hexagon:2611.01095v1

Approximate Coloring Is a Promise CSP

Abstract

The contribution contains:

1. A direct reduction from 3-SAT to approximate coloring without using the PCP theorem. For each fixed $n \geq 3$, the task is to find an n-coloring of a graph promised to be 3-colorable.
2. A new notion of minion homomorphism that explains this reduction. The polymorphism minions of 3-SAT and approximate coloring admit such maps in both directions. Soundness follows from an extension of the Barto–Kozik baby PCP theorem.
3. A combination of existing and new reductions that prove a dichotomy for promise CSPs on pairs of finite digraphs without sources or sinks.
4. An unconditional dichotomy for finite Boolean promise CSPs whose polymorphisms are unate. The tractable cases admit block-symmetric polymorphisms on two arbitrarily large blocks. The hardness proof uses an injective extension of the new baby PCP theorem.
5. NP-hardness for seven families of promise CSPs studied in the literature.

Contributors

Joshua Brakensiek (UC Berkeley), Venkatesan Guruswami (UC Berkeley), Marcin Kozik (Jagiellonian University), and Xuandi Ren (UC Berkeley)

AI disclosure

The vast majority of the proofs in this paper were devised by ChatGPT-6 Astra or in close collaboration with it. We carefully (and painfully) simplified and verified some of these proofs before presenting them here. We verified others without simplifying them. Some still await both verification and simplification. Future versions of this paper will clarify the presentation.

Parallel work

In parallel work, Yuwen Zhou [24, Corollary 1.2] proved that for every fixed odd cycle C and fixed $0 < \delta < 1/3$, it is NP-hard to distinguish graphs mapping to C from graphs with independence ratio below δ.

Introduction

For a fixed integer $n \geq 3$, *approximate graph coloring* asks for an n-coloring of a graph promised to be 3-colorable. We give a direct LOGSPACE reduction from 3-SAT to this problem, together with a LOGSPACE search decoder. Thus a satisfying assignment produces a 3-coloring of the output graph, and any n-coloring of that graph can be decoded into a satisfying assignment. The parameters of the reduction depend on n, which remains fixed.

Fei, Minzer and Wang proved NP-hardness of this decision problem for every fixed $n \geq 3$ [2, Corollary 1.7]. The starting point is the coordinate-list argument in OpenAI's proof of hardness for finding large independent sets in three-colorable graphs [1]. We seek an algebraic condition that

captures the part of that argument needed for coloring and supplies a general reduction theorem for promise constraint satisfaction problems. The conclusion here is the reduction to PCSP(K_3, K_n); we do not recover the stronger independent-set density conclusion of [1]. Brakensiek and Guruswami conjectured decision hardness for every pair of finite loopless undirected graphs $G \to H$ with G nonbipartite [23, Conjecture 1.2].

A chain homomorphism assigns a bounded list of operations to each short minor chain. On a chain with one more position, the lists of its faces have a common value after transport to the last coordinate set. The graph-coloring application satisfies this condition on every chain. For Boolean applications we also assign bounded coordinate marks to each operation. We require the same intersection on every injective chain, where successive maps preserve distinctness on all marks accumulated so far. This gives an injective chain homomorphism.

The reduction evaluates the face lists on local solutions. Its soundness requires a global solution even when some restrictions identify marked local assignments. Two distinct assignments can be kept distinct by retaining one variable on which they differ. The difficulty is that the upper scope determines these variables, whereas the lower tests must already fit inside that scope. At each lower level of our injective baby PCP theorem, an edge-coloring records which lists of tuples agree on a small set of variables. If one level has no pivot, ordinary baby PCP recovers a solution on that level. Otherwise we fix one pivot per level. Sampling the prospective lower scopes in advance lets us retain the variables needed for injectivity while avoiding all earlier lists. When lists depend on whole subchains, the same matrix argument is applied at the splits of a binary tree. Ramsey's theorem then forces a pair of matching faces. The instance map uses logspace. Reconstruction is polynomial-time in general and uses logspace in the ordinary case.

Our hardness source pairs a tight cycle with the relation of nonconstant p-tuples of colors. The continuous proof represents each polymorphism by an equivariant Lipschitz map from a power of the circle to a sphere. For graphs, taking a minor changes a value by at most a right angle. For odd primes, it preserves the minimizing phases in each color column. These conditions keep the convex hull of the row values on a common input away from zero.

To choose a face list, turn one terminal coordinate by $1/p$ and measure the mean change in the row values. Only boundedly many coordinates have appreciable effect. If the face lists had empty intersection, assign every terminal coordinate to a face missing it. The cyclic Borsuk–Ulam argument forces one such block to have appreciable effect, and a Lipschitz estimate then finds a selected coordinate in that block. The combinatorial proof works directly with a finite clock and uses cyclic covering to force a color on every rotation of a cell.

The source maps to a cyclic relation exactly when the corresponding modular equations on ordered faces are inconsistent. For graphs these are the bipartiteness equations. This gives the approximate graph-coloring dichotomy, the smooth-digraph dichotomy through auxiliary constructions, and seven hypergraph-coloring applications. In each hypergraph application, a short periodic word supplies the required tight-cycle homomorphism. For Boolean templates whose polymorphisms are unate, the tractable case has block-symmetric polymorphisms on blocks of sizes $t, t+1$ for every t. A missing width gives a bounded direction certificate at every polymorphism. An injective chain transports these certificates and allows the increasing operations to supply the tight-cycle obstruction. This gives a PCP-free hardness proof covering both monotone polymorphisms and signed linear thresholds.

Section 1 places the algebraic condition among existing reduction criteria and proves the reduction theorem in Section 1.7. Section 2 states the tight-cycle source, proves the sandwich characterization, and derives the applications. Appendix A proves the source theorem. The following appendices contain the auxiliary hypergraph results, the smooth-digraph constructions, and the comparison, composition, and separation results.

Preliminaries

We write $[m] = \{1, \dots, m\}$ and $\mathcal{P}_k(X) = \{U \subseteq X : |U| = k\}$. A *relational signature* is a collection of relation symbols, each with a specified *arity* (number of arguments). All relational signatures in this paper are finite, and all relation symbols have positive arity. Unless stated otherwise, domains are finite and nonempty.

For a tuple x indexed by I and $J \subseteq I$, write $x|_J$ for restriction to J. For a set S of such tuples, its *coordinate projection* onto J is $S|_J = \{x|_J : x \in S\}$.

Definition (Relational structures and promise problems)**.** A *relational structure* $\mathbf{A}$ consists of a set A and a relation $R^{\mathbf{A}} \subseteq A^r$ for each r-ary relation symbol R. A *homomorphism* $u : \mathbf{A} \to \mathbf{B}$ is a map $u : A \to B$ that sends every tuple in $R^{\mathbf{A}}$ to a tuple in $R^{\mathbf{B}}$, for every R. We write $\mathbf{A} \to \mathbf{B}$ when such a map exists.

A *promise template* is a pair $(\mathbf{A}, \mathbf{B})$ of relational structures in the same relational signature with $\mathbf{A} \to \mathbf{B}$. An input to $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is a finite relational structure $\mathbf{X}$ in the same signature, represented by its domain and lists of relation tuples; the template is fixed. The decision problem asks one to distinguish $\mathbf{X} \to \mathbf{A}$ from $\mathbf{X} \not\to \mathbf{B}$, with no requirement on inputs satisfying neither condition. The search problem asks for a homomorphism $\mathbf{X} \to \mathbf{B}$ under the promise $\mathbf{X} \to \mathbf{A}$. We write $\mathrm{CSP}(\mathbf{A}) = \mathrm{PCSP}(\mathbf{A}, \mathbf{A})$.

We identify an input relational structure with a *instance*: its elements are variables, and a tuple in one of its relations is a *constraint*. The *constraint scope* of a constraint is the set of variables occurring in its tuple. For an instance Φ over a value set B, an *assignment* on U is a map $U \to B$. It is a *local solution* if it satisfies every constraint whose constraint scope is contained in U; a *global solution* is a solution on all variables. In particular, a local $\mathbf{A}$-solution on $U \subseteq X$ is a homomorphism $\mathbf{X}[U] \to \mathbf{A}$, where the induced relational structure $\mathbf{X}[U]$ retains exactly the relation tuples contained in U. A *list* is a finite set of candidate operations or assignments; it is unordered, and its *list bound* is an upper bound on its cardinality. A *list assignment* is a rule assigning a list to each operation or minor chain in its domain. When enumerating a list, we specify any repeated entries or extra padding symbols explicitly. A polynomial-time *decision reduction* from $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ to $\mathrm{PCSP}(\mathbf{C}, \mathbf{D})$ is a polynomial-time map T on instances such that

$$\mathbf{X} \to \mathbf{A} \implies T(\mathbf{X}) \to \mathbf{C}, \qquad T(\mathbf{X}) \to \mathbf{D} \implies \mathbf{X} \to \mathbf{B}.$$

A *LOGSPACE* decision reduction is such an instance map computable by a deterministic logspace transducer: on an input of length N, its work tapes use $O(\log N)$ workspace, with read-only input and write-only output. A *search reduction* additionally provides a polynomial-time procedure that constructs $\mathbf{X} \to \mathbf{B}$ from a given $T(\mathbf{X}) \to \mathbf{D}$. This procedure is the *search decoder*. The *complete graph* K_c has vertex set $[c]$ and edge relation $\{(i, j) : i \neq j\}$. Thus a homomorphism to K_c is a proper c-coloring.

Definition (Minions)**.** A *minion* $\mathscr{M}$ consists of sets $\mathscr{M}(I)$, indexed by finite nonempty sets I and not all empty, together with maps $f \mapsto f^\pi$ from $\mathscr{M}(I)$ to $\mathscr{M}(J)$ for every map $\pi : I \to J$, satisfying

$$f^{\mathrm{id}_I} = f, \qquad (f^\pi)^\rho = f^{\rho \circ \pi}.$$

An element of $\mathscr{M}(I)$ is called an I-ary *operation*. Its *coordinate set* is I, its arity is $|I|$, and a map $\pi : I \to J$ is a *coordinate map*. The operation f^π is its *minor* along π. A *surjective minor* is one whose coordinate map is surjective.

For functions $f : A^I \to B$, the minor is defined by $f^\pi(x) = f(x \circ \pi)$ for $x \in A^J$. A *function minion* is a nonempty family of such functions closed under all minors. Its operations are not required to be closed under composition.

The *power* $\mathbf{A}^I$ has domain A^I and coordinatewise relations. A *polymorphism* of $(\mathbf{A}, \mathbf{B})$ with coordinate set I is a homomorphism $\mathbf{A}^I \to \mathbf{B}$. The *polymorphism minion* of a promise template is

$$\mathrm{Pol}(\mathbf{A}, \mathbf{B})(I) = \{f : \mathbf{A}^I \to \mathbf{B} : f \text{ is a homomorphism}\}.$$

We abbreviate $\mathrm{Pol}(\mathbf{A}, \mathbf{A})$ to $\mathrm{Pol}(\mathbf{A})$. A *projection* is the function $e_i(x) = x_i$. The *projection minion* $\mathscr{P}$ has one element e_i for each $i \in I$, with $e_i^\pi = e_{\pi(i)}$; it may be realized by coordinate projections on $\{0, 1\}$. Accordingly, a subset of $\mathscr{P}(I)$ will also be regarded as a subset of I.

For a function $f : A^I \to B$, $i \in I$ is an *essential coordinate* if there are two inputs differing only at i on which f has different values. A coordinate that is not an essential coordinate is a *dummy coordinate*. The minors of a family of functions form the minion generated by that family.

Definition (Minor chains and transport)**.** A *minor chain* of length r is a sequence

$$C = (f_0 \xrightarrow{\pi_{01}} f_1 \xrightarrow{\pi_{12}} \cdots \xrightarrow{\pi_{r-1,r}} f_r), \qquad f_i \in \mathscr{M}(I_i),$$

where $f_{i+1} = f_i^{\pi_{i,i+1}}$. For $i < j$ let π_{ij} be the composite of the intervening maps, and put $\pi_{ii} = \mathrm{id}_{I_i}$. Repetitions and identity maps are allowed; repeated operations still occupy distinct positions. A *subchain* retains a specified set of positions, together with their operations and the corresponding composite maps. A minor chain with h operations has h positions and $h - 1$ arrows. A *face* is a subchain obtained by deleting one position. More generally, an h-face retains exactly h positions. We also write $E \subseteq C$ for a selection of positions in a minor chain and the resulting subchain; then $\min E$ and $\max E$ refer to position indices. For a set, an h-face likewise means an h-element subset.

For $L \subseteq \mathscr{N}(I)$ put $L^\pi = \{g^\pi : g \in L\}$. If a subchain F of C ends at position j and $L \subseteq \mathscr{N}(I_j)$, its *transport* to the end of C is $L \downarrow_C = L^{\pi_{jr}}$. A *surjective minor chain* is a minor chain in which all coordinate maps are surjections. A *set chain* is a sequence of distinct sets ordered by strict inclusion; we order its sets from largest to smallest. Subchains and faces of set chains are defined by retaining or deleting positions in the same way.

Definition (Injective chains)**.** Given a sequence $I_0 \to \cdots \to I_r$ of coordinate sets and maps, with specified subsets $S_i \subseteq I_i$, call it an *injective chain* for these subsets if, for each $i < r$, the next map $\pi_{i,i+1}$ is injective on

$$\bigcup_{j=0}^{i} \pi_{ji}(S_j).$$

Thus coordinates from different rows that are distinct at position i must also stay distinct at position $i + 1$. Every subchain is injective for its retained subsets: their images belong to the tested union at every intervening position.

For a minion, a choice of *coordinate marks* assigns a set $S(f) \subseteq I$ to each $f \in \mathscr{M}(I)$, of size at most one fixed constant and possibly empty. On a minor chain use $S_i = S(f_i)$. The *injective variant* of a minor-preservation condition asks for one choice of coordinate marks and tests the same condition on all injective chains. For a single arrow $f \xrightarrow{\pi} g$ this means testing precisely the maps π injective on $S(f)$. Empty marks recover the test on all chains. One may independently require each coordinate map to be surjective onto its full codomain.

All bounds on list size, lengths of minor chains and set chains, and the sizes of local variable sets are constants independent of arity and instance size. In statements about reductions we use the notation

$$\mathscr{M} = \mathrm{Pol}(\mathbf{C}, \mathbf{D}), \qquad \mathscr{N} = \mathrm{Pol}(\mathbf{A}, \mathbf{B}), \qquad \mathrm{PCSP}(\mathbf{A}, \mathbf{B}) \longrightarrow \mathrm{PCSP}(\mathbf{C}, \mathbf{D}). \tag{1}$$

We assume that the target problem has a NO instance, that is, some finite relational structure has no homomorphism to $\mathbf{D}$. Statements involving arbitrary minions do not implicitly assert that those minions are polymorphism minions of finite promise templates.

Contents

1 Algebraic reductions and the chain condition **7**
- 1.1 Minion homomorphisms and gadget reductions . . . 7
- 1.2 Consistency reductions . . . 8
- 1.3 Exact set-valued homomorphisms . . . 9
- 1.4 The d-homomorphism condition . . . 10
- 1.5 Barto–Kozik homomorphisms . . . 10
- 1.6 Injective baby PCP . . . 10
- 1.7 Chain homomorphisms on injective chains . . . 14

2 Tight-cycle hardness and applications **19**
- 2.1 Approximate graph colouring . . . 21
- 2.2 Digraphs with no sources and no sinks . . . 22
- 2.3 Unate Boolean polymorphisms . . . 23
- 2.4 Rainbow colouring . . . 25
- 2.5 Strong colouring with a spare colour . . . 26
- 2.6 At least a prescribed number of colours . . . 26
- 2.7 Sparse cyclic Boolean relations . . . 26
- 2.8 Linearly ordered triples . . . 27
- 2.9 Capacity-constrained colourings . . . 27
- 2.10 Discrepancy one . . . 28

A Two proofs of Theorem 2.1 **29**
- A.1 Continuous proof . . . 29
- A.2 Combinatorial proof . . . 36

B Auxiliary results for hypergraph colouring **42**

C Constructions for smooth digraphs **44**
- C.1 Transferring hardness through arc digraphs . . . 44
- C.2 Reducing the algebraic length . . . 44
- C.3 Periodic walks and the auxiliary digraph . . . 45
- C.4 A cyclic relation gadget . . . 46

D Proofs for the reduction criteria of Section 1 **48**
- D.1 Exact lists and intersecting lists . . . 48
- D.2 Barto–Kozik homomorphisms . . . 48
- D.3 A separation at the Barto–Kozik condition . . . 49
- D.4 Consistency reductions . . . 50
- D.5 Logspace construction and decoding . . . 52

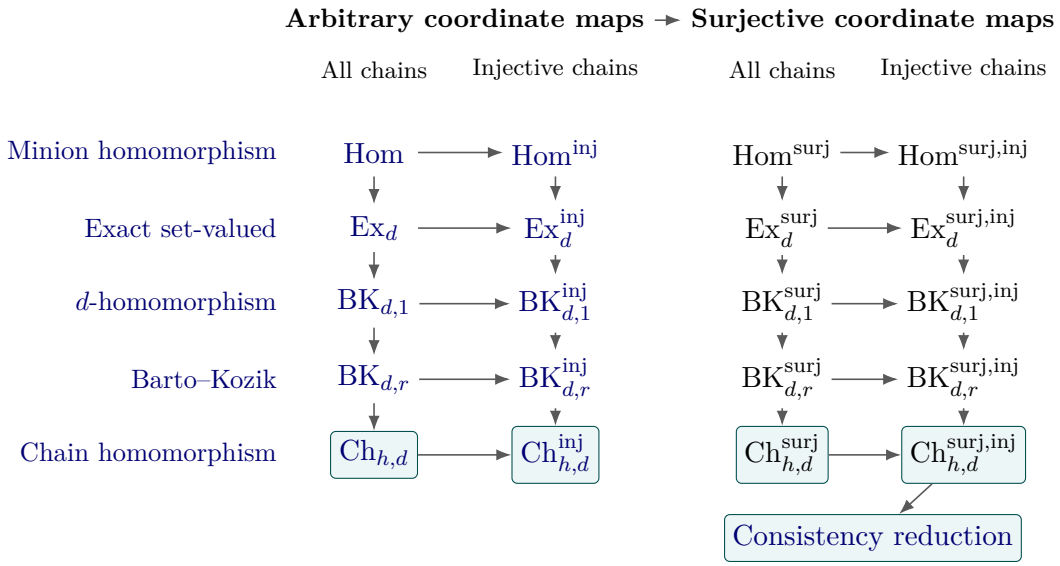
E Surjective variants **55**
- E.1 Surjective-minor homomorphisms and arc-consistency construction . . . 55
- E.2 Surjective Barto–Kozik homomorphisms . . . 56
- E.3 Surjective chain homomorphisms . . . 56

F Comparison and composition of chain homomorphisms **58**

G Proofs for unate Boolean polymorphisms **60**
- G.1 Cancellation and block symmetry . . . 60
- G.2 Rotation for increasing operations . . . 62
- G.3 Direction certificates on an injective chain . . . 63

1 Algebraic reductions and the chain condition

In this section we prove the reduction theorem for injective chain homomorphisms between finite promise templates. Its combinatorial input reconstructs a solution from bounded lists of partial assignments even when restrictions may identify selected local assignments and only the injective chains are tested.



Each entry asserts existence from $\mathscr{M}$ to $\mathscr{N}$ for some finite parameters, in the reduction direction (1). Here Ex_d denotes an exact d-set-valued homomorphism. Each downward arrow weakens the condition on the lists. Within each panel the horizontal arrows allow bounded coordinate marks; the arrow between panel headings restricts every corresponding test to surjective maps. These restrictions are independent. The final arrow is Theorem E.8.

The first three downward implications in the all-chains column for arbitrary maps are strict (Examples D.1, D.2, and D.4). For minion homomorphisms, restriction to surjections can also be strict (Example E.4). At the Barto–Kozik row, the minion ST separates the test on all chains from its injective variant (Example D.5). Even the bottom-right condition is strictly stronger than a consistency reduction for finite templates (Example D.10).

Ordinary minion homomorphisms characterize gadget reductions (Section 1.1), surjective-minor homomorphisms characterize the arc-consistency construction (Theorem E.3), and consistency reductions characterize Datalog-with-unions reductions (Theorem D.8).

1.1 Minion homomorphisms and gadget reductions

Minion homomorphisms are the minor-preserving maps used by Barto, Bulín, Krokhin, and Opršal [3] to characterize gadget reductions between finite promise templates.

Definition 1.1 (Minion homomorphism)**.** A *minion homomorphism* $\xi : \mathscr{M} \to \mathscr{N}$ is a family of maps $\xi_I : \mathscr{M}(I) \to \mathscr{N}(I)$ satisfying

$$\xi_J(f^\pi) = \xi_I(f)^\pi \qquad (f \in \mathscr{M}(I),\ \pi : I \to J).$$

An *injective minion homomorphism* chooses bounded coordinate marks and requires this equality whenever π is injective on the marks of f. We usually omit the coordinate-set subscripts. We write $\mathrm{Hom}(\mathscr{M}, \mathscr{N})$ for the assertion that a minion homomorphism exists.

Precomposition of a list assignment $\Xi : \mathscr{M} \to \mathscr{N}$ by a minion homomorphism $\alpha : \mathscr{L} \to \mathscr{M}$ means applying Ξ to the images under α of its inputs. *Postcomposition* by a minion homomorphism $\beta : \mathscr{N} \to \mathscr{Q}$ means applying β to every output operation. Ordinary function composition preserves the defining equality and has the identity maps as units. A map to $\mathscr{P}$ selects one coordinate of each operation compatibly with every minor.

