygon

Lemma 4.4 (Two partitions of the circle)**.** *Use the offsets j/h, $0 \le j < h$, with $h \ge 3$ and $0 < \rho < 1/(2h(h + 1))$. On every hypergraph satisfying the phase condition of Theorem 1.1, the partitions into $h + 1$ and $h - 1$ equal half-open bins give a strong and a rainbow coloring, respectively.*

Proof. The distance between any two distinct ideal phases on an edge is at least $1/h$. The corresponding actual phases have distance at least $1/h - 2\rho > 1/(h+1)$. A bin of width $1/(h+1)$ cannot contain two such points.

Choose a lift t of the translation on an edge and write its lifts as $t + j/h + e_j$, with $|e_j| \le \rho$. Set $e_h = e_0$ for the next copy of the first vertex. Every successive gap lies in

$$[1/h - 2\rho,\, 1/h + 2\rho] \;\subset\; (0, 1/(h-1)).$$

The positive lower bound gives cyclic order, including the gap across the seam. If a lifted bin $[a, a + 1/(h-1))$ were empty, the last lift before a would be strictly less than a and the next would be at least $a + 1/(h-1)$. The gap bound excludes this. Both partitions use the original phase map on all vertices. □

4.3 A path of legal tuples

Lemma 4.5 (Circle coloring for a relation)**.** *Suppose a symmetric relation $R \subseteq C^h$ contains $a = (a_0, \ldots, a_{h-1})$ and a single-coordinate path in R from a to $(a_1, \ldots, a_{h-1}, a_0)$. There are a rational $0 < \rho_* < 1/(2h)$ and a map $\chi : \mathbb{T} \to C$ such that*

$$\big(\chi(t + j/h + e_j)\big)_{0 \le j < h} \in R \quad \textit{whenever } t \in \mathbb{T}, \quad |e_j| \le \rho_*.$$

Every nonempty symmetric reconfigurable relation has this property.

Proof. Write the path as $a^0, \ldots, a^m$. Schedule its steps at distinct increasing interior times of $[0, 1/h)$. Let $W(t)$ be the current tuple, using the new value at a step. It equals a^0 near zero and a^m near $1/h$. Extend it to the real line by

$$W_j(t + 1/h) = W_{j+1}(t) \qquad (j \bmod h). \tag{5}$$

At the seam $1/h$, the new value of entry j is $a^0_{j+1} = a^m_j$, its value on the left. Rotation gives the same equality at every seam. After h intervals each entry returns to its original position. The extension is periodic with period one. Set $\chi(t) = W_0(t)$. Then $\chi(t + j/h) = W_j(t)$.

If there are steps, their times repeat with period $1/h$. Choose $2\rho_*$ smaller than $1/h$ and smaller than the minimum spacing of successive step times, including the spacing across a seam. A rational choice exists. With no steps, just take $2\rho_* < 1/h$. A closed interval $[t - \rho_*, t + \rho_*]$ meets at most one step. If it meets one, the tuples on its two sides belong to R by symmetry and differ in one coordinate. Sampling entry j at an arbitrary $t + e_j$ therefore produces one of these two tuples. This includes sampling exactly at the step. If the interval meets none, it gives $W(t)$. Equation (5) identifies the sampled entries with $\chi(t + j/h + e_j)$. Finally, symmetry puts the rotation of any $a \in R$ in R, and reconfigurability supplies the path. □

Lemma 4.6 (Merging colors on triples)**.** *A strong four-coloring of a 3-uniform hypergraph gives an* LO_3*-coloring.*

Proof. Send $0, 1, 2, 3$, respectively, to $0, 0, 1, 2$. A triple containing strong color 3 has a unique maximum of 2. A triple omitting it has strong colors $0, 1, 2$ and a unique maximum of 1. □

5 Soundness I. From an independent set to lists

The soundness framework of OpenAI [26] turns functions on tuple tori into lists of Label Cover answers and tests their coordinate influences. We give the argument for configuration edges, using diagonal integration and lists for subchains of one fixed size.

Fix the Label Cover instance and a nonempty independent set $\mathcal{A}$ in the sampling branch. The choice of $\mathcal{A}$ precedes every random choice in the analysis. Its set of locations is $A = \{x_v : v \in \mathcal{A}\} \subseteq \mathscr{X}$. The vertex set retains multiplicities throughout. Only A forgets them.

5.1 A function with zero integral on each diagonal circle

Lemma 5.1 (Translated bump)**.** *Put $L = 2/\rho$. There is a function $f : \mathscr{X} \to [-1, 1]$ constant on zero links such that $f(x) \geq 1/h$ for every $x \in A$. On the grid at each tuple $\mathbf{q}$ it has an L-Lipschitz extension $F_{\mathbf{q}} : \mathbb{T}^{M_i} \to [-1, 1]$ satisfying*

$$\int_{\mathbb{T}} F_{\mathbf{q}}(S_v y)\, dv = 0 \qquad (y \in \mathbb{T}^{M_i}). \tag{6}$$

Along any chain, write $F_i = F_{\mathbf{q}_i}$. For $i \leq j$,

$$|F_i(y \circ \pi_{ij}) - F_j(y)| \leq 2L/D \qquad (y \in \mathbb{T}^{M_j}). \tag{7}$$

Proof. On the whole grid union define

$$b(x) = (1 - \operatorname{dist}(x, A)/\rho)_+, \qquad f(x) = b(x) - \frac{1}{h} \sum_{j=0}^{h-1} b(S_{u_j} x).$$

Take $b = 0$ at infinite distance. Suppose $x \in A$ and every summand $b(S_{u_j} x)$ is positive. For each j, choose $y_j \in A$ with $\operatorname{dist}(y_j, S_{u_j} x) < \rho$. Failure of (1) makes the closed radius-ρ balls at this center disjoint. Thus the y_j are distinct. Choose one selected vertex at each. These h vertices form an edge by (2) with $w = x$. This contradicts independence. At least one summand is zero. Since $b(x) = 1$ and the other summands are at most one, $f(x) \geq 1/h$.