A *gadget reduction* replaces each variable by a copy of a fixed finite structure with a distinguished tuple. Each relation has a fixed finite gadget with one distinguished tuple per argument. For every input constraint, glue a fresh copy of its gadget to the corresponding variable copies along these tuples. For the finite templates in (1), such a reduction exists exactly when $\mathrm{Hom}(\mathscr{M}, \mathscr{N})$ holds [3, Theorem 4.12].

1.2 Consistency reductions

Dalmau and Opršal [7] introduced consistency reductions by pruning local solutions and gluing the corresponding powers. Barto, Hadek, and Zhuk [8] characterize when this construction is sound by a minion homomorphism.

Definition 1.2 (k-consistency construction)**.** Fix $k \geq 1$ and an input relational structure $\mathbf{X}$ with domain V. For each $U \subseteq V$ of size at most k, initially put $F_U = \mathrm{Hom}(\mathbf{X}[U], \mathbf{A})$, the set of all local $\mathbf{A}$-solutions on U. *Consistency pruning* means repeatedly deleting an assignment if one of its restrictions is absent from a smaller family, or if it has no extension in a larger family under consideration. The process stops at the largest subfamilies satisfying

$$\{s|_W : s \in F_U\} = F_W \qquad (W \subseteq U,\ |U| \leq k). \tag{2}$$

Let $\rho_{UW} : F_U \to F_W$ be restriction. Take the disjoint union of the powers $\mathbf{C}^{F_U}$ and identify $x \in C^{F_W}$ with $x \circ \rho_{UW} \in C^{F_U}$ for every such inclusion. Relations on the quotient are the images of the relations on these powers. Empty families are allowed: $\mathbf{C}^{\varnothing}$ is the one-element structure with every relation full. Denote the quotient by $T_k(\mathbf{X})$.

We say that a *consistency reduction* exists in (1) if T_k is a decision reduction for some fixed k.

For fixed templates and k, the construction is polynomial-time computable and satisfies $\mathbf{X} \to \mathbf{A} \implies T_k(\mathbf{X}) \to \mathbf{C}$ (Proposition D.7). Thus the condition required for a reduction is soundness:

$$T_k(\mathbf{X}) \to \mathbf{D} \quad \Longrightarrow \quad \mathbf{X} \to \mathbf{B} \qquad \text{for every input } \mathbf{X}. \tag{3}$$

The algebraic characterization uses the support construction $\omega(\mathscr{M})$ of (16) and the level construction lvl_ℓ recalled in Definition D.6.

Theorem 1.3 (Barto–Hadek–Zhuk)**.** *For the finite templates in* (1)*, a consistency reduction exists if and only if*

$$\mathrm{Hom}(\mathrm{lvl}_\ell(\omega(\mathscr{M})), \mathscr{N}) \qquad \textit{for some } \ell \geq 1.$$

This is [8, Theorem 6.14]. Their formulation uses saturated powers; [8, Theorem 6.7] compares these with families of local solutions. The comparison may change the level, so the k in T_k and the ℓ in the theorem need not agree.

All the algebraic conditions in the diagram imply a consistency reduction for finite templates. Pruning makes the restriction maps surjective; Theorem E.8 applies the injective chain argument to those maps. With arbitrary maps, the instance construction also works without pruning and uses logspace (Theorem 1.15). Consistency reductions compose by the Datalog characterization in Theorem D.8.

Example 1.4 (Arc-digraph construction)**.** For a digraph $\mathbf{X}$, its arc digraph $\delta\mathbf{X}$ has vertices $E^{\mathbf{X}}$ and arcs

$$(x,y) \longrightarrow (y,z) \qquad \text{whenever } (x,y),(y,z) \in E^{\mathbf{X}}.$$

An edgeless input can be sent to the one-vertex edgeless graph. Put $b(t) = \binom{t}{\lfloor t/2 \rfloor}$. For a symmetric digraph $\mathbf{X}$ and $t \geq 2$,

$$\delta\mathbf{X} \to K_t \quad \Longleftrightarrow \quad \mathbf{X} \to K_{b(t)}$$

by [4, Lemma 4.15]. Consequently, for $2 \leq c \leq d$, the arc-digraph construction gives a reduction

$$\mathrm{PCSP}(K_{b(c)}, K_{b(d)}) \longrightarrow \mathrm{PCSP}(K_c, K_d).$$

This is the specialization $(\mathbf{A},\mathbf{B},\mathbf{C},\mathbf{D}) = (K_{b(c)}, K_{b(d)}, K_c, K_d)$ of (1). For example, $c = 4$ and $d = 5$ give $\mathrm{PCSP}(K_6, K_{10}) \to \mathrm{PCSP}(K_4, K_5)$. An arbitrary input relation may first be made symmetric by adding every reversed edge; this preserves its homomorphisms to every K_t.

To see why this is covered by consistency reductions, express the arc construction by the rules

$$V'(x,y) \leftarrow E(x,y), \qquad E'(x,y,y,z) \leftarrow E(x,y) \wedge E(y,z).$$

These rules, together with symmetrization, form a Datalog interpretation. To enforce nonempty output domains, add an isolated copy of the input domain as a second sort and take the disjoint union of the two sorts. This preserves the existence of a homomorphism to each nonempty complete graph and gives a Datalog-with-unions reduction. The logical characterization therefore gives a consistency reduction between the same templates. This is the line-digraph example of [7, Examples 3.4 and 4.4]; the coloring reduction is [4, Lemma 4.16].

1.3 Exact set-valued homomorphisms

Replacing a single output by a bounded nonempty set gives the following relaxation of a minion homomorphism. The set is still required to commute exactly with every minor.

Definition 1.5 (Exact set-valued homomorphism)**.** Let $d \geq 1$. An *exact d-set-valued homomorphism* $\xi : \mathscr{M} \to \mathscr{N}$ assigns to every $f \in \mathscr{M}(I)$ a set $\varnothing \neq \xi(f) \subseteq \mathscr{N}(I)$ of size at most d, and satisfies

$$\xi(f^\pi) = \xi(f)^\pi \qquad (f \in \mathscr{M}(I),\ \pi : I \to J).$$

An *injective exact d-set-valued homomorphism* requires this set equality on the arrows injective on the source marks. An exact list assignment is a minion homomorphism to the minion

$$\mathcal{P}_{\leq d}(\mathscr{N})(I) = \{L \subseteq \mathscr{N}(I) : 1 \leq |L| \leq d\},$$

whose minors are taken elementwise.

Singleton lists recover ordinary minion homomorphisms. For bounded list assignments $\xi : \mathscr{M} \to \mathscr{N}$ and $\eta : \mathscr{N} \to \mathscr{L}$, their *union composition* is

$$(\eta \star \xi)(f) = \bigcup_{g \in \xi(f)} \eta(g).$$

Bounds d and e give bound de. Unions are associative, and $f \mapsto \{f\}$ is the identity. For exact lists, taking a minor commutes with each list assignment and hence with the union. Boolean conjunction on at most two coordinates shows that exact lists are strictly more general than singleton lists (Example D.1).

1.4 The d-homomorphism condition

We next require only that the transported set intersect the set assigned to the minor. This is the one-arrow case of the Barto–Kozik condition [5].

Definition 1.6 (d-homomorphism)**.** Let $d \geq 1$. A *d-homomorphism* $\xi : \mathscr{M} \to \mathscr{N}$ assigns to every $f \in \mathscr{M}(I)$ a nonempty set $\xi(f) \subseteq \mathscr{N}(I)$ of size at most d such that

$$\xi(f)^\pi \cap \xi(f^\pi) \neq \varnothing \qquad (f \in \mathscr{M}(I),\ \pi : I \to J).$$

An *injective d-homomorphism* requires this intersection on every arrow injective on the source marks.

Equal nonempty lists intersect, and intersecting singletons coincide. Thus exact lists give d-homomorphisms, and $d = 1$ recovers ordinary minion homomorphisms. Union composition again works. Given $f \to f^\pi$, choose $g \in \xi(f)$ with $g^\pi \in \xi(f^\pi)$, then $u \in \eta(g)$ with $u^\pi \in \eta(g^\pi)$. The entry u^π witnesses the intersection for the composite. Boolean majority gives a 3-homomorphism to $\mathscr{P}$ and admits no exact set-valued homomorphism to $\mathscr{P}$ (Example D.2).

1.5 Barto–Kozik homomorphisms

Barto and Kozik [5, Definition 3] introduced the following relaxation of minor preservation, in which a compatible pair may occur anywhere along a minor chain of a fixed length.

Definition 1.7 (Barto–Kozik homomorphism)**.** Let $d, r \geq 1$. A *(d, r)-Barto–Kozik homomorphism* from $\mathscr{M}$ to $\mathscr{N}$ assigns to each $f \in \mathscr{M}(I)$ a nonempty list $\xi(f) \subseteq \mathscr{N}(I)$ of size at most d, such that every minor chain $f_0 \to \cdots \to f_r$ has a *compatible pair*: positions $i < j$ with

$$\xi(f_i)^{\pi_{ij}} \cap \xi(f_j) \neq \varnothing. \tag{4}$$

Write $\mathrm{BK}_{d,r}(\mathscr{M}, \mathscr{N})$ for existence of such a list assignment and $\mathrm{BK}(\mathscr{M}, \mathscr{N})$ for existence with some finite d, r. The intersection in (4) is taken at I_j, not only at the terminal coordinate set of the minor chain.

An *injective Barto–Kozik homomorphism* chooses bounded coordinate marks on the source and requires (4) on every injective chain of r arrows. The mark bound may differ from the list bound d.

The case $r = 1$ is a d-homomorphism. For finite templates, the Barto–Kozik combinatorial gap theorem gives the polynomial-time reduction in (1) for every fixed d, r [5, Theorem 12]. These maps compose by union (Proposition D.3). Example D.4 gives a $(1, 2)$-Barto–Kozik homomorphism where no d-homomorphism exists for any d.

For example, $\mathrm{Pol}(K_3, K_5)$ has a Barto–Kozik homomorphism to $\mathscr{P}$. Writing $\mathbf{H}_c$ for the structure on $[c]$ with the nonconstant triples as its relation, there is a minion homomorphism $\mathrm{Pol}(K_3, K_5) \to \mathrm{Pol}(\mathbf{H}_2, \mathbf{H}_c)$ for some finite c [3, Theorem 6.2 and Lemma 6.4]. Wrochna's coordinate lists give the Barto–Kozik homomorphism $\mathrm{Pol}(\mathbf{H}_2, \mathbf{H}_c) \to \mathscr{P}$ [12, Lemmas 4–5 and the argument preceding Theorem 6]. The composition is recorded in [27, Examples 5.6–5.7].

1.6 Injective baby PCP

We prove the layered gap theorem by edge-coloring. If one lower level has no pivot, ordinary baby PCP recovers an assignment from that level alone. If every lower level has a pivot, we construct one injective chain on which every pair of lists disagrees.

Let V be a finite set of variables and B any nonempty set of values.[1] A *partial assignment system* of arity p gives a nonempty list $L(U) \subseteq B^U$ on every p-subset U of V. An *m-solution* of a family of partial assignments is a map $b : V \to B$ such that for every m-set X there is a member of the family whose domain contains X and whose restriction to X equals $b|_X$. For a system the family consists of all its list entries.

On each scope U we may also mark at most s assignments on U, with values in an auxiliary set common to all scopes. Apply the definition of an injective chain to the restriction maps on these assignments. Thus on $U_0 \supset \cdots \supset U_r$ the map to U_{i+1} must distinguish all distinct assignments on U_i obtained by restricting marks from positions $0, \dots, i$. Banakh and Kozik's *injective layered choice* is the injective Barto–Kozik condition with target $\mathscr{P}$ in which the coordinate list itself is marked [10, Section 1, condition *Injective layered choice*]. For this target, both lists consist of coordinates. Adjoining the marked coordinates to the choice list and marking the union preserves every required compatible pair.

Theorem 1.8 (Injective baby PCP)**.** *Fix integers $r, m, k \geq 1$ and $s \geq 0$. There are arities $d_0 > \cdots > d_r \geq m$ with the following property for every finite V with $|V| \geq d_0$ and every nonempty value set B. On level i, let L_i be a partial assignment system of arity d_i, with at most k entries and s marks per scope. Suppose every injective chain $U_0 \supset \cdots \supset U_r$, $|U_i| = d_i$, has a pair $j < i$ with*

$$L_j(U_j)|_{U_i} \cap L_i(U_i) \neq \varnothing.$$

Then some L_i, $1 \leq i \leq r$, has an m-solution. For fixed parameters, a level and a solution can be found in polynomial time from the explicit lists.

Edge-coloring on nested scopes. For $p \geq q$, the *Variable-Set graph* $G_{p,q}(V)$ has left vertices $W \in \mathcal{P}_q(V)$, right vertices $Z \in \mathcal{P}_p(V)$, and an edge (W, Z) precisely when $W \subseteq Z$. For a finite nonempty set C of colors, an *edge-coloring* is a map $\Lambda : E(G_{p,q}(V)) \to 2^C$. A color γ is allowed on (W, Z) if $\gamma \in \Lambda(W, Z)$. We also use questions indexed by a finite set Q, with a map $\sigma : Q \to V$ giving the variable of each question. In this case the left vertices are $W \in \mathcal{P}_q(Q)$ and (W, Z) is an edge precisely when $\sigma(W) \subseteq Z$. Taking $Q = V$ and σ the identity gives the Variable-Set graph above.

Lemma 1.9 (Edge-coloring on nested scopes)**.** *Fix integers $r \geq 1$, $e \geq q \geq 1$, $h \geq 0$ and a finite nonempty set C. There are $p_0 > \cdots > p_r \geq e$ with the following property. Let V, Q be finite, $|V| \geq p_0$, $|Q| \geq e$, and $\sigma : Q \to V$. For $1 \leq i \leq r$, give the edges (W, Z) with $|W| = q$, $|Z| = p_i$ and $\sigma(W) \subseteq Z$ an edge-coloring Λ_i. At least one of the following holds.*

(i) *For some $i \geq 1$ one can choose $f(W) \in C$ for every $W \in \mathcal{P}_q(Q)$ so that each $X \in \mathcal{P}_e(Q)$ has a $Z \in \mathcal{P}_{p_i}(V)$ containing $\sigma(X)$ with $f(W) \in \Lambda_i(W, Z)$ for all $W \in \mathcal{P}_q(X)$.*

(ii) *There are fixed $W_1, \dots, W_r \in \mathcal{P}_q(Q)$ and a p_0-set U_0 containing all their images. At step $i = 1, \dots, r$, after any choice of $\gamma_i \in C$ and $S_i \subseteq U_{i-1}$, $|S_i| \leq h$, one can choose $U_i \in \mathcal{P}_{p_i}(U_{i-1})$ so that*

$$S_i \cup \bigcup_{j=i}^{r} \sigma(W_j) \subseteq U_i, \qquad \gamma_i \notin \Lambda_i(W_i, U_i).$$

The choice of U_i works for all subsequent requests.

For fixed parameters, if (ii) *is impossible, i, f and the sets Z in* (i) *can be found in polynomial time on explicit input.*

[1] Ordinary baby PCP also permits infinite variable sets.

For $r = 1$ the usual local hypothesis on $G_{p_1,q}(V)$ rules out (ii). Indeed, if every p_0-set U and q-set $W \subseteq U$ admit $\gamma \in C$ and $S \subseteq U$, $|S| \le h$, with γ allowed whenever $W \cup S \subseteq Z \subseteq U$ and $|Z| = p_1$, use this request at U_0, W_1. Thus (i) gives the two-level edge-coloring conclusion.

Proof of Lemma 1.9. Put $b = \binom{e}{q}$ and $\delta = 1/(2b)$. For a sufficiently large integer n, take $p_r = n$ and $p_{i-1} = n^2 p_i$ for $i = r, \dots, 1$. Put $a_i = \lfloor p_i/(2b) \rfloor$ for $i \ge 1$. Then $e + q + ba_i \le p_i$. Order V, let F_i consist of its first p_i variables, and write $E_i(B)$ for B filled to size p_i in this order.

At each $i \ge 1$, $W \in \mathcal{P}_q(Q)$ and $\gamma \in C$ use the matrix with rows $S \subseteq V$, $|S| \le a_i$, and columns $T \subseteq V$, $|T| \le p_i - q - a_i$. Its entry is one if

$$\gamma \notin \Lambda_i(W, E_i(\sigma(W) \cup S \cup T)).$$

Apply the finite matrix alternative (Lemma D.12 in Appendix D.5). It returns either a distribution on rows giving the ones in every column mass less than δ, or a distribution on columns giving the ones in every row mass at least $\delta/4$. Call W a *pivot* at level i if no color γ at W admits such a row distribution. At a pivot every color therefore admits the stated column distribution.

No pivot at one level. At such a level i, choose for every W a color $f(W)$ for which the row distribution exists. Fix $X \in \mathcal{P}_e(Q)$, choose a row S_W independently from each of the b distributions for $W \in \mathcal{P}_q(X)$, and put $Z = E_i(\sigma(X) \cup \bigcup_W S_W)$. With all other rows fixed, $T = \sigma(X) \cup \bigcup_{W' \ne W} S_{W'}$ is a column, since its size is at most $e + (b-1)a_i \le p_i - q - a_i$. The probability that Z forbids $f(W)$ is less than δ. Since $b\delta < 1$, some Z allows every chosen color. This gives (i).

A pivot at every level. Choose one such W_i at each $i \ge 1$. Sample n columns from each of the corresponding column distributions, with all samples independent. For $0 \le i \le r$ let R_i be the union, over $j > i$, of $\sigma(W_j)$, F_j, and all columns sampled at level j. Thus $R_r = \varnothing$, and for $i < r$, $|R_i| = O(np_{i+1}) = O(p_i/n)$. In particular $|R_0| \le p_0$ and $|R_i| + h \le a_i$ for $i \ge 1$. Set $U_0 = E_0(R_0)$.

We claim that the samples can be chosen so that for each i, γ and each $S \subseteq U_0$ of size at most h, some column sampled for (i, γ) has a one in row $R_i \cup S$. Each vertex of U_0 can be named by a position in F_0, in a fixed $\sigma(W_j)$, or in a sampled column ordered as in V. There are $O(p_0)$ possible positions. Fix i, γ and at most h positions. Reveal all samples at other levels and the samples at level i named by these positions. This fixes R_i and, if all positions exist, the set S they name. At least $n - h$ samples for (i, γ) remain independent. Each has a one in that row with probability at least $\delta/4$. The probability that all miss it is at most $(1 - \delta/4)^{n-h}$. There are $O(p_0^h)$ descriptions and only $r|C|$ choices of i, γ. Their total failure probability tends to zero, since $p_0 = n^{2r+1}$. Fix samples with the claimed property.

When γ_i, S_i are requested, take an answering column T and put

$$U_i = E_i(\sigma(W_i) \cup R_i \cup S_i \cup T).$$

U_{i-1} contains R_{i-1}, hence $\sigma(W_i)$, R_i, T and F_i. It also contains S_i. Thus $U_i \subseteq U_{i-1}$, U_i contains S_i and $\sigma(W_j)$ for $j \ge i$, and $\gamma_i \notin \Lambda_i(W_i, U_i)$ by the choice of T. The claim applies at every step because $S_i \subseteq U_0$. This gives (ii).

For the algorithm, run Lemma D.12 on the polynomially many matrices, each of polynomial size. If at one level the output supplies a row distribution at every W, read off i, f and enumerate the p_i-sets to find the witnesses Z. Otherwise the output supplies, at every level, a W_i with a column distribution for every color. These are all the second case uses to prove (ii). □

Ordinary overlap. The public combinatorial gap theorem reconstructs an m-solution from ordinary overlap for each fixed alphabet in polynomial time [5, Theorem 8]. The following two-level version has alphabet-independent arities and a logspace decoder. Its pivot proof is included in Appendix D.5.

Theorem 1.10 (Barto–Kozik [6, Theorem 3.1 and Appendices A–B]). *For $m, D \geq 1$ there are $p > q \geq m$ and $n \geq p$ such that the following holds on any variable set V of size at least n. Give every p-set Z and q-set W nonempty lists $J(Z)$ and $K(W)$, respectively, of at most D assignments on that set. If $J(Z)|_W \cap K(W) \neq \varnothing$ whenever $W \subseteq Z$, then one of the two systems has an m-solution. One may take $n = p$. The parameters are independent of the value set. On finite explicit input a solution can be output in LOGSPACE.*

Arbitrary value sets. For Boolean values we apply ordinary baby PCP directly to the tuples read on a set of variables. For general B we read Boolean comparisons at those variables. Fix a finite family of input rows. At each $v \in V$, order the distinct values appearing there as $b_{v,1}, \ldots, b_{v,t_v}$. If none occurs, use a one-element list from B and set $t_v = 1$. A *rank question* (v, j) asks whether the value at v has rank at least j, for $1 \leq j \leq t_v$. Let Q be the set of questions and let $\sigma(v, j) = v$. For $X \subseteq Q$ write $\sigma(X)$ for its underlying variables and $\tau_X(u)$ for the bit tuple answered by a row u defined on $\sigma(X)$. Use the same notation for the set of tuples answered by a list.