On a component meeting A, distance to A is 1-Lipschitz for the path pseudometric. On a component disjoint from A, b is constant. A within-grid link bounds path distance from above by sup circle distance. It follows that b restricted to a grid is $1/\rho$-Lipschitz. Points at zero distance have equal values under b, as do their simultaneous translates. Hence f agrees across zero links.

At the tuple $\mathbf{q}$, extend the bump by the McShane formula and clip to the interval $[0, 1]$ [32]. Explicitly, with $G_{\mathbf{q}}$ denoting this tuple's grid, let

$$B_{\mathbf{q}}(y) = \min \left\{ 1,\ \min_{x \in G_{\mathbf{q}}} \left(b(x) + \rho^{-1} d_\infty(x, y) \right) \right\}.$$

Every term is nonnegative. The inner minimum is $1/\rho$-Lipschitz. At a grid point y, every term is at least $b(y)$ and the term $x = y$ gives equality. The outer minimum preserves this value and the Lipschitz bound. Now put

$$F_{\mathbf{q}}(y) = B_{\mathbf{q}}(y) - \frac{1}{h} \sum_{j=0}^{h-1} B_{\mathbf{q}}(S_{u_j} y).$$

Every required shift preserves the grid, so $F_{\mathbf{q}}$ extends f. Its two terms take values in $[0, 1]$ and each has Lipschitz constant at most $1/\rho$. On the diagonal circle $S_v y$, the integral of each translated bump equals the integral of the first bump by $v \mapsto v + u_j$. This gives (6) at every y.

Finally, round $y \in \mathbb{T}^{M_j}$ to a grid point y' at sup distance at most $1/D$. Pullback preserves this upper bound. The grid points y' and $y' \circ \pi_{ij}$ are joined by a sequence of zero links, so f agrees at the two endpoints. The two rounding errors give (7), independently of the length of the composite. □

5.2 Lists on a subchain

The only subchains that will need lists have $s-1$ layers. They are the possible results of deleting one layer from a sampled s-set. Fix $\gamma > 0$. For $I \in \binom{[r]}{s-1}$ and a fixed coefficient vector $t \in \mathcal{T}_m^I$, consider

$$G_{I,t}(\theta) = F_{\min I}\big(X_I(t,\theta)\big), \qquad \theta \in \prod_{j\in I}[0,1]^{M_j}. \tag{8}$$

The rotation entries now have continuous product law. In one coordinate of X_I, each block contributes one scalar multiplied by a number in $[0,1]$. Hence changing the whole rotation array by sup distance η changes the phase vector by at most $(s-1)\eta$ in circle sup distance. This holds for all projection maps, including those with large fibers. The function $G_{I,t}$ is therefore $(s-1)L$-Lipschitz.

Lemma 5.2 (Local choice of lists)**.** *For each fixed chain and each $I \in \binom{[r]}{s-1}$, there are sets $S_I^j \subseteq M_j$, $j \in I$, with*

$$\sum_{j\in I} |S_I^j| \le d, \qquad d = J((s-1)L,\gamma)(m+1)^{s-1}. \tag{9}$$

For every $t \in \mathcal{T}_m^I$ there is a junta $g_{I,t}$ on entries (j,ℓ) with $\ell \in S_I^j$ such that $\|G_{I,t} - g_{I,t}\|_2 < \gamma$. The lists depend only on the tuples in I, their fixed functions, and the projections between them.

Proof. Apply the junta bound of Section 2 to (8) at every coefficient vector. Order the scalar coordinates once and for all. Among the sets of size at most $J((s-1)L,\gamma)$ whose coordinate conditional expectation has error less than γ, choose the first in a fixed ordering. Use that conditional expectation as $g_{I,t}$. Union the chosen indices within each block over the $(m+1)^{s-1}$ vectors t. This proves the bound and the simultaneous approximation. The choice rule uses only the data defining (8). It is part of the analysis and need not be run by the reduction. □

Project the lists to the answer set at the last layer:

$$A_I = \bigcup_{j\in I} \pi_{jr}(S_I^j) \ \subseteq M_r. \tag{10}$$

Each A_I has size at most d. The projection in this formula can use occurrences outside I. We decode from the selected coordinates *before* this terminal projection.

Lemma 5.3 (Decoding separated lists)**.** *If* $\mathrm{val}(\mathcal{C}) \le \sigma$ *and* $I, J \in \binom{[r]}{s-1}$ *satisfy* $\max I < \min J$*, then*

$$\Pr_{\text{chain}}(A_I \cap A_J \neq \varnothing) \le d^2\sigma.$$

Proof. Put $k = \max I$, and condition on all occurrences except e_k. Each tuple in I has left question p_k at position k. Its remaining entries are fixed. All projections internal to that subchain use occurrences with indices below k. The left list rule consequently depends on e_k only through p_k. On J, position k already has right question q_k and all internal projections avoid e_k. Its list rule depends on e_k only through q_k. The family of functions $F_{\mathbf{q}}$ was fixed for the entire instance in Lemma 5.1. Looking up a function at a tuple respects these dependencies even for different occurrences with the same own question.

On each side collect the entries at position k of all its selected answers. This gives at most d left candidates as a function of the left question and at most d right candidates as a function of the right question. Order and pad the lists to length d with arbitrary answers, allowing an initially empty list. For each pair of list indices, reading those entries at all own questions defines one labeling of the original instance.

If the terminal lists intersect, two selected product answers have equal terminal images. At position k, the left answer was acted on by exactly P_{e_k}, and the right answer was unchanged. Thus one of the d^2 labelings satisfies e_k. Conditional on the background this occurrence still has its original uniform law. Each labeling satisfies it with probability at most σ. A union bound and then averaging over the background prove the claim. $\square$

6 Soundness II. Finite distributional alignment

For $B \in \binom{[r]}{s}$ we seek a map $a : M_r \to B$ such that

$$a(A_{B\setminus\{i\}}) \subseteq B \setminus \{i\} \qquad (i \in B). \tag{11}$$

In fact, (11) is equivalent to the absence of a label common to all s deletion lists. For each label, one may assign any index whose deletion list omits it. Disjointness on separated subchains does not itself give this condition on a prescribed B. For example, with $s = 3$ the same label can occur on all three pairs of a triple without occurring on two separated pairs. The distribution in the next lemma works for every list family.

Lemma 6.1 (Distributional alignment for deletions)**.** *For integers $s \geq 2, d \geq 1$ and every $\eta > 0$, there exist $r \geq s$ and a rational probability law μ on $\binom{[r]}{s}$ with the following property. Let M be any set. For every $I \in \binom{[r]}{s-1}$ give a list $A_I \subseteq M$ of size at most d. Assume*

$$A_I \cap A_J = \varnothing \quad \text{if } \max I < \min J.$$

With probability at least $1 - \eta$ over $B \sim \mu$, a map $a : M \to B$ satisfying (11) exists.

Proof. We may assume $0 < \eta < 1$ by decreasing it. Order the entries in each list. A pattern consists of the occupied slots in these ordered lists and the equivalence relation recording equality of their labels. For a fixed r the number of allowed patterns is finite. Patterns can be realized using their equivalence classes as labels. Both the hypothesis and the desired conclusion depend only on the pattern.

For a fixed pattern and an appearing label x, define an index

$$p(x) = \max\{\min I : x \in A_I\}.$$

Every occurrence $x \in A_I$ satisfies

$$\min I \leq p(x) \leq \max I. \tag{12}$$

The lower bound follows from the definition. If the upper bound failed, a list attaining the maximum would be strictly to the right of I and would share x with it. This proves the upper bound. We will round this one index in the same way for every occurrence of x.

A good set for any distribution of patterns. Choose an integer n with $sd/n < \eta/4$. By the finite Ramsey theorem take r large enough that every coloring of the s-subsets of $[r]$ in at most $(n+1)^{sd}$ colors has a homogeneous set of size $2s - 1$ [35]. This choice precedes the distribution of patterns.

Now give the allowed patterns any probability distribution ν. For an s-set $U = \{v_1 < \cdots < v_s\}$, form sd numbers, one for each $j \in [s]$ and $b \in [d]$. With $I = U \setminus \{v_j\}$, the number is

$$\Pr_{\nu}\{\text{ slot } b \text{ of } A_I \text{ is occupied by } x \text{ with } p(x) \leq v_j \}. \tag{13}$$

Here $p(x)$ is computed on the whole pattern on $[r]$. Color U by the vector of the integer parts of n times these numbers. There are at most $(n+1)^{sd}$ colors.

In a homogeneous set of size $2s - 1$, write the points in order as

$$b_1 < c_1 < b_2 < c_2 < \cdots < c_{s-1} < b_s, \qquad B = \{b_1, \ldots, b_s\}.$$

Use $c_1, \ldots, c_{s-1}$ as boundaries. Round p to b_j when $c_{j-1} < p \leq c_j$, with $c_0 = -\infty, c_s = +\infty$. For each pattern, assign every appearing label by this rounding and send absent labels anywhere in B.

Consider $I = B \setminus \{b_j\}$. If $j = 1$ or s, (12) prevents any label of A_I from rounding to b_j. For $1 < j < s$, the two points c_{j-1}, c_j occupy the same rank when inserted separately into I. The s-sets $I \cup \{c_{j-1}\}$ and $I \cup \{c_j\}$ have the same Ramsey color. For any fixed slot of A_I, subtract its probabilities in (13) at the two thresholds. Both events concern the very same slot in the same random pattern, and the first event is contained in the second. Their difference is exactly the probability that the slot is occupied and its label rounds to b_j. Homogeneity bounds this difference by $1/n$. A union bound over s deletions and at most d slots gives

$$\Pr_{\nu}\{B \text{ admits no assignment}\} \le sd/n < \eta/4. \tag{14}$$

The set B here may depend on ν.