Lemma 1.11. *Suppose $|Q| \geq 2m$. If a Boolean assignment on Q agrees on every $2m$ rank questions with one row of a family $\mathcal{F}$ of input rows, then $\mathcal{F}$ has an m-solution on V.*

Proof. Support on at most $2m$ questions follows by extending to exactly $2m$. At each v the first answer is yes, and support on pairs forces the yes answers to form an initial segment. Choose the value at its last yes. On an m-set of variables, ask for the last yes at each variable and the next answer whenever it exists. A single supporting row answers these at most $2m$ questions, hence takes all the chosen values. □

Proof of Theorem 1.8. For Boolean values take $Q = V$, take σ to be the identity, and let τ_X be restriction to X. For arbitrary B use the rank questions for all input lists. Take $e > q \geq 2m$ from Theorem 1.10 for a $2m$-solution and list bound rk. Fix an order on Q. Let C be the set of all nonempty lists of at most rk Boolean q-tuples, read in that order. Apply Lemma 1.9 with r lower levels and request bound rs. Write $d_i = p_i$ for the arities it supplies. The ordinary theorem applies on Q, since $|Q| \geq |V| \geq d_0 \geq e$. At level $i \geq 1$ the edges are the pairs (W, Z) with $|W| = q$, $|Z| = d_i$ and $\sigma(W) \subseteq Z$. Use the edge-coloring

$$\Lambda_i(W, Z) = \{\gamma \in C : \ \gamma \cap \tau_W(L_i(Z)) \neq \varnothing\}.$$

No pivot at one lower level. Let i be the level and f its choice of colors in the first alternative. At each e-set X of questions choose its common right vertex Z_X and give X the list $\tau_X(L_i(Z_X))$. At each q-set W take $f(W)$, keeping only tuples realized by some $L_i(Z)$ with $\sigma(W) \subseteq Z$. For every $W \subseteq X$ the two lists still meet after restriction to W: the overlap given by $\Lambda_i(W, Z_X)$ survives this deletion. Both lists are nonempty, have size at most rk, and every listed tuple comes from level i. The ordinary theorem gives a $2m$-solution in one of these two systems. It is supported on every $2m$ questions by an actual row of L_i. This gives an m-solution of L_i directly for Boolean values and by Lemma 1.11 for general B.

A pivot at every lower level. Take the fixed question sets $W_1, \ldots, W_r$ and the initial scope U_0 of the second alternative. Just before choosing U_i, all preceding scopes contain $\sigma(W_i)$. Request the color

$$\gamma_i = \bigcup_{j<i} \tau_{W_i}(L_j(U_j)).$$

It has at most $ik \leq rk$ tuples. On U_{i-1} there are at most $is \leq rs$ distinct restrictions of marks from levels $0, \ldots, i-1$. Choose at most rs variables that distinguish them: while two restrictions coincide

on the variables chosen so far, add one on which they differ. Each addition increases the number of distinct traces. Request these variables with γ_i. The returned U_i keeps them, so restriction to U_i is injective on the accumulated marks. It also rejects γ_i. Thus no row in $L_i(U_i)$ agrees on $\sigma(W_i)$ with a row from any earlier list. After r steps the chain is injective and has no agreeing pair, a contradiction.

The edge-coloring procedure therefore finds the first alternative. Enumerate the values occurring at each variable in order of first occurrence in the input, and hence enumerate the questions. All arities are fixed. The support tests and the search for each Z_X take polynomial time. Apply the ordinary decoder. For general B, read off the last yes at each variable. □

Layered Label Cover. A layered functional Label Cover instance has vertex layers $X_0, \dots, X_r$. Each vertex x has a finite label set A_x, and each edge from x to a vertex y of the next layer carries a map $A_x \to A_y$. A solution assigns one label to each vertex and satisfies every edge map. An *injective k-solution* assigns each vertex a nonempty set $D(x) \subseteq A_x$ of at most k labels. On a path $x_0, \dots, x_r$ let π_{ji} be the composite from x_j to x_i, with π_{ii} the identity. Whenever the path is an injective chain for these lists, meaning that each $\pi_{i,i+1}$ is injective on $\bigcup_{j \le i} \pi_{ji}(D(x_j))$, we require

$$\pi_{ji}(D(x_j)) \cap D(x_i) \neq \varnothing \qquad \text{for some } 0 \le j < i \le r.$$

For $r = 1$ this requires a matching pair exactly on edges whose map is injective on the chosen upper labels.

Corollary 1.12 (Layered injective baby PCP)**.** *For fixed $r, k \ge 1$ there is a bound $N(r,k)$ on label-set size for which it is NP-hard to distinguish solvable instances on $r+1$ layers from those admitting no injective k-solution. The reduction from 3-SAT uses logarithmic space.*

Proof. Take the arities of Theorem 1.8 with $m = 3$ and $s = k$. Pad the variable set V with unused variables to size at least d_0. Layer i consists of the d_i-subsets of V. A label on U is a Boolean assignment satisfying every clause with variables in U. Edges are inclusions between consecutive layers, and their maps are restrictions. Thus $N(r,k) = 2^{d_0}$ suffices. A satisfying assignment labels all scopes by restriction. Conversely, an injective k-solution supplies the lists and marks in the theorem. Its lower-level 3-solution satisfies every clause, since its variables extend to a 3-set.

If some scope has no label, return the same construction on a fixed odd cycle of length greater than d_0, with Boolean disequality on each edge. Every scope then has a label, since every proper induced subgraph is a forest. The cycle has no Boolean solution, so the same soundness argument gives a NO instance. For fixed r, k we can enumerate all scopes, labels and inclusions and check the clauses by rescanning the formula in logarithmic work space. □

1.7 Chain homomorphisms on injective chains

In this section we allow each list to depend on a short minor chain. A chain with one additional position tests whether the lists on its faces have a common entry.

Definition 1.13 (Chain homomorphisms)**.** Fix $h, d \ge 1$ and coordinate marks for $\mathscr{M}$. An *injective (h,d)-chain homomorphism* from $\mathscr{M}$ to $\mathscr{N}$ assigns to each h-term minor chain F in $\mathscr{M}$ a nonempty *chain list* $\Xi(F)$ of at most d operations of $\mathscr{N}$ at its terminal coordinate set. Every injective chain C with $h+1$ positions must satisfy

$$\bigcap_{i=0}^{h} (\Xi(C \setminus \{i\}) \downarrow_C) \neq \varnothing. \tag{5}$$

This *intersection of face lists* uses transport to the terminal coordinate set of C. The integers h, d are the height and list bound. Write $\mathrm{Ch}^{\mathrm{inj}}_{h,d}(\mathscr{M}, \mathscr{N})$ if such lists and uniformly bounded marks exist, and $\mathrm{Ch}^{\mathrm{inj}}(\mathscr{M}, \mathscr{N})$ if some h, d suffice.

With empty marks the same definition gives an ordinary (h, d)-*chain homomorphism*. Write $\mathrm{Ch}_{h,d}(\mathscr{M}, \mathscr{N})$ and $\mathrm{Ch}(\mathscr{M}, \mathscr{N})$ for this *chain condition*. Restricting the coordinate maps to surjections gives the surjective chain condition and its injective variant in Appendix E.3.

1.7.1 The reduction theorem

In this section we glue powers indexed by all local solutions. The restriction maps between these index sets may be arbitrary. The injectivity test concerns only the marked local solutions.

Definition 1.14 (Unpruned local-power construction)**.** Fix $k \geq 1$ and an input relational structure $\mathbf{X}$ with domain V. For each $U \subseteq V$ of size at most k, put $F_U = \mathrm{Hom}(\mathbf{X}[U], \mathbf{A})$, retaining every local solution. If some F_U is empty, output a fixed relational structure $\mathbf{Z} \not\to \mathbf{D}$. Otherwise use these unpruned families in the gluing construction of Definition 1.2: identify $x \in C^{F_W}$ with $x \circ \rho_{UW} \in C^{F_U}$, where $\rho_{UW} : F_U \to F_W$ is restriction, and take the images of the relations on the powers. Denote the resulting quotient, including the empty-family case, by $\widetilde{T}_k(\mathbf{X})$.

Theorem 1.15 (Reduction from injective chain homomorphisms)**.** *For finite promise templates as in* (1)*, suppose* $\mathrm{Ch}^{\mathrm{inj}}(\mathrm{Pol}(\mathbf{C}, \mathbf{D}), \mathrm{Pol}(\mathbf{A}, \mathbf{B}))$*. For some fixed k the unpruned local-power construction $\widetilde{T}_k$ gives a LOGSPACE decision reduction from* $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ *to* $\mathrm{PCSP}(\mathbf{C}, \mathbf{D})$*. A homomorphism $\widetilde{T}_k(\mathbf{X}) \to \mathbf{D}$ can be decoded to a homomorphism $\mathbf{X} \to \mathbf{B}$ in polynomial time. For a chain homomorphism with empty marks the decoder runs in LOGSPACE.*

An m-solution supported by local solutions satisfies every constraint of arity at most m, provided $|V| \geq m$: extend its scope to an m-set and take a witnessing list entry.

Theorem 1.16 (Injective chain consistency)**.** *Fix positive integers a, h, d and $c \geq 0$. There are $k_1 > \cdots > k_t \geq a$ and $N_0 \geq k_1$ with the following property for every nonempty value set B. Let Φ be an instance over B of constraint arity at most a on a finite set V, $|V| \geq N_0$. Mark at most c auxiliary assignments on each scope of a prescribed size. For each decreasing h-term chain F at those sizes give a nonempty list $L(F)$ of at most d local solutions on* $\min F$*. Assume* injective chain consistency*, meaning that every injective $(h+1)$-term chain C at those sizes satisfies*

$$\bigcap_{U \in C} L(C \setminus \{U\})|_{\min C} \neq \varnothing.$$

Then Φ has a global solution, computable in polynomial time from the explicit lists. With empty marks this is ordinary chain consistency*. In that case existence also holds for infinite variable sets, and finite explicit input admits a LOGSPACE decoder.*

Proof of Theorem 1.15. Fix an injective (h, d)-chain homomorphism Ξ with mark bound c. Apply Theorem 1.16 to B and the constraint arity bound a of $\mathbf{A}, \mathbf{B}$, and choose $k \geq \max(k_1, N_0, a)$. *Completeness.* If $s : \mathbf{X} \to \mathbf{A}$, then $s|_U \in F_U$ for every U. Thus every family is nonempty, and evaluation at $s|_U$ gives a homomorphism $\mathbf{C}^{F_U} \to \mathbf{C}$. These evaluations agree on the defining identifications, and hence induce $\widetilde{T}_k(\mathbf{X}) \to \mathbf{C}$.

Soundness. Let $q : \widetilde{T}_k(\mathbf{X}) \to \mathbf{D}$. Every F_U is nonempty, since the empty-family case outputs $\mathbf{Z} \not\to \mathbf{D}$. Composing q with the canonical map from each power to the quotient gives

$$f_U \in \mathrm{Pol}(\mathbf{C}, \mathbf{D})(F_U), \qquad f_U^{\rho_{UW}} = f_W \quad (W \subseteq U),$$

because the quotient identifies $x \in C^{F_W}$ with $x \circ \rho_{UW} \in C^{F_U}$. Thus a decreasing set chain C of variable sets gives a minor chain $\widehat{C}$. Restrictions compose, so taking a face commutes with this correspondence. Fix one enumeration of each F_U and apply Ξ and its marking in these coordinates throughout. Mark on U the local assignments corresponding to the marked coordinates of f_U. Consequently every injective set chain gives an injective minor chain, including the test on marks transported from different positions.

For $g \in \mathrm{Pol}(\mathbf{A}, \mathbf{B})(F_U)$, define

$$b_g(v) = g\big((s(v))_{s \in F_U}\big) \qquad (v \in U).$$

This is a local $\mathbf{B}$-solution: for a constraint $R(v_1, \dots, v_m)$ contained in U, every $s \in F_U$ gives a tuple in $R^{\mathbf{A}}$, and preservation by g gives $(b_g(v_1), \dots, b_g(v_m)) \in R^{\mathbf{B}}$. Moreover,

$$b_{g^{\rho_{UW}}} = b_g|_W,$$

since $\rho_{UW}(s)(v) = s(v)$ for $v \in W$. For each h-term chain F of the prescribed sizes give it the list $L(F) = \{b_g : g \in \Xi(\widehat{F})\}$. It has size at most d. On an injective $(h+1)$-term chain the common face entry in (5) evaluates to a common restricted assignment. Thus Theorem 1.16 gives $\mathbf{X} \to \mathbf{B}$ for $|X| \geq N_0$. For $|X| < N_0 \leq k$, a member of F_X composed with the promise map $\mathbf{A} \to \mathbf{B}$ gives the desired solution.

Complexity. Proposition D.11 implements the instance map in LOGSPACE. The operations on the enumerated coordinate sets and their minor maps have bounded tables, so all chain lists at height h can be fixed in the decoder. Given the output homomorphism, evaluating the indicated operations computes the local lists in polynomial time and even in LOGSPACE. Now apply the decoder of Theorem 1.16. Appendix D.5 gives the space bound in the ordinary case. □

1.7.2 Proof of Theorem 1.16

We prove the theorem from injective chain consistency. Here a list may depend on an entire face. At a split we must therefore avoid matches between all designated upper and lower subchains. The next lemma uses the same row and column alternatives as Lemma 1.9. Its requests also record which variables must survive later choices within a child.

Retaining variables. Split an interval of positions recursively into two nonempty consecutive intervals. In each child designate a nonempty collection of nonempty subsets of its positions. Each designated subset selects a subchain F and carries a nonempty list $L(F)$ of at most d assignments on its last scope $\min F$. The list may depend on all scopes in F. Two *separated subchains* $F < G$ (all positions of F precede those of G) *match* if $L(F)|_{\min G} \cap L(G) \neq \varnothing$.

Lemma 1.17 (Retaining variables)**.** *Fix the splits and designated subsets, integers $m, d \geq 1$ and a request bound $g \geq 0$. There are strictly decreasing scope sizes, all at least m, such that for every finite V of at least the largest size, every nonempty value set B and every system of lists as above, one of the following holds.*

(i) *For some right child, the family of entries on its designated subchains, over all choices of scopes, has an m-solution.*

(ii) *The builder can choose the scopes successively so that no designated subchains on opposite sides of any split match. After each choice the opponent may request up to g variables of that scope. Every later scope must contain all variables requested so far.*

For fixed parameters, a child and solution in (i) *or a strategy in* (ii) *can be found in polynomial time from explicit lists.*

Proof. Use the rank questions Q, their support map σ and the bit traces τ_W of Lemma 1.11 for the family of all listed entries. At each split let D_0 bound the total number of designated entries in either child in one play. Take $e > q \geq 2m$ and the threshold n from Theorem 1.10 for diameter $2m$ and list bound D_0. Take the initial size floor at least m and every such n. Thus $|Q| \geq |V| \geq n$ at every split.

For an interval with size floor k, seed budget w and request bound g, the builder must retain a given seed $S \subseteq V$, $|S| \leq w$, throughout. We obtain (i) or a random reserve $R \subseteq V$, independent of S, such that for each fixed S the builder can win inside $S \cup R$ with probability $1 - o(1)$. Winning means honoring the requests and preventing all matches across splits in the interval. The error tends to zero in one integer N, uniformly in V, the lists and the seed. All scope sizes and reserve bounds are polynomial in N. At a leaf of length l, use consecutive sizes ending at $\max(k, w + lg)$. A fixed reserve of the largest size suffices.

At a split let U, D be the left and right children, $l = |U|$, $b = \binom{e}{q}$, $\delta = 1/(2b)$ and $w' = w + lg$. Prescribe D first, with floor k, seed budget $e + bw'$ and request bound bg. Let M_D bound its reserve size, k_D its largest scope size, and put $t = e + bw' + M_D$. There are K subsets of $\{0,1\}^q$ of size at most D_0. Prescribe U with floor $k_D + 1$, seed budget $w + q + KNt$ and request bound g. If either child gives (i), we are done.

Fix $W \in \mathcal{P}_q(Q)$ and $\gamma \subseteq \{0,1\}^W$ of size at most D_0. Form a zero–one matrix with rows $S_0 \subseteq V$ of size at most w' and columns $T \subseteq V$ of size at most t. Its entry is one if the builder can play D inside $\sigma(W) \cup S_0 \cup T$, retaining $\sigma(W) \cup S_0$, honoring g requests per position, preventing the internal matches and ensuring $\tau_W(a) \notin \gamma$ for every designated entry a of D. Apply Lemma D.12 with δ. At a pivot W, no color admits the row distribution, so each γ admits a column distribution with

$$\Pr_T[(S_0, T) \text{ has entry one}] \geq \delta/4 \qquad \text{for every row } S_0. \tag{6}$$

No pivot. For each W take a γ_W and a row distribution giving the ones in each column mass less than δ. Given $X \in \mathcal{P}_e(Q)$, independently draw a row S_W for each $W \in \mathcal{P}_q(X)$ and the reserve R_D. Set $S_* = \sigma(X) \cup \bigcup_W S_W$. For the game at W use the column

$$T_W = \sigma(X) \cup R_D \cup \bigcup_{W' \neq W} S_{W'}.$$

It is independent of S_W. The probability that some entry is one or the strategy for D fails for seed S_* is at most $b\delta + o(1) < 1$. Fix draws for which all entries are zero and the strategy succeeds. Every one of the b games has this same playing area. Run one winning opponent from each game against the strategy for D, passing it the union of their requests at each move. It can honor all bg requests and prevent every internal match. In the resulting common play each opponent therefore forces a designated entry a with $\tau_W(a) \in \gamma_W$.

Use the traces on X of all designated entries of this play as the upper list. On W retain from γ_W just the realized traces of designated entries of D whose domains contain $\sigma(W)$. The common play gives ordinary overlap on every inclusion. Both lists have size at most D_0, and each entry comes from D. By Theorem 1.10 the same family supports a bit assignment on every $2m$ questions. Lemma 1.11 decodes an m-solution in (i).

A pivot. Fix a pivot W. Draw N columns for each γ using (6), all independently, and draw R_U independently. The parent reserve is $R_U \cup \sigma(W) \cup \bigcup T$, where the union runs over all draws. Fix S. With probability $1 - o(1)$ the strategy for U succeeds with enlarged seed $S \cup \sigma(W) \cup \bigcup T$, by

conditioning on the draws. With probability $1 - o(1)$ every set Z of up to lg variables requested during U and every γ have a sampled column for γ with a one in row $S \cup Z$. Indeed, fix R_U and name the variables of Z by positions in S, R_U, $\sigma(W)$ or the sampled columns, ordering each set in advance. There are polynomially many descriptions in N, since lg is fixed. For a fixed description expose only the samples it names. If all positions exist, this fixes $S \cup Z$. The remaining at least $N - lg$ independent draws for γ miss that row with probability at most $(1 - \delta/4)^{N-lg}$. The union bound over descriptions and γ gives error $o(1)$.

Play U with its enlarged seed. Take Z to be its actual requests and γ the traces on W of its designated entries. A successful column for this γ gives the required play of D. Its entire playing area lies in the last upper scope, which contains the seed, the requests and all sampled columns. The size floor gives strict inclusion, and equality of an upper and a lower entry on the lower scope would give equal traces on W. This wins the interval. The recursion for sizes uses only sums and products by constants or N, while the request bounds are independent of N. The tree is fixed, so one sufficiently large N makes all the probability bounds hold. At the root take $w = 0$.

For the algorithm evaluate the matrices from the leaves upwards. At fixed parameters each finite game has a bounded number of moves, each with polynomially many choices. Backward induction computes its value, and Lemma D.12 supplies the distributions. At the first split whose output gives a row distribution for some color at every W, find the common play for each X by enumerating plays in D retaining $\sigma(X)$ and checking the required hits. Decode the two lists as above. If there is no such split, at every split the output supplies a W with column distributions for all colors. The second case then gives (ii), and backward induction in the root game supplies the strategy. □

Matching subchains. An h-face is a subchain at any h selected positions. An assignment f *occurs* on an h-face G if its domain contains $\min G$ and $f|_{\min G} \in L(G)$.

The same matching argument applies to restrictions of assignments and to minors of operations. In a sequence of sets and maps $S_1 \to \cdots \to S_n$, write $\rho_{ij} : S_i \to S_j$ for the composite from i to j, with ρ_{ii} the identity.

Lemma 1.18 (Matching subchains)**.** *Fix $h, d \geq 1$ and $t \geq 2$. For all sufficiently large n, consider any sequence $S_1 \to \cdots \to S_n$. Give each h-set $F \subseteq [n]$ a nonempty list $L(F) \subseteq S_{\max F}$ of size at most d. Suppose that for every $(h+1)$-set A the images of its face lists in $S_{\max A}$ intersect. Then there are h-sets $B_1 < \cdots < B_t$ and entries $g_j \in L(B_j)$ such that $\rho_{\max B_j, \max B_{j+1}}(g_j) = g_{j+1}$ for $j < t$.*

Proof. Order each list. For each $(h+1)$-set choose one common image and a preimage in each face list. Color the set by the indices of these entries, in order of the omitted position. There are at most d^{h+1} colors. By the finite Ramsey theorem [14] there are $h(t+1)$ positions on which all $(h+1)$-sets have one color $(a_0, \ldots, a_h)$. Split these positions into successive h-sets $P, B_1, \ldots, B_t$.

Fix $B = B_j$. Start with $P \cup \{\max B\}$ and end with $\{\min P\} \cup B$, replacing their differing interior positions from right to left. Each replacement keeps the rank of the changed position. Deleting that position gives the same face before and after, with the same entry index. Hence the common image in $S_{\max B}$ stays fixed. At the start it is the image of entry a_h of $L(P)$, and at the end it is entry a_0 of $L(B)$. Thus all the chosen entries of the $L(B_j)$ are images of one entry of $L(P)$. □

Proof of Theorem 1.16. Take Q from Lemma 1.18 for h, d and two separated matching faces. Set $H = Q + d\binom{Q}{h}$ and $t = h2^H$. Starting with t positions, split each interval into equal halves until each leaf has h positions. On each child designate all its h-faces. Use Lemma 1.17 with $m = a$ and $g = tc$.

If the first alternative holds, the resulting a-solution is supported by local solutions of Φ and hence is a global solution. Otherwise play the second alternative, requesting, after each position, variables that distinguish all accumulated marks. At most tc variables suffice. The resulting chain and all its subchains are injective. Hence every $(h+1)$-subchain has the intersection of face lists required by Lemma 1.18. We derive a contradiction. A block of height zero is a strictly decreasing chain of h variable sets. A block of height $j+1$ is a strictly decreasing chain UD whose upper and lower halves U, D are blocks of height j, such that no h-face of U matches an h-face of D. An assignment can therefore occur on h-faces in at most one half. If it occurred on a face in each half, its two listed restrictions would make those faces match.