One rational law for all patterns. Let N_0 be the number of allowed patterns. Start each at weight one. At each round, normalize the current weights to a probability law and choose a set B by (14). Double the weight of every pattern for which that set admits no assignment. The total weight increases by a factor at most $1 + \eta/4$. If a particular pattern fails in m_C of T rounds, its own weight is 2^{m_C}. It follows that

$$\frac{m_C}{T} \le \frac{\log N_0}{T \log 2} + \frac{\log(1 + \eta/4)}{\log 2}.$$

The second term is less than $\eta/2$. Choose T to make the first term less than $\eta/2$. Give each of the T chosen sets mass $1/T$, adding masses for repetitions. This rational law has failure probability less than η for every pattern. Every actual list family induces such a pattern, and assignments transfer back to its labels. □

Call a chain and B *aligned* when (11) holds for their terminal lists. Apply Lemma 6.1 with $\eta = \varepsilon$. For a chain whose separated lists are all disjoint, the conditional failure probability over B is at most ε. There are at most 4^r ordered pairs of subsets of $[r]$. Lemma 5.3 and the union bound give

$$\Pr_{\text{chain},B}\{\text{unaligned}\} \le 4^r d^2 \sigma + \varepsilon. \tag{15}$$

The order of the choices matters here. The single law μ belongs to the reduction. The lists belong to the independent set. An assignment for an aligned pair can be fixed before sampling any coefficients or rotations.

7 Soundness III. The coefficient test

Fix an aligned chain and B, and choose an assignment $a : M_r \to B$ as in (11). Put $b = \min B$. Write θ for independent uniform real rotation entries in $[0, 1)$. Let $z = (z_i)_{i \in B}$ consist of independent Haar phases, independent of all rotations and coefficients. Define

$$\begin{aligned} Y(t, \theta, z)_\kappa &= z_{a(\pi_{br}(\kappa))} + X_B(t, \theta)_\kappa \pmod 1, \qquad \kappa \in M_b, \\ H(t, \theta; z) &= F_b(Y(t, \theta, z)). \end{aligned} \tag{16}$$

The auxiliary phase depends on the terminal answer. Thus coordinates in different grids whose answers have the same terminal image will read the same auxiliary phase.

Lemma 7.1 (Conditional mean zero)**.** *For every fixed t, θ, the function $z \mapsto H(t, \theta; z)$ is L-Lipschitz, has absolute value at most one, and satisfies $\mathbb{E}_z H(t, \theta; z) = 0$.*

Proof. Each coordinate of Y reads one z_i with coefficient one. Thus Y is nonexpanding in z and $Y(t, \theta, z + v\mathbf{1}) = S_v Y(t, \theta, z)$. With t, θ fixed, Haar invariance and (6) give

$$\mathbb{E}_z H(t, \theta; z) = \int_{\mathbb{T}} \mathbb{E}_z F_b(Y(t, \theta, z + v\mathbf{1}))\, dv = \mathbb{E}_z \int_{\mathbb{T}} F_b(S_v Y(t, \theta, z))\, dv = 0.$$

The range and Lipschitz bounds follow from those of F_b. The identity also covers the all-zero coefficient vector. □

Take β, K from Lemma 2.2 at L, ε. In this section assume

$$\begin{aligned} L/m < \beta/2, \qquad \lambda K < \varepsilon, \qquad (1 - p_m)^s < \varepsilon, \\ s(2\gamma)^2/\beta^2 < \varepsilon, \qquad 2L/D < \gamma, \qquad 2Ls/P < \varepsilon. \end{aligned} \tag{17}$$

For fixed t, θ, let

$$V(t, \theta) = \mathbb{E}_z H(t, \theta; z)^2, \qquad D_i(t, \theta) = \|Q_i H(t, \theta; \cdot)\|_{L^2(z)}.$$

Applying Lemma 2.2 pointwise in t, θ gives

$$V(t, \theta) > \varepsilon \implies D_i(t, \theta) \geq \beta \text{ for some } i \in B, \tag{18}$$

$$\#\{i \in B : D_i(t, \theta) \geq \beta/2\} \leq K. \tag{19}$$

The count holds for all coefficient vectors. Its bound was fixed independently of s.

7.1 Deleting a zero coefficient

Call a coefficient *full* when it equals one and let $\mathcal{F}$ be the event that at least one coefficient is full. By independence, $\Pr(\mathcal{F}^c) = (1 - p_m)^s$.

Lemma 7.2. *Fix $t \in \mathcal{F}$ and $i \in B$ with $t_i = 0$. Then*

$$\mathbb{E}_\theta D_i(t, \theta)^2 < (2\gamma)^2.$$

Proof. Choose $j \in B$ with $t_j = 1$. Then $j \neq i$. Put $I = B \setminus \{i\}$ and $c = \min I$. In the rotation array on I, modify only block j by setting

$$\theta'_{j,\ell} = \{\theta_{j,\ell} + z_{a(\pi_{jr}(\ell))}\}, \qquad \ell \in M_j. \tag{20}$$

Here braces mean the representative in $[0,1)$ of a sum modulo one. Conditional on z, these are fixed translations of the individual uniform coordinates. The joint rotation law remains product uniform, including when several entries receive the same shift.

Let $y = X_I(t|_I, \theta')$. Composition of the projections gives for every $\kappa \in M_b$

$$y_{\pi_{bc}(\kappa)} = Y(t, \theta, z)_\kappa.$$

Indeed, the deleted term has coefficient zero. In the translated block, the added phase is $z_{a(\pi_{jr}\pi_{bj}(\kappa))} = z_{a(\pi_{br}(\kappa))}$. The integer removed by taking the representative in (20) disappears in the phase sum because $t_j = 1$. Consequently (7) gives the pointwise bound

$$|H(t, \theta; z) - G_{I,t|_I}(\theta')| = |F_b(y \circ \pi_{bc}) - F_c(y)| \le 2L/D < \gamma.$$

This also treats deletion of the minimum of B.

Take the junta $g_{I,t|_I}$ from Lemma 5.2. Preservation of the rotation law at every fixed z gives

$$\|G_{I,t|_I}(\theta') - g_{I,t|_I}(\theta')\|_{L^2(\theta,z)} < \gamma.$$

Only entries in block j can introduce dependence on z into this junta. Each such selected entry ℓ has $\pi_{jr}(\ell) \in A_I$, hence $a(\pi_{jr}(\ell)) \in I$ by alignment. Therefore $g_{I,t|_I}(\theta')$ is independent of z_i. Apply Q_i, acting on z_i in the joint product space. It kills this junta and has norm at most one. We obtain

$$\left(\mathbb{E}_\theta D_i(t, \theta)^2\right)^{1/2} = \|Q_i H\|_{L^2(\theta,z)} \le \|H - g_{I,t|_I}(\theta')\|_{L^2(\theta,z)} < 2\gamma.$$

□

For each fixed $t \in \mathcal{F}$, Markov's inequality and a union bound over its zero coefficients now apply. After averaging in t,

$$\Pr\{\mathcal{F},\ \exists i \in B :\ t_i = 0,\ D_i(t, \theta) \ge \beta\} \le s(2\gamma)^2/\beta^2 < \varepsilon. \tag{21}$$

7.2 Lowering a positive coefficient

Lemma 7.3. *Under the product laws on t, θ,*

$$\Pr\{\exists i \in B :\ t_i > 0,\ D_i(t, \theta) \ge \beta\} \le \lambda K.$$

Proof. Fix i, the array θ, and the other coefficients. At level $t_i = k/m$, denote $H(t, \theta; \cdot)$ by H_k and its i-influence by d_k. Moving down one level changes each phase coordinate by distance at most $1/m$, since each rotation is in $[0,1)$. Hence

$$|d_k - d_{k-1}| \le \|Q_i(H_k - H_{k-1})\|_2 \le L/m < \beta/2 \qquad (1 \le k \le m).$$

If $d_k \ge \beta$, the preceding level has $d_{k-1} \ge \beta/2$. Its probability is $p_{k-1} = p_k/\lambda$. It follows that

$$\sum_{k=1}^{m} p_k \mathbf{1}_{\{d_k \ge \beta\}} \le \lambda \sum_{k=0}^{m-1} p_k \mathbf{1}_{\{d_k \ge \beta/2\}} \le \lambda \sum_{k=0}^{m} p_k \mathbf{1}_{\{d_k \ge \beta/2\}}.$$

Average over the fixed data, use the union bound over i, and apply (19). The right side becomes at most $\lambda \mathbb{E}\,\#\{i : D_i(t, \theta) \ge \beta/2\} \le \lambda K$. Lowering the only full coefficient is allowed because (19) holds at the resulting vector as well. □

7.3 Absorption and the discrete law

Proposition 7.4 (Mean square on an aligned pair)**.** *Under* (17)*, for every aligned chain and* B*,*

$$\mathbb{E}_{t,\widehat{\theta}} F_b\big(X_B(t,\widehat{\theta})\big)^2 < 6\varepsilon.$$

Proof. By (18), if $V(t,\theta) > \varepsilon$, then one of three events occurs. There is no full coefficient, a zero coefficient is influential in the presence of a full one, or a positive coefficient is influential. Their probabilities are bounded by $(1-p_m)^s$, $s(2\gamma)^2/\beta^2$, and λK, respectively, by (21) and Lemma 7.3. All three are less than ε. Since $0 \le V \le 1$,

$$\mathbb{E}_{t,\theta,z} H(t,\theta;z)^2 = \mathbb{E}_{t,\theta} V(t,\theta) \le \varepsilon + \Pr\{V > \varepsilon\} < 4\varepsilon. \tag{22}$$

To remove z, fix $t \in \mathcal{F}$ and choose any one full block j. Translate its rotations by (20), now retaining all blocks. The same coordinate computation gives $X_B(t,\theta') = Y(t,\theta,z)$. The conditional rotation law is unchanged for each z, so

$$\mathbb{E}_{\theta,z} H(t,\theta;z)^2 = \mathbb{E}_{\theta} F_b\big(X_B(t,\theta)\big)^2 \qquad (t \in \mathcal{F}).$$

Multiply by the probability of each such t and sum. The resulting contribution to the unshifted square is below 4ε by (22). On $\mathcal{F}^c$ the square is at most one, which adds less than ε.

Finally couple the discrete rotations by $\widehat{\theta}_{j,\ell} = \lfloor P\theta_{j,\ell} \rfloor / P$. They have exactly the independent uniform law of the construction. Each phase coordinate in (4) moves by at most s/P. Since $|F_b| \le 1$, the difference of the two squared function values is at most

$$2L\, d_\infty\big(X_B(t,\theta), X_B(t,\widehat{\theta})\big) \le 2Ls/P < \varepsilon.$$

Adding these contributions proves the proposition. □

8 Parameters and proof of the configuration theorem

First fix $L = 2/\rho$ and a rational $\varepsilon > 0$ such that

$$(6h^2 + 2)\varepsilon < \min\{\delta, 1\}.$$

The following order places the source alphabets after the list size and the number of layers.

1. Take β, K from Lemma 2.2 for L, ε. Choose m with $L/m < \beta/2$, then rational $0 < \lambda < 1$ with $\lambda K < \varepsilon$. Equation (3) now fixes $p_m > 0$.

2. Choose $s \geq 2$ with $(1 - p_m)^s < \varepsilon$. Choose $\gamma > 0$ with $s(2\gamma)^2/\beta^2 < \varepsilon$. Take a positive multiple P of the offset denominators large enough that

$$2L/(mP) < \gamma, \qquad 2Ls/P < \varepsilon.$$

   Set $D = mP$. This gives all of (17).