The play would be a block of height H. To rule this out, select one position in the upper half and continue in the lower half. For every *selected h-face* F completed by this position and every $g \in L(F)$, descend one more level into a half where g has no occurrence. All later positions will lie in that half. Repeat until Q positions have been selected. The selections use Q levels, and the descents for list entries use at most $d\binom{Q}{h}$ more, since each selected h-face is completed once. If $F < G$ are two selected h-faces, every entry of $L(F)$ was excluded before the first position of G was chosen. Hence F and G do not match, contrary to Lemma 1.18.

Thus the polynomial-time decoder of Lemma 1.17 applies. With empty marks the deterministic construction in Appendix D.5 gives the stronger value-set and space bounds. □

1.7.3 Composition and comparison

Composition. Chain homomorphisms are closed under composition: the parameters of a composite depend only on those of the two given maps. Precomposition and postcomposition by minion homomorphisms preserve the original height and list bound. Postcomposition is used in the approximate graph coloring application to pass from the projection minion to an arbitrary target minion. Existence of a chain homomorphism is therefore reflexive and transitive. Theorem F.2 gives the composite. Strict associativity of this construction on individual chain homomorphisms is not established.

For an injective chain homomorphism, precomposition and postcomposition by ordinary minion homomorphisms preserve both the parameters and the mark bound. More generally, the second map in Theorem F.2 may be any ordinary chain homomorphism when the first map is injective.

Position in the hierarchy. Every (d, r)-Barto–Kozik homomorphism gives an (r, rd)-chain homomorphism. The same implication holds with the same marks for the injective variants. At height one, chain homomorphisms are exactly d-homomorphisms; at list bound one, they are exactly minion homomorphisms. Proposition F.1 also gives the corresponding two equivalences when all three conditions use the injective variant. Whether the chain condition implies the BK condition with some finite parameters remains unresolved. The implication from chain homomorphisms to consistency reductions is proper even for finite templates (Example D.10).

2 Tight-cycle hardness and applications

In this section we derive hardness results for coloring from finite tight cycles. For a cyclic relation of prime arity p, a system of equations over $\mathbb{F}_p$ determines whether a tight cycle of the required length maps to it. For Boolean templates with unate polymorphisms, a missing block symmetry gives bounded direction certificates. These allow the same tight cycle to be used on injective chains.

Definition (Tight cycle and coloring template)**.** For positive integers n, q, write $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$ and $[q] = \{1, \dots, q\}$. Fix an integer $r \geq 2$. For $q \geq 2$, the *tight cycle* $\mathrm{CYCLE}_r[q]$ has vertex set $\mathbb{Z}_q$. Its relation consists of all cyclic rotations of $(t, t+1, \dots, t+r-1)$ for $t \in \mathbb{Z}_q$. For $q \geq 1$, the *coloring template* $\mathrm{NAE}_r[q]$ has vertex set $[q]$ and its relation consists of all nonconstant r-tuples. We also write $\mathrm{NAE}_r[q]$ for this relation.

For $r = 2$, the neighbors of x in $\mathrm{CYCLE}_2[q]$ are $x-1$ and $x+1$ modulo q, and two colors are adjacent in $\mathrm{NAE}_2[q]$ exactly when they differ. Thus $\mathrm{CYCLE}_2[q] = C_q$ for $q \geq 3$, and $\mathrm{NAE}_2[q] = K_q$ up to naming the vertices.

Theorem 2.1 (Tight-cycle source)**.** *Let p be prime, $L > 1$ satisfy $L \equiv 1 \pmod p$, and suppose $\mathrm{CYCLE}_p[L] \to \mathrm{NAE}_p[q]$. There is a finite $d = d(p, q, L)$ such that*

$$\mathrm{Ch}_{q(p-1),d}(\mathrm{Pol}(\mathrm{CYCLE}_p[L], \mathrm{NAE}_p[q]), \mathscr{P}).$$

Consequently 3-SAT $\leq_{\mathrm{L}}$ $\mathrm{PCSP}(\mathrm{CYCLE}_p[L], \mathrm{NAE}_p[q])$ with a LOGSPACE search decoder.

Appendix A gives two proofs for every prime p. Both proofs assign a coordinate list to each face and show that, in every partition of the terminal coordinates among the faces, some block meets its own face list. The continuous proof detects a block by a change in a sphere value. The combinatorial proof detects it by a change in the distance to a color class on the finite clock.

We can test whether a cyclic relation receives one of these sources by solving equations over $\mathbb{F}_p$. For a p-ary relation R, introduce one variable $z_{a_1,\dots,a_{p-1}}$ for each ordered $(p-1)$-tuple that occurs as an initial or final segment of a tuple of R. Let $\mathcal{F}_p(R)$ be the system

$$z_{a_2,\dots,a_p} - z_{a_1,\dots,a_{p-1}} = 1 \quad \text{in } \mathbb{F}_p, \qquad (a_1, \dots, a_p) \in R. \tag{7}$$

For graphs this is precisely the system specifying a bipartition. For higher arities its variables belong to ordered faces, not vertices.

Theorem 2.2 (Cyclic relations)**.** *Let p be prime and let $R \subseteq A^p$ be finite, nonempty, and invariant under the right rotation $\tau(a_1, \dots, a_p) = (a_p, a_1, \dots, a_{p-1})$. The following conditions are equivalent:*

(i) *$\mathcal{F}_p(R)$ has no solution;*

(ii) *$\mathrm{CYCLE}_p[L] \to (A; R)$ for some $L > 1$ with $L \equiv 1 \pmod p$;*

(iii) *some tuple $a \in R$ can be joined to τa by a finite sequence of tuples of R with at least one step, changing at most one coordinate at each step; a step may leave the tuple unchanged.*

If these conditions hold and $(A; R) \to (B; S)$ is a finite promise template with S containing no constant tuple, then, for $q = |B|$ and some finite d,

$$\mathrm{Ch}_{q(p-1),d}(\mathrm{Pol}((A; R), (B; S)), \mathscr{P}).$$

In particular 3-SAT $\leq_{\mathrm{L}}$ $\mathrm{PCSP}((A; R), (B; S))$ with a LOGSPACE search decoder.

Proof. Form the digraph whose vertices are the variables of (7), with an arc from $(a_1, \dots, a_{p-1})$ to $(a_2, \dots, a_p)$ for each tuple of R. Rotation invariance puts every arc on a closed directed walk of length p, so no arc joins different strongly connected components. The equations are solvable exactly when every closed directed walk has length divisible by p. Indeed, within one component choose a starting vertex and assign to each vertex the length modulo p of a directed path to it. A return path shows that this is well defined under the stated divisibility condition.

A closed directed walk is a periodic word all of whose length-p windows belong to R. If its length ℓ is not divisible by p, repeat it k times with $k\ell \equiv 1 \pmod p$, increasing k by a multiple of p if necessary to obtain $L = k\ell > 1$. The resulting word defines the homomorphism in (ii). Conversely, the image of a tight cycle of length $L \equiv 1 \pmod p$ is such a closed walk, which contradicts solvability of the equations. This proves (i)⇔(ii).

For (ii)⇒(iii), write w_t for the images of the vertices, periodically modulo L. Place w_{t+j} in position $t + j \bmod p$, for $0 \le j < p$. These tuples belong to R, consecutive tuples differ only in position $t \bmod p$, and after L steps the tuple has rotated one place to the right.

Conversely, take a path of $m \ge 1$ steps as in (iii). Reserve time slot t for changing position $t \bmod p$, carry out each change in the next available slot, and wait in the intervening slots. The changes fit in pm slots. Pad to $L = pm + 1$ slots and repeat the schedule with every position rotated after each L slots. If a^t is the tuple before slot t, then $a^{t+L} = \tau a^t$. The letters $w_t = a^t_{t \bmod p}$ therefore have period L. Each length-p window is a rotation of a^t: its successive positions have not yet been updated when their letters are read. Hence $t \mapsto w_t$ gives (ii).

Finally identify B with $[q]$. The source homomorphism gives a sandwich

$$\mathrm{CYCLE}_p[L] \longrightarrow (A; R) \longrightarrow (B; S) \longrightarrow \mathrm{NAE}_p[q].$$

Precomposing polymorphisms with the first map and renaming their outputs gives an ordinary minion homomorphism $\mathrm{Pol}((A; R), (B; S)) \to \mathrm{Pol}(\mathrm{CYCLE}_p[L], \mathrm{NAE}_p[q])$. Pull back the chain lists from Theorem 2.1, then apply Theorem 1.15. □

2.1 Approximate graph colouring

At arity two, the face equations are exactly the bipartiteness equations, and the sources are odd cycles.

Theorem 2.3 (Graph dichotomy)**.** *Let $G \to H$ be finite undirected graphs, with H loopless. If G is bipartite, then $\mathrm{PCSP}(G, H)$ is solvable in polynomial time. Otherwise 3-SAT $\le_{\mathrm{L}}$ $\mathrm{PCSP}(G, H)$ with a LOGSPACE search decoder, and for some finite d,*

$$\mathrm{Ch}_{|V(H)|,d}(\mathrm{Pol}(G, H), \mathscr{P}).$$

Proof. If G has no edge, test whether the input is edgeless and, if so, map all its vertices to any vertex of H. Otherwise, when G is bipartite, test and construct a bipartition of the input and map its two parts to the ends of a fixed edge of H. If G is not bipartite, it contains an odd cycle. Apply Theorem 2.2 with $p = 2$; the height is $|V(H)|$. □

In particular, hardness of $\mathrm{PCSP}(K_3, K_q)$ for every fixed $q \ge 3$ is the case $p = 2$, $L = 3$ of Theorem 2.1, since $\mathrm{CYCLE}_2[3] = K_3$ and $\mathrm{NAE}_2[q] = K_q$.

Corollary 2.4 (Reduction to approximate graph coloring)**.** *For every finite promise template $(\mathbf{A}, \mathbf{B})$ and fixed $n \ge 3$, there is a LOGSPACE decision reduction*

$$\mathrm{PCSP}(\mathbf{A}, \mathbf{B}) \le_{\mathrm{L}} \mathrm{PCSP}(K_3, K_n)$$

with a LOGSPACE search decoder. In particular this gives a direct LOGSPACE reduction from 3-SAT to approximate graph coloring. The polymorphism minions of 3-SAT and $\mathrm{PCSP}(K_3, K_n)$ admit chain homomorphisms in both directions.

Proof. A fixed $e : \mathbf{A} \to \mathbf{B}$ gives the minion homomorphism $\mathscr{P} \to \mathrm{Pol}(\mathbf{A}, \mathbf{B})$ sending e_i to $(x_j)_j \mapsto e(x_i)$. Postcompose the chain homomorphism of Theorem 2.3 with this map. This preserves the chain lists. Apply Theorem 1.15; K_{n+1} is a target NO instance.

Take the Boolean structure $\mathbf{S}$ for 3-SAT with the eight signed ternary clause relations. The graphs of conjunction and negation are definable by conjunctions of clauses of at most three literals, so every polymorphism of $\mathbf{S}$ preserves both operations. It is therefore a Boolean algebra homomorphism $\{0,1\}^I \to \{0,1\}$. Exactly one atom maps to 1, so the operation is the projection onto that coordinate. Thus $\mathrm{Pol}(\mathbf{S}) = \mathscr{P}$. Now Theorem 2.3 gives a chain homomorphism from $\mathrm{Pol}(K_3, K_n)$ to $\mathscr{P}$. In the other direction, the inclusion $\iota : K_3 \hookrightarrow K_n$ gives the minion homomorphism $\mathscr{P} \to \mathrm{Pol}(K_3, K_n)$ sending e_i to $(x_j)_j \mapsto \iota(x_i)$. Its singleton lists form a $(1,1)$-chain homomorphism. □

For a 3-CNF input, use $\widetilde{T}_k$ from Definition 1.14: retain all local satisfying assignments, place a power of K_3 at each local variable set, and identify vertices along restriction maps. No consistency pruning is performed. A satisfying assignment gives compatible evaluation coordinates. An n-coloring gives the local lists decoded by Theorem 1.16. Use K_{n+1} as the fixed NO instance, and also output it if the quotient has a loop; neither case occurs on a satisfiable input. All bounds depend only on n, and the entire instance construction, including the loop check, is in LOGSPACE by Proposition D.11.

2.2 Digraphs with no sources and no sinks

Barto, Kozik and Niven [16, Section 3 and Theorem 8.1] classified the CSPs of finite digraphs without sources or sinks. We consider pairs $G \to H$ of such digraphs.

A digraph is *smooth* if every vertex has an incoming and an outgoing arc. A *weak component* is a component after forgetting arc directions. A walk may traverse an arc forwards or backwards; its *algebraic length* is

$$\mathrm{al}(W) = \#\{\text{forward steps of } W\} - \#\{\text{backward steps of } W\}.$$

The same convention applies to an oriented path. Vertices may repeat in a walk, and a closed walk has equal initial and terminal vertices. For a nonempty weak component A of a finite smooth digraph, let d_A be the positive gcd of the algebraic lengths of its closed walks. Equivalently, these lengths at any fixed vertex form the subgroup $d_A\mathbb{Z}$. This number is the *algebraic length of* A.

Write $\vec{C}_d$ for the directed cycle on $\mathbb{Z}/d\mathbb{Z}$, with arcs $i \to i+1$; $\vec{C}_1$ is a loop. A *cycle sandwich* is a pair of homomorphisms $G \to C \to H$, where C is a finite disjoint union of directed cycles. For fixed digraphs $G \to H$, the input to $\mathrm{PCSP}(G, H)$ is a finite digraph X, given by its vertices and arcs. The decision problem distinguishes $X \to G$ from $X \not\to H$; the search problem asks for a homomorphism $X \to H$, promised $X \to G$.

The *arc digraph* δX has the arcs of X as vertices and arcs $(u,v) \to (v,w)$. Thus vertices of $\delta^t X$ are directed walks of length t in X. We write $X^{[d]}$ for the digraph on $V(X)$ with an arc $u \to v$ whenever X has a directed walk of exactly d steps from u to v. Finally, K_n is the loopless symmetric digraph on n vertices.

Theorem 2.5 (Smooth-digraph dichotomy)**.** *Let G, H be fixed finite nonempty smooth digraphs with $G \to H$. If H has a closed directed walk of length d_A for every weak component A of G, then*

$$G \longrightarrow C := \coprod_{A \text{ a weak component of } G} \vec{C}_{d_A} \longrightarrow H,$$

and both decision and search $\mathrm{PCSP}(G, H)$ are polynomial-time solvable through $\mathrm{CSP}(C)$. This condition is equivalent to the existence of a cycle sandwich.

Otherwise there is a reduction

$$3\text{-}SAT \leq_{\mathrm{L}} \mathrm{PCSP}(G, H).$$

It is a composition of a chain-homomorphism reduction, fixed gadget reductions, and finitely many arc-digraph reductions. Its instance map and search decoder use LOGSPACE, and no consistency-pruning algorithm is run.

When $G = H$, Barto, Kozik and Niven [16, Section 3] characterized tractability by the core being a disjoint union of directed cycles. Their proof uses a loop lemma for smooth digraphs of algebraic length one with a weak near-unanimity polymorphism [16, Theorem 8.1]. For a promise pair the condition depends instead on which closed walks exist in H. The chain homomorphism below comes from Theorem 2.2. The transfer through arc digraphs is isolated in Lemma C.1; its construction is the one in Example 1.4.

Proof of Theorem 2.5. For a weak component A, signed length from a fixed vertex, taken modulo e, gives

$$A \to \vec{C}_e \quad \Longleftrightarrow \quad e \mid d_A.$$

Consequently a cycle sandwich exists exactly under the stated condition. Indeed, a map from A to one cycle $\vec{C}_e$ forces $e \mid d_A$, and a closed directed e-walk in H can be repeated to length d_A. Conversely, the displayed union C gives the sandwich. To solve $\mathrm{CSP}(C)$, test each weak input component against each of the finitely many cycles: assign a phase at one vertex and propagate the congruences $c(v) = c(u) + 1 \pmod{d_A}$ along its arcs. These tests also construct a homomorphism to C, whose composition with $C \to H$ solves the promised search problem.

Suppose the condition fails for a weak component A, and put $d = d_A$, $B = H^{[d]}$. Then B is loopless. By Lemma C.2, a cyclic class of A, with arcs given by directed walks of length d, is a smooth weakly connected digraph Q of algebraic length one. The given map $G \to H$ restricts to $Q \to B$.

The construction in Subsection C.3, applied to Q, provides a finite smooth digraph D and an integer t with $\delta^t D \to Q$. By Lemma C.1, it suffices to prove hardness of $\mathrm{PCSP}(D, K_M)$ for every sufficiently large M. By Lemma C.5, there is a prime-arity promise template (R_D, R_{K_M}) satisfying the hypotheses of Theorem 2.2, together with a fixed gadget reduction to $\mathrm{PCSP}(D, K_M)$. Applying that theorem and then the gadget gives

$$\text{3-SAT} \leq_{\mathrm{L}} \mathrm{PCSP}(R_D, R_{K_M}) \leq_{\mathrm{L}} \mathrm{PCSP}(D, K_M).$$

The transfer lemma now gives hardness of $\mathrm{PCSP}(Q, B)$. Finally replace every input arc of $\mathrm{PCSP}(Q, B)$ by a fresh directed path of d edges. A map to Q extends along realizing walks to $A \subseteq G$, while a map of the subdivision to H restricts on its original vertices to a map to $H^{[d]} = B$. Hence $\mathrm{PCSP}(Q, B) \leq_{\mathrm{L}} \mathrm{PCSP}(G, H)$. All gadgets and iteration counts depend only on the fixed templates. Their instance maps are computable in LOGSPACE; their decoders group gadget coordinates, form the outgoing colour sets, or restrict to original vertices. Composing these with the decoder in Theorem 2.2 proves the stated LOGSPACE guarantee without consistency pruning. □

2.3 Unate Boolean polymorphisms

In this section we prove a dichotomy for Boolean promise templates whose polymorphisms are unate. A Boolean operation is *increasing* (or monotone) if changing any one input from zero to one never decreases the output. It is *unate* if each coordinate separately is either increasing or

decreasing, with its direction independent of the other inputs. A *linear threshold operation* has the form $\mathbf{1}[\sum_i w_i x_i \geq c]$ for real weights w_i and c. Its coordinates have the directions of their weights. A *symmetric threshold* is an increasing Boolean operation whose value depends only on the number of ones. An operation is *block-symmetric* on specified blocks if their union partitions its coordinates and it is invariant under all permutations within each block.

Theorem 2.6 (Unate Boolean dichotomy)**.** *Let $\mathscr{M} = \mathrm{Pol}(\mathbf{A}, \mathbf{B})$ for a finite Boolean promise template, and suppose all its operations are unate. If for every $t \geq 1$ there is an operation in $\mathscr{M}$ block-symmetric on two blocks of sizes $t, t+1$, then the decision problem $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is in P. Otherwise $\mathrm{Ch}^{\mathrm{inj}}(\mathscr{M}, \mathscr{P})$ and the decision problem is NP-hard.*

The decision dichotomy includes two earlier results. For ordered Boolean templates, the relation $\leq$ on both sides makes every polymorphism increasing. Brakensiek, Guruswami and Sandeep [17, Theorem 1.1] proved their dichotomy assuming the Rich 2-to-1 Conjecture. For templates with only linear threshold polymorphisms, Banakh and Kozik [10, Theorem 1.2 and Appendix B] proved an unconditional dichotomy using the PCP-based hardness of smooth Label Cover. Here both hardness conclusions follow from the tight-cycle construction and injective chain homomorphisms. For general unate polymorphisms, Banakh and Michno [11, Theorems 5.11, 5.14 and 5.21] gave a sufficient condition for tractability and two sufficient conditions for hardness under Rich 2-to-1. Their Question 1 asks whether their analytic criteria cover every unate template. Our criterion is block symmetry.

The difficulty is that identifying coordinates of opposite directions can change the direction of their common input. For a Boolean operation f, call a set S of coordinates increasing if changing its entries together from zero to one never decreases f, for any fixed values outside S. Define a decreasing set by the reverse inequality. A *direction certificate* for f consists of disjoint coordinate sets A, D and a direction, increasing or decreasing, such that every S containing A and disjoint from D has that direction. Its size is $|A \cup D|$. A minor map injective on $A \cup D$ transports the certificate to its image pair, since taking preimages preserves both containment and disjointness.

Proposition 2.7. *For each $t \geq 1$ there is b such that, if a minion of unate Boolean operations contains no operation block-symmetric on blocks of sizes $t, t+1$, every operation has a direction certificate of size at most b.*

We prove Proposition 2.7 in Appendix G. Failure of bounded certificates gives a minor in which each input, after orienting the cube, can be replaced by many disjoint pairs of other inputs. A random identification then gives the missing block symmetry. For the increasing operations, the following consequence of rotation on a tight cycle will supply the hardness source.

Lemma 2.8. *For each $m \geq 2$ there are a prime p and $d \geq 1$ such that, for $n = pd + 1$, every increasing $g : \{0,1\}^{\mathbb{Z}_n} \to \{0,1\}$ having no symmetric threshold minor of arity m or $m+1$ satisfies*

$$g(S) > g\big((S + d) \cup (S + d + 1)\big) \qquad \text{for some } S \subseteq \mathbb{Z}_n.$$

Here sets denote their indicator vectors and translations are modulo n.

Proof of Theorem 2.6. The tractable case is the combined basic linear programming and affine integer programming algorithm of Brakensiek, Guruswami, Wrochna and Živný [18, Theorems 4.1 and 5.1]. The two block sizes are their exact criterion for this decision algorithm.