3. Fix the integer d in (9). Apply Lemma 6.1 with $s, d, \eta = \varepsilon$ to obtain r, μ. Choose rational $0 < \sigma < 1$ such that $4^r d^2 \sigma < \varepsilon$. Only now take the alphabets of Theorem 2.1.

4. Use any $q \geq h$ with $(h - 1)/q < \delta$ for the fallback.

Every object used to construct the output is a fixed integer or a finite rational table, apart from the input Label Cover instance. In particular, the law μ is independent of the terminal answer set and of the independent set later considered. The real functions, conditional expectations, and decoding maps occur only in the proof. For each fixed real $\delta > 0$ the strict inequalities above can be met with finite construction data.

Proof of Theorem 1.1. Apply Theorem 2.1 at the chosen σ, then use Section 3. Proposition 3.1 gives a deterministic polynomial-time reduction in the binary input length. In the satisfiable case, Lemma 4.1 gives the phase map and excludes the fallback.

Suppose the input is unsatisfiable. The fallback has independence number $h - 1$, hence independence ratio less than δ. On the sampling branch, fix a nonempty independent set $\mathcal{A}$. For every output vertex v, with F the function on its location's grid, Lemma 5.1 gives

$$\mathbf{1}_{\mathcal{A}}(v) \leq h^2 F(x_v)^2. \tag{23}$$

Indeed, a selected vertex has grid value at least $1/h$. At an unselected vertex the inequality holds also when a different copy at the same location was selected.

On an aligned pair, Proposition 7.4 bounds the expected right side of (23) by $6h^2\varepsilon$. On an unaligned pair bound the indicator itself by one. By (15) and the choice of σ, the probability of the latter event is less than 2ε. The vertex law is uniform by Proposition 3.1. Thus

$$\frac{|\mathcal{A}|}{|V(H)|} = \Pr\{v \in \mathcal{A}\} < 6h^2\varepsilon + 2\varepsilon < \delta.$$

The empty independent set satisfies the strict inequality as well. Taking the maximum over independent sets proves soundness. □

9 Proofs of the coloring corollaries

9.1 Graphs

Proof of Corollary 1.2. In Theorem 1.1 take $h = 2$, offsets $0, 1/2$, and $\rho = 1/(8g)$. On every satisfiable edge, subtracting the two phases introduces error at most $2\rho = 1/(4g)$ relative to the half-turn. Lemma 4.2 gives a homomorphism to C_g. The unsatisfiable bound is the density assertion of the theorem.

Every odd cycle maps to K_3, hence to K_c for $c \geq 3$. Take $0 < \delta < \min\{1/3, 1/c\}$. A proper c-coloring would have an independent class of density at least $1/c$. This proves the odd-cycle versus clique gap and hardness of finding such a coloring under the odd-cycle promise, by checking the output of a proposed polynomial-time algorithm.

For general $A \to B$ as in the corollary, some component of A contains an odd cycle C_g. Choose $0 < \delta < \min\{1/3, 1/|V(B)|\}$. In the satisfiable case compose $X \to C_g \to A$. A homomorphism $X \to B$ in the unsatisfiable case would have an independent fiber of density at least $1/|V(B)|$, since B is loopless. This is impossible.

For completeness, the other finite undirected templates give polynomial-time problems. A loop at the target permits constant acceptance. A loop in the source forces a target loop. For a bipartite source with an edge, $K_2 \to A \to B$. Testing bipartiteness separates the promises. For a nonempty edgeless source, accept the edgeless inputs. For the empty source, accept only the empty input. □

Remark 9.1 (Circular chromatic number). The circular clique $K_{p/q}$ joins residues modulo p whose cyclic distance is at least q, where $p \geq 2q$. For $p = 2k + 1, q = k$ its steps are precisely $\pm k$, so it is isomorphic to C_{2k+1}. Given $\xi > 0$, take $1/k \leq \xi$. Corollary 1.2 then gives hardness of distinguishing $\chi_c(X) \leq 2 + \xi$ from independence ratio below any fixed $0 < \delta < \min\{1/3, 1/(2 + \xi)\}$. Indeed, q consecutive residues are independent in $K_{p/q}$, so a uniformly translated such interval pulls back under any circular coloring to an independent set of expected density q/p. Taking the supremum over admissible circular colorings gives independence ratio at least $1/\chi_c(X)$ for graphs with an edge. Edgeless graphs have independence ratio one.

9.2 Strong, rainbow, and linearly ordered colorings

Proof of Corollary 1.3. Use the regular configuration $u_j = j/h$ and rational $0 < \rho < 1/(2h(h + 1))$. Lemma 4.4 reads both colorings from the map supplied by Theorem 1.1. For $k = h \geq 3$, take $\delta < 1/c$. The rainbow colors can be grouped into two nonempty parts to give a proper two-coloring, so the search problem on the simultaneous promise has a solution for every $c \geq 2$. A proper c-coloring on an unsatisfiable input would have an independent color class of density at least $1/c$. Testing a returned coloring on every edge therefore distinguishes the cases. Hardness on instances satisfying both promises gives hardness under each promise separately. For $k = 2, c \geq 3$, strong three-colorability is ordinary three-colorability. This is the case $g = 3$ of Corollary 1.2. □