Suppose the block sizes $t, t+1$ are missing. No operation can have a constant diagonal $x \mapsto f(x, \ldots, x)$, since its constant minors have every block symmetry. Normalize each f by composing its output with the inverse of its diagonal. The resulting minion $\mathscr{M}'$ consists of unate

operations that send the all-zero input to zero and the all-one input to one. Normalization commutes with minors and preserves block symmetry in both directions. By Proposition 2.7 each operation of $\mathscr{M}'$ has a bounded certificate. If (A, D) is decreasing then (D, A) is increasing. Indeed for any decreasing set S, $f(S^c) \geq f(\mathbf{1}) = 1 > f(\mathbf{0})$, and identifying S^c gives a unate minor. This one strict increase fixes its direction in all contexts.

Put $m = t(t+1)$. A symmetric threshold of arity m would give the missing block symmetry by taking $t+1$ copies of every first-block input. At arity $m+1$ take t copies of each first-block input and one of each second-block input. Thus we may use p, n, d from Lemma 2.8. Write E for the tuples of $\mathrm{CYCLE}_p[n]$ and ρ_a for the map taking entry a, $0 \leq a < p$. An increasing $F \in \mathscr{M}'(E)$ cannot have all F^{ρ_a} equal. Otherwise the minors $F^{d\rho_a}$ also coincide, say with value g as an operation on $\mathbb{Z}_n$. Along two successive entries of each tuple, multiplication by d gives an increment of d or $d+1$. Hence for every $S \subseteq \mathbb{Z}_n$,

$$(d\rho_0)^{-1}(S) \subseteq (d\rho_1)^{-1}((S+d) \cup (S+d+1)).$$

Applying F contradicts Lemma 2.8. The cycle argument with the bounded certificates (Lemma G.2, Appendix G) now gives $\mathrm{Ch}^{\mathrm{inj}}(\mathscr{M}', \mathscr{P})$. Pull back its marks and face lists along normalization to obtain $\mathrm{Ch}^{\mathrm{inj}}(\mathscr{M}, \mathscr{P})$. Any bijection from the vertices to $[n]$ is a homomorphism $\mathrm{CYCLE}_p[n] \to \mathrm{NAE}_p[n]$. Composing projections with it gives a minion homomorphism from $\mathscr{P}$ to the polymorphisms of this template. Postcompose the face lists with that map. By Theorem 1.15 the resulting injective chain homomorphism reduces $\mathrm{PCSP}(\mathrm{CYCLE}_p[n], \mathrm{NAE}_p[n])$ to $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$. The source is NP-hard by Theorem 2.1. For each bit c failure of its constant polymorphism supplies a nonempty strict relation forbidding the constant c tuple on the weak side. Imposing both relations on one variable gives the required fixed NO instance. □

For increasing polymorphisms, the criterion can be stated using symmetric thresholds alone, and the tractable case also admits polynomial-time search. If symmetric thresholds have bounded arity, the proof just given uses Lemma 2.8 directly, with the empty certificate for every operation. Indeed the maps $f \mapsto (x \mapsto f^{dx})$ give an ordinary minion homomorphism to $\mathrm{Pol}(\mathrm{CYCLE}_p[n], \mathrm{NAE}_p[|\mathscr{M}(\mathbb{Z}_n)|])$ by the same inclusion of preimages. The unary identity supplies the promise map. Thus Theorem 2.1 gives the ordinary chain homomorphism to projections. If symmetric thresholds have unbounded arity, take the relative-interior rational basic LP solution w and the affine integer solution a on its support used in [18, Section 3]. Sort the variables by their label-one marginals (w_v, a_v) in lexicographic order and try setting each suffix to one. Let L clear the denominators of w. For a sufficiently large admitted threshold arity N, put $r = N \bmod L$. The tuple multiplicities $(N-r)w + ra$ are nonnegative integers summing to N at each constraint, and their label-one marginals are nondecreasing in the chosen order. Applying the threshold to these tuples gives one of the tested suffixes and satisfies every weak constraint.

2.4 Rainbow colouring

All seven hypergraph hardness statements below follow from chain homomorphisms with finite parameters, and the reductions have LOGSPACE search decoders. They also hold for edges with distinct vertices, by Lemma B.3.

A rainbow c-colouring uses all c colours on every edge. Guruswami and Sandeep asked whether a proper two-colouring is hard to find for a rainbow six-colourable, seven-uniform hypergraph [21, Section 5]. Write $U_{r,c}$ for the rainbow relation.

Corollary 2.9. *Let $2 \leq c \leq r$ and $q \geq 2$. Then $\mathrm{PCSP}(U_{r,c}, \mathrm{NAE}_r[q])$ is in P when $c = r$ and is NP-hard when $c < r$.*

Any nonconstant map from $[c]$ to $[q]$ witnesses validity. When $c = r$, each colour occurs once, so Lemma B.2 applies. When $c < r$, take a word of length $r-1$ containing all c colours and repeat it periodically. Every r-window still contains all c colours, and Lemma B.1 gives hardness. For $(r, c, q) = (7, 6, 2)$ the word 123456 gives the source map and height 12. Guruswami and Sandeep proved the two-colour-output cases through uniformity six and exhibited polymorphisms with no fixing set of bounded size at $(7, 6, 2)$ [21, Theorem 2 and Proposition 17]. For even $r \geq 4$, Guruswami and Lee proved hardness for rainbow $(r/2)$-colourable inputs and every fixed output palette [32, Corollary 1.2]. Brakensiek and Guruswami obtained hardness for the $(r-1)$-rainbow promise for every $r \geq 3$ and fixed output palette under their induced V Label Cover conjecture [28, Theorem 1.2 and footnote 6].

2.5 Strong colouring with a spare colour

A strong c-colouring assigns distinct colours to the r vertices of each edge. Write $I_{r,c}$ for this relation.

Corollary 2.10. *For $c \geq r \geq 2$, the template $(I_{r,c}, \mathrm{NAE}_r[q])$ is valid exactly when $c \leq q(r-1)$. Among valid parameters, proper q-colouring is in P at $c = r$ and is NP-hard at $c > r$.*

Each fibre of a colour map must have size at most $r-1$, and this is sufficient. At $c = r$, each colour occurs once, so use Lemma B.2. At $c > r$, periodically repeat $r+1$ distinct colours; every r-window is injective. Apply Lemma B.1. In particular the word 123456 proves the five-uniform, strong-six-colour case, with height 8 for Boolean output. The one-spare-colour cases were posed by Brakensiek and Guruswami [22, Conjecture 1.2]. Their Theorem 1.3 gives a bound of $\lceil 3r/2 \rceil$ strong colours for Boolean output when $r \geq 4$ and proves the $r+1$ bound for $r = 3, 4$. At $r = 2, c = 3, q = 2$ the promise fails the fibre test.

2.6 At least a prescribed number of colours

Austrin, Bhangale and Potukuchi studied colourings with at least t distinct colours on every edge and proved hardness in quantitative ranges of the parameters [25, Theorem 1.7 and Corollary 1.8]. Write $D_{r;c,t}$ for this strong relation.

Corollary 2.11. *For $2 \leq t \leq \min(c, r)$, the template $(D_{r;c,t}, \mathrm{NAE}_r[q])$ is valid exactly when $c \leq q(t-1)$. Among valid parameters, $c = t = r$ gives a polynomial algorithm and every other case is NP-hard.*

A fibre of size t can contain the colours of an entire strong edge; fibres of size at most $t-1$ therefore characterize validity. If $t < r$, periodically repeat a word of length $r-1$ containing t colours. If $t = r < c$, periodically repeat $r+1$ distinct colours. In either case all r-windows are legal, so Lemma B.1 applies. At $c = t = r$, each colour occurs once and Lemma B.2 applies. For example, $D_{5;4,3}$ already contains the known hard rainbow promise $U_{5,4}$ for Boolean output [21, Theorem 2].

2.7 Sparse cyclic Boolean relations

Let $r = 2m + 1 \geq 5$. The strong relation C_r consists of the rotations of $1^m 0^{m+1}$ and its bitwise complement. Its $2r$ words specify one cyclic interval of ones, of length m or $m+1$.

Corollary 2.12. *For every odd $r \geq 5$ and $q \geq 2$, $\mathrm{PCSP}(C_r, \mathrm{NAE}_r[q])$ is NP-hard.*

Repeat $0^m 1^m$ periodically. Each length-r window duplicates its first letter at the cyclic join and is therefore in C_r. Its period $r-1$ is coprime to r, so Lemma B.1 applies. For the ten-tuple

relation $R_5 = C_5$, at $q = 2$ the height is 8. Austrin, Håstad and Martinsson discuss both the folded SAT problem and this ordinary ten-tuple PCSP in [29, Section 8.1]. Banakh, Barsukov and Nakajima derive hardness for the latter under Rich 2-to-1 [19, discussion following Corollary 1.14]. The corollary concerns the ordinary PCSP.

The same argument allows more general restrictions on runs. Let the allowed Boolean words have Hamming weight in $[\ell, u] \subseteq [1, r-1]$, and let the number of cyclic runs of ones belong to $J \subseteq \{1, \dots, \lfloor r/2 \rfloor\}$. Assume the relation is nonempty and put $t = \min J$. The feasible integer weights are exactly $[\ell, u] \cap [t, r-t]$, since j runs require at least j zeros and j ones, and these bounds suffice. If only one weight occurs, use Lemma B.2. Otherwise choose consecutive feasible weights $k, k+1$. Arrange k ones and $r-1-k$ zeros in a circular word with t runs of each, and repeat it. Every r-window duplicates one letter, preserves the run count, and has weight k or $k+1$. Lemma B.1 proves hardness.

2.8 Linearly ordered triples

Let LO_c consist of the triples on $\{0, \dots, c-1\}$ whose maximum occurs once. One may ask for either a proper colouring or a colouring with a unique maximum on every edge.

Corollary 2.13. *For $3 \le c \le q$, both $\mathrm{PCSP}(\mathrm{LO}_c, \mathrm{NAE}_3[q])$ and $\mathrm{PCSP}(\mathrm{LO}_c, \mathrm{LO}_q)$ are NP-hard, with height $2q$.*

Repeat the word 0012. Each three-letter window omits one position of this word: omitting 2 leaves a unique 1, and every other omission leaves a unique 2. Thus the period-four word gives the source map for both targets, so Theorem 2.1 applies. The order-preserving injection into $[q]$, with its usual order, witnesses validity. For proper output, $c = 2$ is polynomial by Lemma B.2, since each triple contains exactly one occurrence of the higher label. For $c \ge 2$, the triples (a, a, b), $a < b$, force a colour map to be injective, so validity requires $q \ge c$.

Filakovský et al. proved the case $\mathrm{LO}_3 \to \mathrm{LO}_4$ [30, Theorem 1.1]. Krokhin and Vagnozzi proved hardness for $\mathrm{LO}_2 \to \mathrm{LO}_3$ [20, Theorem 1]. Iterating the gadget in their proof of Corollary 29 gives hardness for $\mathrm{LO}_c \to \mathrm{LO}_{c+1}$ for every $c \ge 2$ [19, Theorem 1.13]. For $c \ge 4$ and ordered output there is also a reduction from $\mathrm{PCSP}(K_{c-1}, K_{q-1})$ [30, p. 34:3, footnote 2]:

$$\exists z \; R(z, z, x) \wedge R(z, z, y) \wedge R(x, y, z)$$

defines the complete graph on the nonminimum colours of LO_c, with the minimum colour isolated. Since $K_3 \to K_{c-1}$, the graph hardness of Fei, Minzer and Wang [2, Corollary 1.7] already gives the ordered-output decision result for $c \ge 4$. At $c = 3$ the graph is K_2, and for a proper-colouring target the same formula permits $x = y$. The period-four source applies in both situations. Banakh, Barsukov and Nakajima obtained the ordered-output hardness under Rich 2-to-1 [19, Corollary 1.14]. For every uniformity $r \ge 4$ and $q \ge 3$, Nakajima et al. proved hardness of finding a conflict-free q-colouring of a linearly ordered three-colourable hypergraph [31, Theorem 17]. This also implies hardness for ordered output in those uniformities whenever $3 \le c \le q$.

2.9 Capacity-constrained colourings

Write $R^{(c)}_{k:r}$ for the relation in which a colour occurs at most $b = k - 1$ times. Assume $1 \le b < r$. The relation is nonempty exactly when $r \le cb$.

Corollary 2.14. *In the nonempty case, $(R^{(c)}_{k:r}, \mathrm{NAE}_r[q])$ is valid exactly when $c \le q\lfloor (r-1)/b \rfloor$. Among valid parameters, $r = cb$ gives a polynomial algorithm and $r < cb$ gives NP-hardness.*

A fibre containing j colours can fill an edge exactly when $jb \geq r$; this proves the validity criterion. At $r = cb$, every colour occurs exactly b times, so Lemma B.2 applies. If $r < cb$, choose $r + 1$ letters with each colour occurring at most b times, and repeat this word. Every r-window deletes one position and therefore respects the capacities. Apply Lemma B.1. The case $k = 2$ is strong colouring. Banakh, Barsukov and Nakajima use this colouring template in their dichotomy for vertex-coloured monochromatic cliques. Their hardness direction assumes Rich 2-to-1 [19, Lemma 1.6 and Theorem 1.7].

2.10 Discrepancy one

For $r = 2m + 1 \geq 3$, let W_r have vertex set $\{0, 1\}$ and the single r-ary relation $\{x \in \{0, 1\}^r : \sum_{i=1}^{r} x_i \in \{m, m + 1\}\}$. Thus the two colour counts differ by one.

Corollary 2.15. *For every odd $r \geq 3$ and $q \geq 2$, $\mathrm{PCSP}(W_r, \mathrm{NAE}_r[q])$ is NP-hard.*

The periodic word $0^m 1^m$ again has only legal length-r windows, now with arbitrary orders also allowed in the relation. Apply Lemma B.1. A single prescribed weight $0 < k < r$, including exact balance at even arity, instead gives the polynomial algorithm of Lemma B.2. For Boolean output, the hardness was proved by Austrin, Guruswami and Håstad [33, Theorem 1.2]. For $r = 3$, hardness for every fixed output palette is the theorem of Dinur, Regev and Smyth [34]. This case also admits a Barto–Kozik certificate [12, proof of Theorem 1][27, Example 5.6]. For triples, Fei, Minzer and Wang proved the stronger statement that for every fixed $0 < \delta \leq 1/2$ it is NP-hard to distinguish a two-colourable hypergraph from one with no independent set of relative size at least δ [2, Corollary 1.10]. Guruswami and Lee proved hardness for even-uniform discrepancy two and stated the odd-uniform discrepancy-one extension under a d-to-1 assumption, omitting its proof [32, Corollary 1.3 and footnote 1]. The perfect-completeness 4-to-1 theorem supplies this assumption with $d = 4$ [2, Theorem 1.6]. Bhattiprolu, Guruswami and Lee obtained a conditional hardness result for near-discrepancy-one promises at large odd arities and an unconditional one for discrepancy two at even arities, both with Boolean output [26, Theorems 2 and 3].

A Two proofs of Theorem 2.1

In this section we prove the tight-cycle theorem for every prime p. The continuous proof represents polymorphisms by equivariant Lipschitz maps from powers of the circle to a fixed sphere. The combinatorial proof chooses terminal coordinates on a finite clock and proves their intersection by cyclic covering.

A.1 Continuous proof

In this section we factor the chain homomorphism of Theorem 2.1 through Lipschitz maps. We give the factorization first and prove the coordinate-selection theorem at the end of the section.

The graph case. Give the circle $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ its shortest-arc distance, each power $\mathbb{T}^I$ the maximum of the coordinate distances, and the unit sphere $S^q \subseteq \mathbb{R}^{q+1}$ Euclidean distance. Fix a Lipschitz constant C for all arities. Consider all C-Lipschitz maps $F : \mathbb{T}^I \to S^q$ satisfying the antipodal identity

$$F(x + \tfrac{1}{2}\mathbf{1}) = -F(x).$$

Coordinate substitution makes these maps a minion, and this minion admits a chain homomorphism to projections.

To use this for (C_L, K_k), we would like to assign such a map to each polymorphism and preserve minors exactly. For $k \geq 5$ no such minion homomorphism exists (Proposition A.8). Keep the same maps and allow G to be a π-minor of F whenever their values make an angle of at most $90°$, that is,

$$\langle F(x \circ \pi), G(x) \rangle \geq 0 \qquad \text{for every } x \in \mathbb{T}^J, \quad \pi : I \to J.$$

For the first arrow, place vertex $2j$ mod L at j/L on the circle, as in Figure 1 for C_5. At x take the normalized difference of the color distributions near x and near its antipode. Their supports are disjoint, and taking a minor can only remove available colors. Thus the difference changes by at most a right angle, giving a homomorphism to the relaxed minion. For the second arrow, put on each face list those terminal coordinates whose individual half-turn changes a sphere observation appreciably on average. If the lists had empty intersection, we could assign every terminal coordinate to a face missing it. Borsuk–Ulam on the boundary of a cube forces one of these blocks to move an observation. A Lipschitz estimate then puts a coordinate of that block on its own face list.

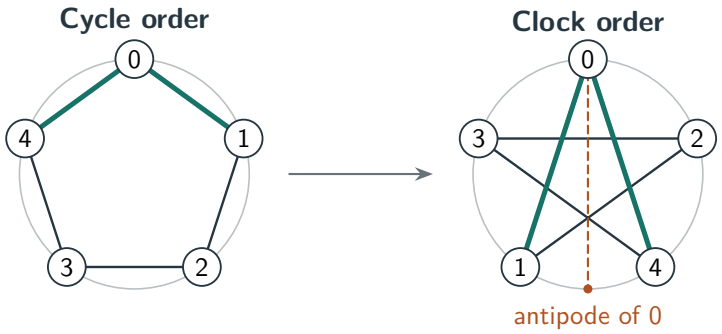


Figure 1: The highlighted edges in both drawings of C_5 join 0 to its neighbors 1 and 4. In clock order, the dashed diameter ends at the midpoint of the short arc between those neighbors.

What changes for odd primes. Replace the half-turn by simultaneous rotation through $1/p$ of a turn. A sphere value now has one row for each of the p phases and one column for each color. Rotation cycles the rows. Under a minor, every phase minimizing a given column must still minimize that column. The vectors along a chain share a minimizing phase in each column, which keeps their convex hull away from zero. The half-turn test becomes a $1/p$-turn test, and the interval of signed depths becomes a star with p arms.

A.1.1 Lipschitz maps and their minors

In this section we state the geometric condition that gives a chain homomorphism for both relaxations. Throughout, q is the dimension of the target sphere and C is independent of arity.

Definition A.1 (Lipschitz maps)**.** A map $F : X \to Z$ is C-Lipschitz if $d_Z(F(x), F(y)) \le C\, d_X(x, y)$ for all $x, y \in X$. Fix a prime p and an orthogonal transformation T with $T^p = \mathrm{id}$ such that T^a has no fixed point on S^q for $1 \le a < p$. For finite nonempty I, let $\mathcal{Lipshit}_{\mathcal{Z}}[p, q \mid C](I)$ consist of the C-Lipschitz maps $F : \mathbb{T}^I \to S^q$ satisfying

$$F(x + \tfrac{1}{p}\mathbf{1}) = TF(x).$$

We use the metrics above and take C for which the family is nonempty.

Proposition A.2 (Ordinary minors)**.** *The maps $\mathcal{Lipshit}_{\mathcal{Z}}[p, q \mid C]$ form a minion under $F^\pi(x) = F(x \circ \pi)$.*

Proof. Coordinate substitution is 1-Lipschitz in the maximum metric, commutes with simultaneous rotation, and respects identity and composition. □

Corollary A.3 (Ordinary chains)**.** *If $h > q$, then $\mathcal{Lipshit}_{\mathcal{Z}}[p, q \mid C]$ admits an (h, D)-chain homomorphism to projections for some finite D.*

Proof. Evaluating an ordinary chain on the pullbacks of a common terminal input gives a single unit vector. Apply Theorem A.6 below with distance one from zero. □

Definition A.4 (Relaxed minion)**.** A relaxed minion has a set of operations for each finite nonempty coordinate set I and, for each $\pi : I \to J$, a relation $f\ [\pi]\ g$ between I-ary and J-ary operations. We require $f\ [\mathrm{id}]\ f$. A homomorphism preserves arities and these relations. In an ordinary minion, $f\ [\pi]\ g$ means $g = f^\pi$.

A relaxed chain requires $f_i\ [\pi_{ij}]\ f_j$ for every $i < j$, where π_{ij} is the composite coordinate map. Thus every subchain is again a relaxed chain. A chain homomorphism to projections assigns bounded lists to chains of h rows, with intersecting transported face lists on every relaxed chain of $h + 1$ rows.

Definition A.5 (Relaxed minors of Lipschitz maps)**.** Write $\mathcal{Lipshit}_{\mathcal{Z}}^{\mathrm{rel}}[p, q \mid C]$ for the same sphere maps with the following minor relations. For $p = 2$ take $T = -\mathrm{id}$ and put $F\ [\pi]\ G$ if $\langle F(x \circ \pi), G(x)\rangle \ge 0$ for every x.

For odd p and $k \ge 2$ colors, take the unit sphere of the real p-by-k matrices whose row and column sums are zero. Its dimension is $q = (p-1)(k-1) - 1$. Cyclic row shift acts freely, since a matrix fixed by a nontrivial shift has equal rows and hence is zero. Put $F\ [\pi]\ G$ if for every x and every color c,

$$\operatorname*{arg\,min}_{a \in \mathbb{Z}_p} F(x \circ \pi)_{a,c} \ \subseteq\ \operatorname*{arg\,min}_{a \in \mathbb{Z}_p} G(x)_{a,c}.$$

Theorem A.6 (Relaxed chains)**.** *If $h > q$, then $\mathcal{L}ip\mathcal{S}hit_{\mathcal{Z}}^{\mathrm{rel}}[p, q \mid C]$ admits an (h, D)-chain homomorphism to projections for some finite D. More generally, the conclusion holds for any minor relations on $\mathcal{L}ip\mathcal{S}hit_{\mathcal{Z}}[p, q \mid C]$ satisfying the following condition. There is a $\rho > 0$ such that for every relaxed chain of at most $h + 1$ rows, with last row r, and every $x \in \mathbb{T}^{I_r}$,*

$$\operatorname{dist}(0, \operatorname{conv}\{F_j(x \circ \pi_{jr}) : j \text{ in the chain}\}) \geq \rho. \tag{8}$$

Here D depends only on p, q, C, h, ρ. The list can be chosen by measuring the mean change when one terminal coordinate turns by $1/p$.

The bound $h > q$ suffices for hardness.[2]

For the stated relaxed minors, the margin follows directly. For $p = 2$ the transported unit vectors have pairwise nonnegative inner products. Any convex combination with weights λ_j has squared norm at least $\sum_j \lambda_j^2 \geq 1/(h + 1)$.

For odd p, fix a terminal input and choose a minimizing phase a_c in each column of the first matrix. It minimizes that column in every transported matrix U. The linear functional $U \mapsto -\sum_c U_{a_c,c}$ therefore agrees on these matrices with $-\sum_c \min_a U_{a,c}$. The latter function is strictly positive on the unit sphere: every column sums to zero and some column is nonzero. By compactness it has a positive minimum. The linear functional has norm at most $\sqrt{k}$, so the same positive distance bound holds for every chain, after dividing this minimum by $\sqrt{k}$.

A.1.2 From the tight cycle to the sphere

In this section we construct the first homomorphism by placing the tight-cycle vertices on the circle. The colors attainable near each input determine its sphere value and control how it changes under a minor.

Proposition A.7. *Let $L > 1$ satisfy $L \equiv 1 \pmod p$ and suppose $\mathrm{CYCLE}_p[L] \to \mathrm{NAE}_p[k]$. With $q = (p - 1)(k - 1) - 1$ and some $C = C(p, k, L)$, there is a homomorphism of relaxed minions*

$$\mathrm{Pol}(\mathrm{CYCLE}_p[L], \mathrm{NAE}_p[k]) \longrightarrow \mathcal{L}ip\mathcal{S}hit_{\mathcal{Z}}^{\mathrm{rel}}[p, q \mid C].$$

Proof. Place vertex $pj \bmod L$ at the point j/L of $\mathbb{T}$. Measure distances temporarily in units of the spacing $1/L$. Every point is within $1/2$ of a vertex. Choose a radius $r = (1 + 1/p)/2$, and write $B(t)$ for the cycle vertices at distance strictly less than r from t in these units.

Any choice of one vertex from each $B(t - a/p)$, in phase order $a = 0, \ldots, p - 1$, forms a tight-cycle tuple. Indeed, write $L = p\ell + 1$. The ideal backward step is $L/p = \ell + 1/p$ spacings. The actual step is an integer differing from it by strictly less than $2r = 1 + 1/p$, so it is ℓ or $\ell + 1$. Around one turn the steps sum to L, including the return step, so exactly one is longer. The labels thus increase by 1 except for a single increment $1 - p$.

For a polymorphism f, its color classes $f^{-1}(c)$ are now finite subsets of $\mathbb{T}^I$. Give color c the weight $\max\{0, r - \operatorname{dist}_\infty(x, f^{-1}(c))\}$, using the same spacing units and weight zero for an empty class. A tuple of nearest vertices gives some color weight at least $r - 1/2$. Each weight is L-Lipschitz in the unit-circle metric, so normalizing them to sum to one gives a probability vector $P_f(x)$ with an arity-independent Lipschitz bound. Its support is exactly the palette $f(\prod_i B(x_i))$.

Let P have rows $P_f(x - \frac{a}{p}\mathbf{1})$, $a \in \mathbb{Z}_p$. Every column has a zero. A color present at every phase would have witnesses forming a tight-cycle tuple in each coordinate, with constant output. Subtract

[2] Using the boundary of a product of $h + 1$ disks gives the sharper bound $2h + 1 > q$. Keeping the same $1/p$-turn test then requires an additional degree argument.

the mean row from each row and normalize to unit norm. This gives $\Phi(f)(x)$ on the sphere of matrices with zero row and column sums. The centered norm is bounded away from zero: on the compact set of such probability tables it could vanish only if all rows were equal, which is incompatible with a zero in each column. Thus $\Phi(f)$ is Lipschitz with a bound depending only on p, k, L. Translation by $1/p$ shifts row a to row $a+1$. For $p = 2$ the rows are negatives of one another, so this shift is the antipodal map.

For a minor f^π, every witness of a color at x pulls back to a witness of the same color at $x \circ \pi$. Thus every zero in the table for f persists in the table for f^π. For odd p, the zero entries in each column are exactly its minimizing phases, and centering and normalization preserve those minimizers. For $p = 2$ the two rows of P have disjoint supports. In their difference, a surviving color therefore has the same sign before and after taking the minor. The products of corresponding entries of the two sphere values are nonnegative, so their dot product is nonnegative. □

Continuous proof of Theorem 2.1. Take $h = k(p-1) > q = (p-1)(k-1) - 1$. Proposition A.7 and Theorem A.6 give the two arrows

$$\mathrm{Pol}(\mathrm{CYCLE}_p[L], \mathrm{NAE}_p[k]) \xrightarrow{\text{homomorphism}} \mathscr{L}ip\mathscr{s}\mathscr{h}it_{\mathscr{Z}}^{\mathrm{rel}}[p, q \mid C] \xrightarrow{\text{chain homomorphism}} \mathscr{P}.$$

Assign each source chain the list of its image. Taking a face commutes with the first arrow and leaves the transport maps unchanged, so these lists intersect as required. The constant-tuple constraint is a fixed NO instance for $\mathrm{NAE}_p[k]$, and Theorem 1.15 gives the stated reduction and search decoder. □

A.1.3 The obstruction to preserving minors exactly

In this section we prove the obstruction motivating the relaxed relation. Exact preservation of minors and a common Lipschitz bound would identify the images of the two binary projections.

Proposition A.8. *Suppose $p = 2$ and $k \geq 5$, or p is an odd prime and $k \geq 3$. For the lengths L of Theorem 2.1 there is no ordinary minion homomorphism from $\mathrm{Pol}(\mathrm{CYCLE}_p[L], \mathrm{NAE}_p[k])$ to equivariant C-Lipschitz maps $\mathbb{T}^I \to S^q$ with ordinary coordinate substitution, for any fixed C, q and free cyclic action on the sphere.*

Proof. We first join the two binary projections by recoloring one off-diagonal entry at a time in a smaller template. For $p = 2$, precomposition with a 3-coloring of C_L and inclusion of five colors into K_k let us use $\mathrm{Pol}(K_3, K_5)$. View its binary operations as 3-by-3 tables, numbering the colors from zero. Starting with the row projection, recolor the off-diagonal cells of row zero with color 3 and those of row one with color 4, one cell at a time. Now recolor columns zero, one, and two in order with their column indices. Each change is proper, since the new color is confined to the row or column being filled.

For $p \geq 3$, alternating two colors around the cycle gives a map $\mathrm{CYCLE}_p[L] \to \mathrm{NAE}_p[2]$: even across the closing edge, three consecutive vertices cannot all have the same color. Including three colors into $[k]$ now lets us use $\mathrm{Pol}(\mathrm{NAE}_p[2], \mathrm{NAE}_p[3])$. Keep the values at $00, 11$ equal to $0, 1$ and change the pair of values at $01, 10$ along

$$(0,1), \quad (2,1), \quad (2,0), \quad (1,0).$$

Each table is a polymorphism. A tuple of Boolean pairs nonconstant in both coordinates contains either $00, 11$ or $01, 10$, and each opposite pair receives distinct colors.

Let a, b be consecutive tables differing at (u, v), where $u \neq v$. Define $H_n(x, y, z_1, \ldots, z_n)$ to use table b when an odd number of the switches z_j equal v, and table a otherwise. This is a

polymorphism in either case. A constant output on a relation tuple would already be constant for one table: away from (u, v) they agree, and at every occurrence of (u, v) the chosen color would be the same. In particular $H_1(x, y, x) = a(x, y)$ and $H_1(x, y, y) = b(x, y)$.

An exact homomorphism as in the statement would, by composition, give one on the smaller template. Denote it by Φ. Equal switch pairs cancel in H_n. Fix the first two circle inputs and subdivide an arc from z_0 to z_m into m equal steps. By exact minor preservation, evaluating $\Phi(H_{2m+1})$ with the two switch lists

$$(z_0, z_1, z_1, \ldots, z_m, z_m), \qquad (z_0, z_0, \ldots, z_{m-1}, z_{m-1}, z_m)$$

gives $\Phi(H_1)$ at z_0 and at z_m, respectively. The lists are within one arc step in the maximum metric. The common bound C and $m \to \infty$ force the two values to agree. Thus the image of H_1 ignores its switch, and the two displayed minor identities give $\Phi(a) = \Phi(b)$.

Following the recolorings identifies the images of the row and column projections. Both are minors of the unary color inclusion, so its image takes the same value at every two circle points. A constant equivariant sphere map would give a fixed point of the free action. □

A.1.4 From relaxed chains to projections

In this section we prove that the face lists intersect whenever $h > q$. Topology detects a block of coordinates whose turn changes a sphere value. A Lipschitz estimate detects a single coordinate in that block.

Definition A.9 (Score of a block of coordinates)**.** Fix a mesh of depths $0, 1/N, \ldots, 1$, with N chosen below. On an h-row chain E, draw a uniform tuple Y on the torus of its last row. In each earlier row draw independent offsets $\boldsymbol{o}_i \in [0, 1]^{I_i}$ with uniform entries, independently of Y. Fix one mesh depth t_i for each earlier row. Starting with $X_{\max E} = Y$, set, for consecutive rows $i < j$ of E,

$$X_i = X_j \circ \pi_{ij} + t_i \boldsymbol{o}_i \pmod 1. \tag{9}$$

One test H is one component of $F_i(X_i)$ in a fixed orthonormal basis of $\mathbb{R}^{q+1}$, for one row i and one choice of depths. For a set B of terminal coordinates, let H^B be the same test with Y replaced by $Y + \frac{1}{p}\mathbf{1}_B$, keeping all offsets fixed. Put

$$I_E(B) = \max_H \mathbb{E}\,|H^B - H|.$$

The number of tests depends only on h, q, N. The score depends only on E and B, so the same face receives the same scores in every chain.

Lemma A.10 (Some face detects its block)**.** *Consider an $(h+1)$-row chain whose prefixes satisfy* (8)*. Partition its terminal coordinates into blocks indexed by the rows, allowing empty blocks. Delete row j to obtain E_j, and let B_j be the inverse image of block j in the last row of E_j. If $h > q$, then $I_{E_j}(B_j) \geq c$ for some j, where $c > 0$ and N depend only on p, q, C, h, ρ.*

Proof. First fix arbitrary offsets on *all* rows. Form a star with p arms indexed by $\mathbb{Z}_p$, each subdivided at the mesh depths, and take the product of $h + 1$ such stars with at least one factor at its tip. In row j write t_j for the depth and a_j for the arm. At depth zero the arm can be chosen freely. This space is the join of $h + 1$ copies of $\mathbb{Z}_p$, with join weights $t_j / \sum_k t_k$. For $p = 2$ it is the boundary of an $(h + 1)$-cube. Adding one to every arm acts freely.

Pull the block labels back to each row, writing $\beta_i(v)$ for the block of $\pi_{i,h+1}(v)$. For completed arm choices, the input of row i has entries

$$x_i(v) = \frac{a_{\beta_i(v)}}{p} + \sum_{k=i}^{h+1} t_k o_{k,\pi_{ik}(v)} \pmod 1. \tag{10}$$

At each vertex let i be the first positive-depth row. Copy its arm to all zero-depth rows and label the vertex by $F_i(x_i)$. Simultaneous rotation of the arms applies T to the label.

We claim that, for some vertex and completion, increasing one zero-depth arm a_j by one changes $F_i(x_i)$ by at least a fixed $\delta > 0$. To see this, take a product cell and choose one arm for each factor, using its arm in the cell whenever that factor can have positive depth. At the lower-depth corner let k be the first positive row and $y = x_k$ its input on the chosen arms. At any vertex of the cell its first positive row satisfies $i \le k$. Change the copied arms of its zero-depth rows to the chosen ones. This takes at most $h\lfloor p/2 \rfloor$ single turns in either direction. A backwards turn has the same absolute change as the forward turn from its endpoint. Then lower the depths to the corner. Each argument changes by at most $(h+1)/N$, since every offset lies in $[0,1]$. The resulting input in row i is $y \circ \pi_{ik}$, as all depths before k now vanish.

Choose δ and N so that $h\lfloor p/2 \rfloor \delta + C(h+1)/N < \rho$. If every single turn had change less than δ, each vertex label would be within ρ of its corresponding $F_i(y \circ \pi_{ik})$. By the common-input margin (8), the convex hull of the cell's labels would miss zero. Multilinearly interpolate the vertex labels on each product cell and normalize. These maps agree on common faces and give an equivariant map from the join to S^q. The join is $(h-1)$-connected, and S^q is free of dimension $q < h$, contrary to Dold's theorem [39] (see also [37, Theorem 5.14]). This proves the claim.

Now draw all offsets independently and uniformly. There are boundedly many choices of depths, arms, turned row j, and scalar component. For every draw one of these choices gives an absolute change at least $\delta/\sqrt{q+1}$. Hence one fixed choice has mean change at least some $c > 0$ depending only on the stated parameters.

For this choice let k be the first row of depth one. It is at or after the observed row i. In (10), x_k is a uniform tuple independent of the offsets before k, since it contains the full independent offset $\boldsymbol{o}_k$. Turning a_j shifts this same tuple by $1/p$ precisely on block j in row k. Delete row j, whose depth is zero. Use the prescribed depths in (9) up to row k, and choose any unused depths. If k is now terminal, $X_k = Y$ is uniform. Otherwise $t_k = 1$, so its full offset again makes X_k uniform. Turning B_j in Y induces exactly the same shift in row k, and both evaluations use the same earlier offsets. Thus the pair of values at row i has the law of (H, H^{B_j}) on the face E_j. Its mean change is at least c. □

Lemma A.11 (A block turn detects a coordinate)**.** *Let H be real and L-Lipschitz on a finite product of unit circles and unit intervals with the maximum metric. Give the circles circumference one and all coordinates independent uniform measure. For $p \ge 2$ an integer and B a set of circle coordinates, put*

$$I_B = \mathbb{E}\big|H(Y + \mathbf{1}_B/p) - H(Y)\big|, \qquad b_j = I_{\{j\}}.$$

Then $\sum_j b_j \le L/p$. For every $c > 0$ there is $\tau = \tau(p, L, c) > 0$ such that $I_B \ge c$ implies $b_j > \tau$ for some $j \in B$.

The threshold assertion also follows from Austin's junta theorem [35, Theorem 1.1]. Pull circle coordinates back to $[0,1]$ and approximate H in L^1 to error less than $c/4$ by a function of a bounded set J of coordinates. Turning $B \setminus J$ changes H in mean by less than $c/2$, so $I_B \le c/2 + \sum_{j \in B \cap J} b_j$. We give a direct proof of both estimates.

Proof. The maximum metric gives $\sum_j |\partial_j H| \le L$ almost everywhere by testing the signs of the derivatives. Integrating along each coordinate turn gives $b_j \le p^{-1}\mathbb{E}|\partial_j H|$, proving the sum bound. We may take $L > 0$ and $B \neq \varnothing$.

Choose a multiple m of p with $4L/m < c/2$. In each coordinate divide into m equal bins. Move one bin left or right with probability $1/4$ each, keeping the bin otherwise and retaining outward steps at interval endpoints. Average uniformly within the resulting bin. The product averaging operator S moves no coordinate farther than $2/m$, so $\|H - SH\|_\infty \le 2L/m$.

We claim $\|Sf\|_2 \le \|f\|_r$ for some $1 < r < 2$ depending only on m. The two bin matrices K are symmetric, lazy and connected, so their L^2 norms on mean-zero vectors are bounded by some $\sigma < 1$. By positivity and homogeneity it suffices in one coordinate to consider bin values $v \ge 0$ of mean one, hence $v \le m$. Taylor's inequality for x^r at 1, using $(x^r)'' \ge r(r-1)m^{r-2}$ on $(0, m]$, gives

$$\|v\|_r^2 = (\mathbb{E}v^r)^{2/r} \ge 1 + (r-1)m^{r-2}\operatorname{Var}(v) \ge \|Kv\|_2^2$$

when r is close enough to 2 that $(r-1)m^{r-2} \ge \sigma^2$. Bin averaging contracts L^r by Jensen. Iterating the one-coordinate bound using the mixed-norm Minkowski inequality proves the claim.

Write T_j for the coordinate turn, $P_j = p^{-1}\sum_{a=0}^{p-1} T_j^a$ for its orbit average, and $G = SH$. Since $p \mid m$, T_j commutes with S. Telescoping the powers of T_j gives

$$2\|G - P_j G\|_2 \le (p-1)\|G - T_j G\|_2 \le (p-1)\|H - T_j H\|_r \le (p-1)(L/p)^{1-1/r} b_j^{1/r}.$$

The P_j are commuting orthogonal projections. Their product over B is fixed by the block turn, and its successive averaging differences are orthogonal. Consequently

$$I_B \le 4L/m + 2\Big(\sum_{j\in B} \|G - P_j G\|_2^2\Big)^{1/2} \le 4L/m + (p-1)(L/p)^{3/2-1/r}\big(\max_{j\in B} b_j\big)^{1/r-1/2},$$

where we used $\sum_j b_j \le L/p$. The last exponent is positive, so choosing τ sufficiently small proves the assertion. □

Proof of Theorem A.6. Each scalar test in (9) is hC-Lipschitz for the maximum metric on its independent inputs. By Lemma A.11, there is a uniform threshold $\tau > 0$ such that any block of score at least c contains a coordinate v with $I_E(\{v\}) > \tau$. Only boundedly many singletons exceed this threshold for each test, and there are boundedly many tests. Assign E the list of all such coordinates in its last row.

If the transported lists on an $(h+1)$-row chain had empty intersection, assign each terminal coordinate to a face whose transported list misses it. The resulting blocks have B_j disjoint from the list on E_j for every j, contradicting Lemma A.10. Every list is nonempty as well, by appending a copy of the last row with the identity minor. This gives the required chain homomorphism. □

A.2 Combinatorial proof

In this section we prove the chain condition using only the finite clock. Each face tests how far one must move to see a given color. Cyclic covering forces a distance to change when one of the blocks turns, so that block must meet the list of its face. For $p = 2$ the turn is antipodal and the covering argument is Tucker's lemma.

A.2.1 Reduction to the clock

In this section we map the clock to the tight cycle and deduce the restriction on the colors of rotated boxes that we will use throughout the proof.

Fix a prime p and $L > 1$ with $L \equiv 1 \pmod p$, and assume $\mathrm{CYCLE}_p[L] \to \mathrm{NAE}_p[q]$. Write $s = q(p-1)+1$, the number of rows on which the face lists must intersect. Arrange N positions in L equal consecutive blocks around a circle, with the block size a multiple of p. The clock H has vertex set $\mathbb{Z}_N$. A tuple $(x^0, \dots, x^{p-1})$ belongs to its relation if every step, including the last one back to the start, satisfies

$$x^{a+1} - x^a \equiv -N/p + e \pmod N \qquad \text{for some integer } e \text{ with } |e| \le N/(pL).$$

Here a is taken modulo p. Each step is a backward $1/p$ turn with an error of at most $1/p$ of a block. For $p = 2$ this joins each position to the closed arc of one block length centered at its antipode.

Lemma A.12 (Clock sandwich)**.** *Number the blocks $0, \dots, L-1$ in cyclic order. Sending block j to $pj \bmod L$ gives a sandwich $H \to \mathrm{CYCLE}_p[L] \to \mathrm{NAE}_p[q]$.*

Proof. Write $L = p\ell + 1$. A backward $1/p$ turn moves $\ell + 1/p$ block lengths. With the allowed error it moves between ℓ and $\ell + 2/p \le \ell + 1$ block lengths, so the block index decreases by ℓ or $\ell + 1$. The label therefore increases by 1 or $1 - p$ modulo L. If m of the p steps have increment $1 - p$, closure gives $p(1-m) \equiv 0 \pmod L$. Since p is coprime to L and $0 \le m \le p < L$, we have $m = 1$. Starting just after that step, the labels are p consecutive residues, as required. The second homomorphism is the hypothesis of the theorem. □

Take at least $4p^2s^2$ positions per block.[3] For $x \in \mathbb{Z}_N^I$, let $B(x)$ be the box of radius $N/(2pL)$ about x in the maximum circle distance. Every clock polymorphism $f : H^I \to \mathrm{NAE}_p[q]$ satisfies

$$\bigcap_{a \in \mathbb{Z}_p} f(B(x - aN/p)) = \varnothing, \tag{11}$$

where the shift acts in every coordinate. Indeed, choose one input from each box. In every coordinate the successive steps, including the closing step, differ from $-N/p$ by at most twice the box radius, namely $N/(pL)$. Each column is therefore a clock tuple, so the p outputs cannot all have the same color.