Proof of Corollary 1.5. Take $h = 3$ in Corollary 1.3 and $\delta < 1/d$. In the satisfiable case, Lemma 4.6 gives an LO_3-coloring. An increasing inclusion into the c-element palette preserves unique maxima. In the unsatisfiable case, each class of an LO_d-coloring would be independent, as a monochromatic triple has three maxima. Its largest class would have density at least $1/d$. □

9.3 Reconfiguration and monochromatic cliques

Proof of Corollary 1.4. For the first assertion choose the regular h-gon and the radius of Lemma 4.5 for R. Apply Theorem 1.1 with $\delta < 1/d$. Make the output a relational structure by including every

ordering of every hyperedge as a tuple of its h-ary relation. In the satisfiable case give vertex v the color $\chi(\varphi(v))$. Lemma 4.5 puts the edge tuple in R in its witnessing order. Symmetry gives membership in every order. In the unsatisfiable case any homomorphism to $\mathrm{NAE}_h^{(d)}$ would give a proper d-coloring, contrary to the density bound.

For the second assertion, use the finite relations in the monochromatic-clique equivalence of Banakh, Barsukov, and Nakajima, Lemma 6 [6]. Its source $R_{k:\ell}^{(c)}$ consists of the length-ℓ words on c colors in which each multiplicity is at most $k-1$. Its target $R_{\ell:\ell}^{(d)}$ is $\mathrm{NAE}_\ell^{(d)}$. The palette map from the containment hypothesis sends source words to target words. Indeed, a legal word colors an ℓ-clique without a monochromatic k-clique, so its recoloring cannot make the whole clique monochromatic. Containment implies $k \le \ell$, since otherwise a monochromatic ℓ-clique would itself contradict containment.

Suppose first that $\ell < c(k-1)$. Then $k \ge 2$ and $\ell \ge 2$. Give each color $k-1$ distinguishable slots. The legal words are precisely the projections to colors of injections of the ℓ positions into the $c(k-1)$ slots. There is at least one injection and at least one spare slot. To reach a prescribed injection, fix the positions in turn. If another position occupies the desired slot, move that occupant to a vacant slot. Move the current position into its desired slot if necessary. Every previously fixed position occupies its own distinct target, so these moves preserve it. Projecting the sequence to colors and omitting repetitions gives a single-coordinate path of legal words. The relation is symmetric and reconfigurable. Thus the first assertion gives the hard finite PCSP used in the proof of Theorem 23 of Banakh et al. [6].

For these particular hard instances one can also see the transfer to graphs directly. This verifies the direction of the equivalence needed for hardness without a sparsification theorem. Use the circle coloring for $R_{k:\ell}^{(c)}$ just obtained from Lemma 4.5, with tolerance ρ_*. Run the configuration reduction with regular ℓ-gon offsets and a positive rational radius ρ satisfying $2\rho < \rho_*$ and $6\rho < 1/\ell$. On the same vertices form the graph joining two vertices whenever they lie in a common hyperedge.

In the satisfiable case consider any k-clique in this graph. Fix the phase t of one vertex. Each other phase is within 2ρ of $t + j/\ell$ for some nonzero $j \bmod \ell$, by applying the phase condition to a hyperedge shared with the fixed vertex. These assigned indices are distinct. Two vertices assigned the same index would be at distance at most 4ρ, whereas sharing any hyperedge gives distance at least $1/\ell - 2\rho > 4\rho$. Include the fixed vertex at index zero and fill unused indices by their ideal phases. Lemma 4.5 sends the resulting ℓ-tuple into $R_{k:\ell}^{(c)}$. Thus the k-clique cannot be monochromatic. The coloring used on the graph is the single map $\chi \circ \varphi$. In the unsatisfiable case, every hyperedge is an ℓ-clique in the graph. A d-coloring avoiding monochromatic ℓ-cliques would therefore be a proper d-coloring of the hypergraph. Choosing $\delta < 1/d$ excludes it. The graph construction is deterministic and polynomial.

Finally suppose $\ell \ge c(k-1)$. We spell out the affine integer argument of Banakh et al., Theorem 22 [6]. In their relational formulation an input may contain loops, and a looped vertex carries a monochromatic image of every clique; reject such inputs first. On a simple graph, enumerate all ℓ-cliques, which takes polynomial time for fixed ℓ. A source coloring gives a legal source word on each such clique. If $\ell > c(k-1)$, no such word exists. Accept exactly when there are no ℓ-cliques. The same test works when the target relation is empty, since a graph avoiding all ℓ-cliques has a target coloring and containment excludes any ℓ-clique in a yes instance.

It remains to treat $\ell = c(k-1)$ with nonempty target relation. Then $d \ge 2$ and $0 < k-1 < \ell$, using $k \le \ell$. Every source word uses each color exactly $k-1$ times. Introduce one integer variable per vertex and impose

$$\sum_{v \in Q} z_v = k-1 \qquad \text{for every } \ell\text{-clique } Q.$$

A source coloring gives a solution by taking the indicator of one fixed color. In any integer solution, each clique has a positive entry since the sum is positive. It also has a nonpositive entry since the sum is strictly below ℓ. The two sign classes give a target coloring. Thus testing integer feasibility separates the promises. This is a system of linear equations over the integers, whose feasibility is decidable in polynomial time by integer normal form. This proves the tractable side, including equality. □

Use of AI

The proof presented in this paper was developed through an interactive dialogue with GPT-6.2 Astra (OpenAI) and Claude Opus 5.5 (Anthropic). The author used them to explore multiple possible proof strategies and directions, and the AI generated the initial draft under the author's guidance. The author then critically examined the generated proof, verified each step for correctness and rigor, and revised it where necessary. The author takes full responsibility for the final content.