Theorem A.13 (Clock)**.** *The polymorphisms of $(H, \mathrm{NAE}_p[q])$ admit an $(s-1, d)$-chain homomorphism to the projections, with d depending only on p, q, N.*

Proof of Theorem 2.1. Precomposition with the sandwich map $H \to \mathrm{CYCLE}_p[L]$ preserves minors, so Theorem A.13 gives the required chain homomorphism, and Theorem 1.15 gives the reduction and search decoder since a constraint on the tuple $(x, \dots, x)$ is a fixed NO instance for $\mathrm{NAE}_p[q]$. □

[3] To obtain the clock theorem for a coarser clock, double each block as often as needed. Rounding z down to $\lfloor z/2 \rfloor$ sends each doubled clock to its predecessor by a homomorphism, since both the prescribed step and its integral error bound double. Precomposition transfers the lists.

A.2.2 One coordinate list per face

In this section we choose a bounded list of terminal coordinates for each face. Turning a block of coordinates outside the list rarely changes the distance to any color class by more than two clock steps per row.

Lemma A.14 (Predicting a nearby color)**.** *Let $m \geq 2$, $q \geq 1$ and $\varepsilon > 0$. For every finite index set I and coloring $g : [m]^I \to [q]$, there are a set $J \subseteq I$ of size bounded in terms of m, q, ε alone and a* predictor $\lambda : [m]^J \to [q]$ *such that, for uniform $z \in [m]^I$,*

$$\Pr[\ \exists z' \in [m]^I : \quad |z'_i - z_i| \leq 1 \text{ for all } i \in I, \quad g(z') = \lambda(z|_J)\ ] \ \geq\ 1 - \varepsilon.$$

Consider a face with $h = s - 1$ rows, numbered $1, \ldots, h$ within the face. Write I_i for the coordinates of row i and $\pi_{ij} : I_i \to I_j$ for the composite minor maps, with π_{ii} the identity. For each color c let $d_{i,c}(x)$ be the distance from x to $f_i^{-1}(c)$ in the maximum of the circular coordinate distances, measured in clock steps. Give an absent color distance N. Moving every argument by at most one step changes each distance by at most one.

Draw a uniform terminal tuple $Y \in \mathbb{Z}_N^{I_h}$ and independent uniform offsets $o_{j,w} \in \{0, \ldots, N-1\}$ for $j < h$ and $w \in I_j$. Give each earlier row a *depth* $t_j \in \{0, \ldots, N-1\}$. Starting from $X_h = Y$, work backwards in the clock by

$$(X_j)_v = (X_{j+1})_{\pi_{j,j+1}(v)} + \min\{t_j, o_{j,v}\} \pmod N.$$

At zero depth a row contributes nothing. At full depth $N - 1$ it adds an independent uniform tuple.

Lemma A.15 (A face test)**.** *For every $\varepsilon > 0$, the face has a list $\Xi \subseteq I_h$ of size bounded in terms of s, N, q, ε. For every block $B \subseteq I_h$ missing Ξ, compare the experiment with the same one after adding N/p to Y on B. Except with probability 2ε, the two readings of $d_{i,c}(X_i)$ differ by at most $2h$, simultaneously for every depth profile, row i and color c.*

Proof. Color the independent parameters (Y, o) by the table of all the readings, including all choices of the depths. Its entries belong to $\{0, \ldots, N\}$ and its length is at most hqN^{h-1}, independent of the arities. Apply Lemma A.14 once to this finite coloring, representing the clock positions by $0, \ldots, N-1$. At a *witness* for the predicted table, each parameter moves at most one path step. In row i the terminal entry and at most $h - 1$ offset terms each change by at most one, so the predicted and actual distances differ by at most h.

Keep in Ξ just the terminal entries read by the predictor, adding one arbitrary terminal coordinate if needed. Turning B preserves the uniform law and leaves all predictor reads unchanged, including its earlier offset reads. With probability at least $1 - 2\varepsilon$ both tables are within h of the same prediction in every entry. □

The proof of the nearby-color lemma is in Section A.2.5.

A.2.3 Intersection of the face lists

In this section we show that every partition of the terminal coordinates has a block whose turn changes a color distance. The face test then forces that block to meet its list.

The *star* S is the union of p paths of $N - 1$ edges with one endpoint of each path identified to a common center. The paths are its *arms*, indexed by $\mathbb{Z}_p$. The center has depth zero. On arm a, write (a, t) for the vertex at depth $t \in [N-1]$, its distance from the center. The vertex $(a, N-1)$ is the tip. A *cell of* S^s is a product whose factors are vertices or edges of S. Its vertices are obtained

by choosing a vertex in each factor. Let $\mathcal{X}$ consist of the cells with at least one factor equal to a single tip. Rotation T adds one to every arm index and fixes each center. For $p = 2$, the star is a path and $\mathcal{X}$ is the boundary of a subdivided s-cube, with T acting by negation about its center.

Lemma A.16 (Cyclic covering)**.** *For every coloring of the vertices of $\mathcal{X}$ by $[q]$, some cell $D \in \mathcal{X}$ has a color occurring at a vertex of each of $D, T^{-1}D, \dots, T^{-(p-1)}D$.*

We prove the cyclic covering lemma in Section A.2.4.

Conclusion of the proof of Theorem A.13. Choose the lists of Lemma A.15 with $2\varepsilon < (sN^s p^s)^{-1}$. If their transported intersection is empty, label each terminal coordinate by a row whose deletion gives a list missing that coordinate. Let $\beta_i : I_i \to [s]$ be the pullback of these labels to row i.

Give each row independent uniform offsets $o_{j,w}$ as in the face test, now including the last row. At a vertex of $\mathcal{X}$, an *arm completion* $\alpha \in \mathbb{Z}_p^s$ keeps its positive-depth arms and chooses an arm at each center. For the first positive row i, the prescribed input is

$$x_i(t,\alpha)_v = \alpha_{\beta_i(v)} N/p + \sum_{j=i}^{s} \min\{t_j, o_{j,\pi_{ij}(v)}\} \pmod N. \tag{12}$$

Changing the arm at a center turns just its block in this same row.

Fix the offsets. Suppose adding one to the arm at any center changes each color distance by at most $2(s-1)$. At each vertex choose one completion and color it by the actual value $c = f_i(x_i(t,\alpha))$. Its distance to color c is zero. Reach any other completion by changing the missing arms one generator step at a time. At most $(s-1)(p-1)$ steps are needed, so at the new input the distance to c is at most $2(s-1)^2(p-1)$.

By Lemma A.16, one cell has a color c on every rotation. Fix arms containing that cell, also at its centers, and use the rotated arms on its rotations. In each rotation take a vertex colored c and pull its color-c witness back along π_{1i} to get the same color in row one. Its distance from the prescribed row-one input is no greater, since pullback just repeats coordinates and all depths before i are zero. With the arms fixed, each depth varies by at most one across the cell, so the prescribed inputs vary by at most s steps. On rotation by T^{-a} all prescribed inputs shift by $-aN/p$. Thus for a row-one input x in the chosen cell, color c occurs within

$$2(s-1)^2(p-1) + s \;\le\; 2ps^2 \;\le\; N/(2pL)$$

of $x - aN/p$ for every a. This contradicts (11).

Consequently every draw of the offsets has a turn changing some distance by more than $2(s-1)$. There are at most $sN^s p^s$ choices of the moved row, depths and starting arms. One fixed choice has such a change with probability at least $(sN^s p^s)^{-1}$.

Delete the moved row, whose depth is zero, and keep the prescribed depths from i through the first full row k. In the last row of the face, take the inverse image of the moved block under its map to row s. The inputs in (12) before and after the turn have exactly the joint law of the two row-i inputs in the face test. Indeed the full row supplies a uniform tuple on I_k, absorbing all later contributions and the starting phases conditionally on them. If k ends the face, the terminal tuple Y supplies that uniform tuple directly. The offsets before k remain independent, and the relative turn in both experiments is the pullback of the moved block. This block misses the face list by its label. Lemma A.15 bounds the probability of such a change by 2ε, a contradiction. □

A.2.4 Cyclic covering

In this section we prove Lemma A.16 using the cyclic form of Tucker's lemma. If each color misses a rotation of each cell, the colors give a labeling that violates Tucker's conclusion.

Proof. The join of s copies of $\mathbb{Z}_p$ has vertices $(j, a) \in [s] \times \mathbb{Z}_p$. Its simplices are the sets containing at most one vertex from each copy. Rotation adds one to a. On the label set $\mathbb{Z}_p \times [s-1]$, rotation adds one to the first entry and fixes the second. A map is *equivariant* if it commutes with these rotations.

The cyclic form of Tucker's lemma states that on any rotation-invariant subdivision of the join, every equivariant vertex labeling by $\mathbb{Z}_p \times [s-1]$ has an edge whose labels differ in their first entries and agree in their second. Otherwise the labeling would give an equivariant simplicial map to the join of $s-1$ copies of $\mathbb{Z}_p$. Such a map is excluded by [37, Corollary 5.10], applied with $d = s-2$ and $m = s-1$, since a sequence of $s-1$ labels has at most $s-2$ changes. For $p = 2$ this is ordinary Tucker's lemma. The subdivision used here refines the octahedral sphere and hence preserves the flag of coordinate hemispheres in [38, Theorem 2].

Use the barycentric subdivision of $\mathcal{X}$. Its vertices are the nonempty cells, and its simplices are chains of cells under inclusion. To see that it subdivides the join, realize each arm as $[0, N-1]$. A point of $\mathcal{X}$ then has depths $t_j \geq 0$ with $\max_j t_j = N-1$. Give vertex (j, a) of the join weight $t_j / \sum_k t_k$ if the point lies on arm a in factor j, and zero otherwise. Zero depth gives zero weight on every arm. The inverse rescales the weights so that the largest is $N-1$. These maps commute with rotation. On each simplex, normalization is radial projection from a hyperplane with one depth fixed at $N-1$ to the hyperplane where the weights sum to one. It sends the simplex onto the convex hull of its vertex images.

Suppose no cell has a color on all its rotations. For each cell D put

$$P(D) = \{(i, b) \in \mathbb{Z}_p \times [q] : \text{some vertex of } T^{-i}D \text{ has color } b\}.$$

The set is nonempty. Each color misses at least one rotation, so $|P(D)| \leq q(p-1) = s-1$. Rotation of D adds one to the first entries of $P(D)$. If a nonzero shift fixed this set, primality of p would put any color present in all p rotations. Every orbit of these sets therefore has size p. In each orbit, number one set 0 and its successive shifts $1, \ldots, p-1$. Label D by the number of $P(D)$ and by $|P(D)|$. This labeling is equivariant. Adjacent vertices of the subdivision represent nested cells, so their sets $P(D)$ are nested. If their sizes agree, the sets and their numbers agree. This contradicts Tucker's lemma. □

A.2.5 Predicting a nearby color

In this section we prove Lemma A.14 by placing the input in a random product of vertices and edges of a path. Every point of this product is within one step of the input in each coordinate, so it suffices to predict a color that occurs in the product. The lemma also follows from Austin's junta theorem [35, Theorem 1.1] by representing the m symbols as midpoints of equal intervals and approximating the distance to each nonempty color class on the unit cube.

Proof. Let $g : [m]^I \to [q]$ be the given coloring. When I is empty, g is constant, so assume $n = |I| \geq 1$. Give $[m]$ the path edges $\{j, j+1\}$ for $j \in [m-1]$. Let P consist of these edges and the singletons $\{j\}$ for $j \in [m]$. A *box* is a product of factors $z_i \in P$ for $i \in I$. We identify it with the tuple $z \in P^I$. We adapt Russo's estimates for pivotal sets of increasing Bernoulli events and the sequential coupling in his proof of Proposition 2 [36, Sections 3–5].

Choose a uniform path vertex v and then a uniform incident edge e. Each incident pair has probability at least $1/(2m)$. *Opening a vertex* replaces $\{v\}$ by e. Open with probability t and otherwise keep $\{v\}$. Make these choices independently in each coordinate to obtain a random box Z. Write $\Pr_t$ and $\mathbb{E}_t$ for probability and expectation under this choice. For $b \in [q]$, let H_b be the event that the box contains a vertex of color b. These events are increasing under inclusion and cover all boxes.

Finding the coordinates. For $z \in H_b$, a *pivotal coordinate* for b is an index i at which some replacement of z_i by an element of P removes the box from H_b. Let $K_b(z)$ be the set of such indices, and put $K_b(z) = \varnothing$ outside H_b. Once a growing box enters H_b, K_b can only shrink. Indeed, if every replacement at i leaves color b present, this remains true when the other factors grow. We will show that for each $\delta > 0$ there are $t \in [1/4, 1/2]$ and $J \subseteq I$, with size bounded in terms of m, q, δ, such that

$$\sum_{b \in [q]} \Pr_t(K_b \not\subseteq J) \le \delta. \tag{13}$$

For $G \subseteq P^I$, define $I_t(G)$ as follows. At coordinate i draw the incident pair (v, e) above, independently of the other factors of Z. Sum over i the probability that changing $\{v\}$ to e changes membership in G.

First, suppose the events in a finite family $\mathcal{G}$ change membership at most k times in total along any increasing sequence of boxes. Then

$$\operatorname*{avg}_t \sum_{G \in \mathcal{G}} I_t(G) \le 8k, \qquad t = \frac{1}{4} + \frac{j}{4n}, \quad 0 \le j < n. \tag{14}$$

At each coordinate independently choose an incident pair and an opening round uniformly in $[4n]$, independently of the pair. After round $4nt$ each vertex has opened independently with probability t. Coordinate i opens in the next round with probability $1/(4n)$. Conditional on the other factors at round $4nt$, each still-closed coordinate opens in the next round with probability $1/(4n(1-t)) \le 1/(2n)$. Thus with probability at least $1/2$ none of the other coordinates opens. Counting only rounds when i opens alone and summing over i gives at least $\sum_G I_t(G)/(8n)$ expected changes in this round. The total over the n indicated rounds is at most k.

Second, writing $a = \Pr_t(G)$ and $0 \log(1/0) = 0$, we have

$$a \log(1/a) \le 2m I_t(G). \tag{15}$$

Fix a coordinate i. For $u \in P$ let a_u be the probability of the section of G obtained by setting that coordinate to u. Then $a = \mathbb{E}_t[a_{Z_i}]$. If $a > 0$,

$$a \log(1/a) - \mathbb{E}_t[\, a_{Z_i} \log(1/a_{Z_i})\,] \le a \log\bigl(\max_u a_u / a\bigr) \le \max_u a_u - \min_u a_u.$$

Join a minimizing factor to a maximizing factor along the path of alternating vertices and edges. Across each incident pair the difference of section probabilities is at most the probability that their indicators differ on the same remaining factors. Each pair has probability at least $1/(2m)$, so the range is at most $2m$ times this coordinate's contribution to $I_t(G)$. Apply the same bound inductively to each section. Averaging over u gives the contributions of all remaining coordinates, proving (15). The empty product has probability zero or one for every event, which starts the induction.

With all other factors fixed, a pivotal coordinate has a nonempty proper section of H_b. On that section at least one incident pair changes membership. Hence $\mathbb{E}_t|K_b| \le 2m I_t(H_b)$. A growing box enters H_b at most once, so this expectation averages at most $16m$ by (14).

For an integer $r \geq 1$, put $a_{b,i} = \Pr_t(i \in K_b,\ |K_b| \leq r)$. At the first box in H_b with $|K_b| \leq r$, at most r indices remain in K_b. Each can leave only once. Thus for fixed b the events defining $a_{b,i}$ change membership at most $2r$ times in total. Equations (14) and (15) show that $\sum_i a_{b,i} \log(1/a_{b,i})$ averages at most $32mr$.

For some common t in (14), summing these two estimates gives

$$\sum_b \mathbb{E}_t|K_b| + \frac{1}{r}\sum_{b,i} a_{b,i}\log(1/a_{b,i}) \leq 48mq.$$

For $0 < \tau < 1$, put i in J if $\Pr_t(i \in K_b) \geq \tau$ for some b. The first sum gives $|J| \leq 48mq/\tau$. A box with $K_b \not\subseteq J$ either has at least $r+1$ pivotal coordinates or has $i \in K_b \setminus J$ with $|K_b| \leq r$. For each such i we have $a_{b,i} < \tau$, so $a_{b,i}\log(1/a_{b,i}) \geq a_{b,i}\log(1/\tau)$. Summing over colors gives

$$\sum_b \Pr_t(K_b \not\subseteq J) \leq 48mq\left(\frac{1}{r+1} + \frac{r}{\log(1/\tau)}\right).$$

Choose r large and then τ small to obtain (13).

Predicting the color. Given the input on J, open its vertices independently with probability t. For a uniform input the resulting factors have the same distribution as $Z|_J$. Given their values z, choose a color b maximizing $\Pr_t(H_b \mid Z|_J = z)$. This probability is at least $1/q$, since every box contains a color. Condition Z on H_b and $Z|_J = z$, and couple its remaining factors with uniform input vertices outside J. On J the input already lies in the chosen factors.

For one coordinate, order P as $\{1\}, \{1,2\}, \{2\}, \ldots, \{m\}$. Partition a list of equally likely tickets into consecutive blocks of sizes proportional to the probabilities of these factors before conditioning. Partition the same list into m equal blocks for the input vertices. The cut between input vertices j and $j+1$ lies inside the block for $\{j, j+1\}$. The mass of that block on either side of the cut comes from opening the incident vertex to this edge, so it is at least $t/(2m) \geq 1/(8m)$ on each side. Now change the factor law. If the mass of each initial segment of the ordered set P changes by at most $1/(8m)$, every cut still lies in its edge block, possibly at an endpoint. Thus a single uniform ticket gives a factor with the specified law and a uniform input vertex contained in that factor. All laws used here are rational, so a finite common number of tickets suffices.

Expose the factors of Z outside J in a fixed order. Given the factors already exposed, let u_i be the conditional probability that the next index i belongs to $K_b(Z)$. With all other factors fixed, the factor at i has its original law conditioned on their section of H_b. This section is nonempty. It is proper exactly when $i \in K_b(Z)$, so only in this case can the law change. Averaging over the unexposed factors, the probability of any initial segment of P differs from its unconditioned probability by at most u_i. Use a fresh ticket for this law and the uniform input vertex. The new input vertex is uniform independently of all previous draws. It belongs to its factor whenever $u_i \leq 1/(8m)$.

At the first prefix with $u_i > 1/(8m)$, the conditional probability of $K_b \not\subseteq J$ exceeds $1/(8m)$. These first prefixes are disjoint. Their total probability is therefore at most $8m \Pr_t(K_b \not\subseteq J \mid H_b, Z|_J = z)$. In every other outcome the whole input lies in Z, which contains a vertex of color b. Averaging over z and using $\Pr_t(H_b \mid Z|_J = z) \geq 1/q$ bounds the error by $8mq\sum_b \Pr_t(K_b \not\subseteq J) \leq 8mq\delta$. Take $\delta = \varepsilon/(8mq)$. Finally, for each fixed input on J, choose the color most likely to occur within one step in every coordinate of the full uniform input conditional on these values. This deterministic predictor succeeds at least as often as the randomized choice above. □

B Auxiliary results for hypergraph colouring

The first lemma extends the tight-cycle criterion to composite arities. A periodic word need not have the specified period as its least period.

Lemma B.1 (Coprime periods)**.** *Let $r \geq 2$ and let $T \subseteq A^r$ be cyclically invariant, let $T \to S \subseteq B^r$ be a finite promise template with no constant tuple in S, and suppose a periodic word of period s has all its length-r windows in T, where $\gcd(r, s) = 1$. Then 3-SAT $\leq_{\mathrm{L}}$ PCSP(T, S) with a LOGSPACE search decoder. Moreover, for every prime $p > \max\{2r, r + (r-1)s\}$, there is a finite d with $\mathrm{Ch}_{|B|(p-1),d}(\mathrm{Pol}(T, S), \mathscr{P})$. If r itself is prime, the height can instead be $|B|(r-1)$.*

Proof. We first construct a periodic sequence of length $n = p - r$ with all its r-windows in T. Form a digraph on the tuples of T by joining $(a_1, \ldots, a_r)$ to $(a_2, \ldots, a_r, b)$ whenever the latter tuple belongs to T. The windows of the given sequence form a closed walk of length s. At the same starting tuple its successive rotations form a closed walk of length r. Choose $0 \leq j < r$ with $js \equiv n \pmod r$. Since $n > (r-1)s$, concatenating j of the first walks and $(n - js)/r$ of the second gives a closed walk of length n. Let a_i, indexed modulo n, be the first letter of its i-th tuple. Overlap of successive tuples gives $(a_i, \ldots, a_{i+r-1}) \in T$ for all i.

For any r-ary relation U, put

$$V_p(U) = \{x : (x_j, x_{j+1}, \ldots, x_{j+r-1}) \in U \text{ for each } j \in \mathbb{Z}/p\mathbb{Z}\},$$

where x ranges over length-p words on the domain of U. Every p-letter window $x = (a_i, \ldots, a_{i+p-1})$ lies in $V_p(T)$. Ordinary r-windows inside it are legal by construction. A window crossing the end consists of the last k entries of x and its first $r - k$, for some $0 < k < r$. Since $p \equiv r \pmod n$, these entries form a rotation of $(a_i, \ldots, a_{i+r-1})$.

Align successive p-letter windows by placing the letter at position j in slot $j \bmod p$. Consecutive tuples of $V_p(T)$ now differ in at most one coordinate. After n steps the tuple is rotated by n places. Since $\gcd(n, p) = 1$, repeating the walk gives a one-place rotation. Both $V_p(T)$ and $V_p(S)$ are cyclically invariant, the original homomorphism takes the first into the second, and $V_p(S)$ contains no constant tuple. Apply Theorem 2.2 to this pair. Every polymorphism of (T, S) preserves the consecutive-window conjunction, hence belongs to $\mathrm{Pol}(V_p(T), V_p(S))$. Restricting the certificate proves the claimed bound. Equivalently, replace each V_p-constraint by its p windows on the same variables. This fixed LOGSPACE replacement preserves both strong and weak solutions, including the values needed by the search decoder.