References

[1] P. Austrin, A. Bhangale, and A. Potukuchi. Improved inapproximability of rainbow coloring. arXiv:1810.02784, 2018.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM* 45 (1998), 501–555.

[3] T. Austin. On the failure of concentration for the ℓ_∞-ball. *Israel J. Math.* 211 (2016), 221–238. arXiv version 4.

[4] S. Avvakumov, M. Filakovský, J. Opršal, G. Tasinato, and U. Wagner. Hardness of 4-colouring G-colourable graphs. STOC 2025.

[5] L. Barto, J. Bulín, A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM* 68 (2021), article 28.

[6] D. Banakh, A. Barsukov, and T.-V. Nakajima. Towards infinite PCSP: a dichotomy for monochromatic cliques. LICS 2026, LIPIcs 380, article 13. Full version.

[7] J. Brakensiek and V. Guruswami. New hardness results for graph and hypergraph colorings. CCC 2016, LIPIcs 50, article 14.

[8] J. Brakensiek and V. Guruswami. Promise constraint satisfaction: algebraic structure and a symmetric Boolean dichotomy. ECCC TR16-183, 2016, revised 2017 and 2021. Conference version, *Promise constraint satisfaction: structure theory and a symmetric Boolean dichotomy*, SODA 2018.

[9] J. Brakensiek and V. Guruswami. The quest for strong inapproximability results with perfect completeness. APPROX/RANDOM 2017, LIPIcs 81, article 4.

[10] M. Braverman, S. Khot, N. Lifshitz, and D. Minzer. An invariance principle for the multi-slice, with applications. Full version, arXiv:2110.10725v2, 2025.

[11] J. Bulín, A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. STOC 2019.

[12] Open Mathematics Courses. Large independent sets in three-colorable graphs: Label Cover and the graph of the reduction. Exercise 6.5, October 13, 2026.

[13] I. Dinur, S. Khot, G. Kindler, D. Minzer, and M. Safra. Towards a proof of the 2-to-1 games conjecture? *Theory of Computing* 21 (2025), article 11.

[14] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.* 39 (2009), 843–873.

[15] I. Dinur, O. Regev, and C. Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica* 25 (2005), 519–535.

[16] I. Dinur and D. Steurer. Analytical approach to parallel repetition. STOC 2014, 624–633.

[17] M. Filakovský, T.-V. Nakajima, J. Opršal, G. Tasinato, and U. Wagner. Hardness of linearly ordered 4-colouring of 3-colourable 3-uniform hypergraphs. STACS 2024, LIPIcs 289, article 34.

[18] Y. Fei, D. Minzer, and S. Wang. On the hardness of 4-to-1 games with perfect completeness. ECCC TR26-179, September 14, 2026. arXiv:2610.12378v2.

[19] OpenAI. Perfect completeness for 2-to-1 games. OpenAI Math Release, item 105, September 23, 2026.

[20] M. R. Garey and D. S. Johnson. The complexity of near-optimal graph coloring. *J. ACM* 23 (1976), 43–49.

[21] V. Guruswami and S. Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.* 18 (2004), 30–40.

[22] V. Guruswami and E. Lee. Strong inapproximability results on balanced rainbow-colorable hypergraphs. ECCC TR14-043, 2014.

[23] V. Guruswami and S. Sandeep. d-to-1 hardness of coloring 3-colorable graphs with $O(1)$ colors. ICALP 2020, article 62.

[24] V. Guruswami and S. Sandeep. Rainbow coloring hardness via low sensitivity polymorphisms. *SIAM J. Discrete Math.* 34 (2020), 520–537.

[25] P. Hell and J. Nešetřil. On the complexity of H-coloring. *J. Combin. Theory Ser. B* 48 (1990), 92–110.

[26] OpenAI. Hardness of finding large independent sets in three-colorable graphs. OpenAI Math Release, item 106, September 24, 2026.

[27] S. Khanna, N. Linial, and S. Safra. On the hardness of approximating the chromatic number. *Combinatorica* 20 (2000), 393–415.

[28] A. Krokhin and J. Opršal. The complexity of 3-colouring H-colourable graphs. FOCS 2019, 1227–1239.

[29] A. Krokhin, J. Opršal, M. Wrochna, and S. Živný. Topology and adjunction in promise constraint satisfaction. *SIAM J. Comput.* 52 (2023), 38–79.

[30] A. Krokhin and D. Vagnozzi. Approximating 1-in-3 SAT by linearly ordered hypergraph 3-colouring is NP-hard. ICALP 2026, LIPIcs 374, article 184.

[31] OpenAI. Hardness of coloring three-colorable graphs. OpenAI Math Release, Lean formalization scope.

[32] E. J. McShane. Extension of range of functions. *Bull. Amer. Math. Soc.* 40 (1934), 837–842.

[33] T.-V. Nakajima, Z. Verwimp, M. Wrochna, and S. Živný. Complexity of approximate conflict-free, linearly-ordered, and nonmonochromatic hypergraph colourings. arXiv:2501.12062, 2025.

[34] T.-V. Nakajima and S. Živný. Linearly ordered colourings of hypergraphs. *ACM Trans. Comput. Theory* 14, nos. 3–4 (2023).

[35] F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.* (2) 30 (1930), 264–286.

[36] R. Raz. A parallel repetition theorem. *SIAM J. Comput.* 27 (1998), 763–803.

[37] M. Wrochna and S. Živný. Improved hardness for H-colourings of G-colourable graphs. SODA 2020, 1426–1435.