If r is prime, repeat the given word to a length congruent to one modulo r and apply Theorem 2.1 directly. This gives the height $|B|(r-1)$. □

Lemma B.2 (A fixed colour count)**.** *Suppose every tuple of $R \subseteq A^r$ has exactly k entries in a fixed set $C \subseteq A$, where $0 < k < r$. Then both decision and search for PCSP$(R, \mathrm{NAE}_r[q])$, $q \geq 2$, are solvable in polynomial time, with two output colours.*

For finite templates the decision statement also follows from the symmetric polymorphism criterion [18, Theorem 3.2]. For each $n \equiv 1 \pmod r$, output 1 when more than nk/r inputs belong to C, and 0 otherwise. On n tuples of R the r column counts average nk/r, a noninteger, so some column gives each output. These symmetric polymorphisms have unbounded arity. The following proof also constructs the colouring.

Proof. For every input constraint e impose $\sum_{v \in e} y_v = 0$ over $\mathbb{Q}$, counting repeated variables with their multiplicities. A strong colouring supplies a solution whose entries are $r - k$ on C and $-k$

elsewhere. Conversely, any kernel vector with all coordinates nonzero has both signs on every constraint, so its signs give the required two-colouring.

Compute a rational kernel basis $b_1, \dots, b_d$. If some coordinate is forced to zero, reject; such an input has no strong colouring. Otherwise, on $n > 0$ variables, one of the first $n(d-1)+1$ positive integers t makes every entry of $\sum_{j=1}^{d} t^{j-1} b_j$ nonzero, since each coordinate polynomial excludes at most $d-1$ choices. Gaussian elimination and these evaluations have polynomial bit complexity. An input without variables is immediate. □

Lemma B.3 (Distinct vertices)**.** *Let $R \to S \subseteq B^r$ be a fixed finite promise template, with $r \geq 2$ and $B \neq \varnothing$. There is a LOGSPACE reduction from* $\mathrm{PCSP}(R, S)$ *to its restriction to instances in which every constraint uses pairwise distinct variables, with a LOGSPACE search decoder.*

Proof. Use the variable-copy construction of Brakensiek and Guruswami [23, Theorem 6.15]. Replace each vertex by $q(r-1)+1$ copies, where $q = |B|$. Replace each constraint by all versions using pairwise distinct copies in its ordered positions. Strong colourings extend by copying values. In a weak colouring, choose for each original vertex a colour occurring on at least r of its copies. Every original constraint can then be read on distinct copies of those colours. This construction and decoder use LOGSPACE for fixed parameters. □

C Constructions for smooth digraphs

The proof of Theorem 2.5 uses two constructions. First, a path power reduces one weak component to algebraic length one. We then construct a digraph D that both admits the cyclic certificate from Theorem 2.2 and maps to the given component after finitely many arc-digraph iterations.

C.1 Transferring hardness through arc digraphs

The palette bound is due to Harner and Entringer [15]. Its adjunction formulation is [4, Observation 4.14]. We give the decoder explicitly.

Lemma C.1. *Let D, Q, B be finite nonempty digraphs with D smooth, B loopless, and $\delta^t D \to Q \to B$ for some fixed $t \geq 0$. If $\mathrm{PCSP}(D, K_M)$ is hard for every sufficiently large M, then $\mathrm{PCSP}(Q, B)$ is hard. More precisely, a LOGSPACE reduction from 3-SAT with a LOGSPACE search decoder to each sufficiently large-palette problem transfers to $\mathrm{PCSP}(Q, B)$.*

Proof. Choose $n \geq \max(2, |V(B)|, |V(D)|)$ large enough, set $n_t = n$, and put $n_i = 2^{n_{i+1}}$ for $i = t-1, \ldots, 0$. If $f : \delta X \to K_c$, assign to each vertex v of X the set of colors on arcs leaving v. For an arc $u \to v$, its color belongs to the set at u and to none of the colors on arcs leaving v. The two sets therefore differ, giving $X \to K_{2^c}$. Together with $X \to D \Rightarrow \delta X \to \delta D$, this proves

$$\begin{aligned}\mathrm{PCSP}(D, K_{n_0}) &\leq_{\mathrm{L}} \mathrm{PCSP}(\delta D, K_{n_1}) \leq_{\mathrm{L}} \cdots \\ &\leq_{\mathrm{L}} \mathrm{PCSP}(\delta^t D, K_n) \leq_{\mathrm{L}} \mathrm{PCSP}(Q, B).\end{aligned}$$

The last reduction is the identity, using $\delta^t D \to Q$ and $B \to K_n$. All intermediate pairs are promise templates: from $\delta^t D \to Q \to B \to K_n$, the same color-set argument gives $\delta^i D \to K_{n_i}$ for each i. To keep domains nonempty, each arc-digraph step may add an isolated vertex, which does not affect any homomorphism to the nonempty targets. All iteration counts and palettes are fixed. The instance maps enumerate walks of fixed length, and the decoder repeatedly takes outgoing color sets, so both use LOGSPACE. □

C.2 Reducing the algebraic length

Barto, Kozik and Niven [16, Lemma 5.1 and the proof of Theorem 3.1] use the same path power to obtain algebraic length one. We record the statement on a single cyclic class.

Lemma C.2 (Reduction of algebraic length)**.** *Let A be weakly connected and smooth with algebraic length d. Partition its vertices into classes $V_0, \ldots, V_{d-1}$ using a map to $\vec{C}_d$. Let P have vertex set V_0 and an arc for each pair joined in A by a directed walk of exactly d steps. Then P is weakly connected, smooth, and has algebraic length one.*

Proof. For $d = 1$ this is immediate. Suppose $d > 1$. Smoothness follows by taking d steps forwards and backwards. For $x \in V_j$, $j > 0$, let $E(x)$ be the set of endpoints of $(d-j)$-step walks from x. They all lie in V_0. Take a j-step walk from some $z \in V_0$ to x, using smoothness. Every $y \in E(x)$ is a P-outneighbor of z. Thus these endpoints belong to a common weak component of P and have equal images in any homomorphism $f : P \to \vec{C}_e$. Put $E(x) = \{x\}$ on V_0. For an arc $x \to x'$ in A with $x \notin V_0$, the sets $E(x), E(x')$ intersect. When $x \in V_0$, x has a P-arc to every endpoint in $E(x')$. It follows from weak connectedness of A that P is weakly connected.

Given $f : P \to \vec{C}_e$, define a map to $\mathbb{Z}_{de}$ on A by putting $F(x) = df(x)$ on V_0 and

$$F(x) = df(y) - (d-j), \qquad x \in V_j,\ j > 0,\ y \in E(x).$$

It is well defined because all such $f(y)$ agree. It advances by one on each arc. Inside the nonzero classes use a common endpoint. On entry to V_0 the endpoint is the vertex entered, and on departure from V_0 one has $f(y) = f(x) + 1$ along the corresponding P-arc. Thus $A \to \vec{C}_{de}$, forcing $de \mid d$, or $e = 1$. The algebraic length of P is therefore one. □

C.3 Periodic walks and the auxiliary digraph

Fix a finite smooth weakly connected loopless digraph A of algebraic length one. Let $L \geq 2$ be divisible by the lengths of all simple directed cycles of A. A *track* is an L-periodic directed walk $c : \mathbb{Z} \to A$. Write $c \rightsquigarrow d$ if some bi-infinite directed walk in A agrees with $c(t)$ for all sufficiently negative t and with $d(t)$ for all sufficiently positive t. The same time coordinate t fixes the phase on both tracks.

Lemma C.3. *For every track c and integer q, there is a zigzag of tracks from $c(t)$ to $c(t - q)$ in which consecutive tracks are related by $\rightsquigarrow$ in one direction or the other.*

Proof. Give $F = A \times \vec{C}_L$ the arcs $(x, r) \to (y, r + 1)$ for $x \to y$ in A. This lift is weakly connected. In fact at a fixed vertex of A a closed signed walk of algebraic length one shifts the second coordinate by one, and such walks together with lifts of connecting walks reach every point of F.

Each track determines a directed cycle in F, comprising $(c(t), t \bmod L)$. To find such a track in a strongly connected component of F containing a directed cycle, take a vertex (x, r) on that cycle. Its projection puts x on a positive closed walk, hence on some simple directed cycle through x. Traverse that simple cycle repeatedly for exactly L steps. The lifted closed walk is in the same component. Call these components recurrent.

For two specified tracks c, d, the condition $c \rightsquigarrow d$ is equivalent to directed reachability from their respective recurrent components in F. In the forward direction lift the two-sided walk using the time coordinate. In the other direction find a directed path starting on the chosen c-orbit and ending on the chosen d-orbit, using strong connectivity inside the two components. The time coordinate increases by one on every edge. Extending the path to the left by c and to the right by d thus gives precisely the prescribed phases.

The recurrent components form a connected graph when comparable components are joined by an undirected edge. To see this, every vertex v of F has a recurrent predecessor and a recurrent successor, obtained by tracing walks backwards and forwards using smoothness. Every recurrent predecessor of v is comparable with every recurrent successor of v, so all these recurrent components belong to one connected class. An arc $v \to w$ identifies the classes of v and w by comparing a recurrent predecessor of v with a recurrent successor of w. The assertion follows by weak connectedness of F. Apply it to the components containing the two specified tracks, with arbitrary tracks in intermediate recurrent components. □

Fix a residue $q \pmod L$ containing arbitrarily large primes. Such a residue exists by infinitude of primes and the pigeonhole principle, and $\gcd(q, L) = 1$. Choose a track c_0, and use Lemma C.3 to obtain a zigzag with last term $c_h(t) = c_0(t - q)$. Record a sign for each step, according as its allowed transition follows the indexing order or goes against it. We may assume the sequence of signs alternates, including at the cyclic seam. Indeed, at any junction of equal signs, duplicate the intermediate track and insert a step of the opposite sign, using the periodic walk from the track to itself. At the last-to-first junction do the same with a duplicate of c_0 at the start, or of its shifted copy at the end. If needed insert further pairs of duplicates so $h \geq 4$. The result has even h and unchanged net displacement. Extend the tracks to all $j \in \mathbb{Z}$ by $c_{j+h}(t) = c_j(t - q)$.

The auxiliary digraph. Define a finite digraph D with h distinct copies of $\vec{C}_L$. Denote its vertices using

$$\pi(j,z) = (j \bmod h,\ z - \lfloor j/h \rfloor q \bmod L) \qquad (j,z \in \mathbb{Z}).$$

In particular $\pi(j+h,z) = \pi(j,z-q)$. There is always an arc $\pi(j,z) \to \pi(j,z+1)$. Write $\epsilon_j = +1$ at the transition between $j-1$ and j if its sign is forward, so that $c_{j-1} \rightsquigarrow c_j$. Write $\epsilon_j = -1$ otherwise, so that $c_j \rightsquigarrow c_{j-1}$. At each transition add, for every z,

$$\begin{cases} \pi(j-1,z) \to \pi(j,z+1), & \epsilon_j = +1, \\ \pi(j,z) \to \pi(j-1,z+1), & \epsilon_j = -1. \end{cases}$$

These definitions are consistent under $j \mapsto j+h$, $z \mapsto z+q$. In the alternating index ring, every index is a source or a sink for transitions between distinct tracks. A directed walk in D therefore changes tracks at most once. This generalizes the tambourine to an alternating ring of cycles with a phase change at the closing seam.

Recall that vertices of $\delta^t X$ are directed walks of length t, with adjacency given by shifting one step along a walk of length $t+1$.

Lemma C.4. *For some fixed t, $\delta^t D \to A$.*

Proof. On a two-sided D-walk without a switch, the state $\pi(j,z)$ is read as the vertex $c_j(z)$. The phase convention makes this independent of the representation of the state. For each actual cross-edge e fix one two-sided A-walk W_e with the two required periodic tails, placing the D-edge e at times $0,1$. Such a walk exists by the oriented connection between the specified tracks and a common shift of their time parameter. For example, if e goes from $\pi(j-1,z)$ to $\pi(j,z+1)$, the required tails are $c_{j-1}(z+k)$ at time $k \ll 0$ and $c_j(z+k)$ at time $k \gg 0$. There are finitely many edges, so choose one T so large that all these walks already agree with the required tails at $|k| \geq T-1$.

Decode a two-sided D-walk with its unique switch e at times $u, u+1$ by the translated walk $W_e(k-u)$ at time k. If there is no switch use its periodic track. The value at time k depends only on the window $[k-T, k+T]$. A visible switch specifies its connector. If none is visible, the periodic value is correct. Every finite D-walk extends in both directions since its endpoints lie on directed cycles. Assign each length-$2T$ walk the decoded value at its center. Consecutive length-$2T$ walks extend to one common two-sided walk, so the assigned values form an A-arc. Take $t = 2T$. □

C.4 A cyclic relation gadget

Lemma C.5 (Cyclic relation gadget)**.** *For the digraph D constructed above and every $M \geq |V(D)|$, there are a prime p and a finite promise template (R_D, R_{K_M}) of arity p with the following properties:*

(i) *R_D is invariant under cyclic rotation and contains a tuple joined to its rotation by single-coordinate changes inside R_D;*

(ii) *R_{K_M} contains no constant tuple;*

(iii) *a fixed gadget gives a LOGSPACE reduction $\mathrm{PCSP}(R_D, R_{K_M}) \leq_{\mathrm{L}} \mathrm{PCSP}(D, K_M)$, with a LOGSPACE decoder.*

Proof. Set $m = M+1$, $\lambda = 2L-1$ and $C = \lambda + 1$. Pick $P > 4Cm + 2\lambda + 2$ and a prime $p \equiv q \pmod{L}$ with $p > hP$. The gadget Y has vertices (i,r), where $0 \leq i < m$ and $r \in \mathbb{Z}_p$, and arcs

$$(i,r) \longrightarrow (k, r-\lambda) \qquad (i<k).$$

For any digraph B, define a p-ary relation R_B on B^m by taking precisely the p columns of every homomorphism $Y \to B$. The permutation $r \mapsto r+1$ shows cyclic invariance. A constant tuple in R_{K_M} would give a homomorphism from the transitive tournament on m rows to K_M, impossible as $m = M+1$. A relation constraint can be replaced by its displayed arc gadget, with m vertices per variable. An assignment to the relation instance is exactly a homomorphism from the expanded digraph to B. This works for $B = D$ and for $B = K_M$, hence gives the required reduction.

We construct a path to a rotation in R_D. Put $u_j = jP$ for $0 \le j < h$, extend by $u_{j+h} = u_j + p$, and for each row i set

$$b_j(i) = u_j - \epsilon_j C i, \qquad J_i(x) = j \quad \text{when} \quad b_j(i) \le x < b_{j+1}(i).$$

The thresholds are strictly ordered and $J_i(x+p) = J_i(x) + h$. For any integer a define a labeling of Y by

$$w^a(i,r) = \pi\big(J_i(s-a), s\big), \qquad s \in \mathbb{Z}, \quad s \equiv r \pmod{p}.$$

It is independent of the lift s, since adding p raises the track index by h and the phase by $p \equiv q \pmod{L}$, which cancel under π.

Consider an arc from row i to row $k > i$ and use the lifts s and $s - \lambda$. Its phases advance by one since $-\lambda \equiv 1 \pmod{L}$. To decide which tracks are used, compare the two ordered arrays of thresholds $b_j(i)$ and $b_j(k) + \lambda$ at $s - a$. The possible motion of each threshold is confined to $[u_j - Cm, u_j + Cm + \lambda]$. These intervals are disjoint, so at most one threshold can change the index. Its displacement is

$$b_j(k) + \lambda - b_j(i) = \lambda - \epsilon_j C(k-i).$$

It is negative at a forward transition, allowing only a change from $j-1$ to j, and positive at a backward transition, allowing only a change from j to $j-1$. Each possible change is an edge of D. Thus $w^a : Y \to D$ for every a.

When a increases by one, the only columns that can change are $r \equiv a + b_j(i) \pmod{p}$. For a fixed j, distinct such columns are at distance at least $C > \lambda$; for different j, the cyclic gaps are also greater than λ. Consequently there is no Y-edge between two columns that change in this step. Changing those columns separately and in any order therefore preserves every edge, which keeps either its old assignment or its new one. This gives a reconfiguration inside R_D.

Choose a positive a with $a \equiv 1 \pmod{p}$ and $a \equiv 0 \pmod{L}$. The Chinese remainder theorem applies. Now $s - a$ lifts $r - 1$, and its phase equals that of s, so $w^a(i,r) = w^0(i, r-1)$. Concatenating the steps from 0 to a gives exactly the right rotation. This proves property (i). The map $D \to K_M$ induces $R_D \to R_{K_M}$, so the pair is a promise template. The gadget has fixed size, and its decoder groups the m vertex colours belonging to each relation variable. This proves the LOGSPACE guarantees in (iii). □

D Proofs for the reduction criteria of Section 1

D.1 Exact lists and intersecting lists

Boolean conjunction separates ordinary from exact set-valued homomorphisms. Boolean majority separates exact set-valued homomorphisms from d-homomorphisms.

Example D.1 (A proper weakening)**.** Let $\mathscr{C}_2(I)$ consist of the Boolean functions $f_S(x) = \bigwedge_{i\in S} x_i$, where $S \subseteq I$ and $1 \le |S| \le 2$. Since $f_S^\pi = f_{\pi(S)}$, these functions form a minion. The list assignment

$$\xi(f_S) = \{e_i : i \in S\}$$

is an exact 2-set-valued homomorphism to $\mathscr{P}$. Binary conjunction is fixed by the transposition of its two coordinates. Neither binary projection is fixed, so there is no ordinary minion homomorphism to $\mathscr{P}$.

Example D.2 (Majority separates the two list conditions)**.** Let $\mathscr{T}$ be the minion generated by the Boolean majority function $m(x, y, z)$, which returns the value occurring at least twice. Every operation in $\mathscr{T}$ is a projection or majority on three distinct coordinates. Selecting all essential coordinates gives a 3-homomorphism $\mathscr{T} \to \mathscr{P}$: a minor is nonconstant and every essential coordinate of the minor is the image of an essential coordinate of the original operation.

There is no exact set-valued homomorphism $\mathscr{T} \to \mathscr{P}$, for any finite bound. Such an exact set-valued homomorphism must send each projection e_i to $\{e_i\}$, since projections are minors of the unary projection. The invariance of m under all permutations forces its nonempty list to contain all three projections. Under the map $\pi(1) = \pi(2) = 1$, $\pi(3) = 2$, that list has image $\{e_1, e_2\}$. But $m^\pi = e_1$, whose list must be $\{e_1\}$, a contradiction.

D.2 Barto–Kozik homomorphisms

Barto and Kozik observed that ordinary minion homomorphisms compose on either side of a (d, r)-homomorphism and asked about composing two (d, r)-homomorphisms [5, footnote 3 after Definition 3].

Proposition D.3 (Composition)**.** *The union composition of a (d, r)-Barto–Kozik homomorphism ξ and an (e, s)-Barto–Kozik homomorphism η has parameters (de, R), where $R = r(d(s-1)+1)$.*

Proof. Put $\chi = \eta \star \xi$. Suppose a source minor chain of $R+1$ operations has no χ-compatible pair. Fix a position i and $g \in \xi(f_i)$. There are at most $s-1$ later positions j with $g^{\pi_{ij}} \in \xi(f_j)$. Otherwise g, followed by its images at s such positions, forms an actual minor chain in $\mathscr{N}$: all images come from the same g and the coordinate maps compose. Applying η to this minor chain produces a χ-compatible pair, a contradiction.

Call $j > i$ a neighbour of i if $\xi(f_i)^{\pi_{ij}} \cap \xi(f_j) \neq \varnothing$. Each position has at most $k = d(s-1)$ later neighbours. Select the earliest remaining position and delete it together with its later neighbours. Each selection removes at most $k+1$ positions. From $r(k+1)+1$ positions this gives at least $r+1$ selected positions with no compatible pair. The source subchain on these positions contradicts the (d, r) property of ξ. □

Example D.4 (Two layers are strictly more general than one)**.** We use the minion $\mathscr{Q}$ and its two-state selection from the project note [13]. Its elements are the following proper six-colorings.

Take an upward-closed family $W \subseteq 2^I$ satisfying $A \in W \iff I \setminus A \notin W$, and $i \in I$. Define $f_{W,i} : \{0,1,2\}^I \to \{0,1,\dots,5\}$ by

$$f_{W,i}(x) = \begin{cases} a, & \{j : x_j = a\} \in W \quad (a \in \{0,1,2\}), \\ 3 + x_i, & \text{if no color class belongs to } W. \end{cases}$$

Any two members of W intersect, so at most one color class belongs to W. Adjacent inputs cannot receive the same color in the first case; equal colors in the second case would force agreement at coordinate i. Thus $f_{W,i}$ is proper. Its minors obey

$$f^{\pi}_{W,i} = f_{W^{\pi},\pi(i)}, \qquad W^{\pi} = \{A \subseteq J : \pi^{-1}(A) \in W\}.$$

The family W^{π} has the same two properties as W, so $\mathscr{Q}$ is a minion.

The family W is recovered from two-color inputs: the value on an input using only $0, 1$ is zero exactly when its zero set belongs to W. If $W = \{A \subseteq I : j \in A\}$ for some j, then $f_{W,i} = e_j$. Otherwise the function determines i uniquely: any two distinct coordinates can be separated by a partition of I into three sets outside W, and the corresponding input distinguishes their functions. To construct such a partition when one coordinate i belongs to a minimal member A of W, use $\{i\}, A \setminus \{i\}, I \setminus A$. All three sets lie outside W: in particular, a singleton in W would force W to have the excluded form. If neither coordinate belongs to any minimal member of W, move the two coordinates into different parts of a partition just constructed. These moves preserve nonmembership in W, since they do not change which minimal members of W are contained in a part.

Select j for a projection e_j, and i for a nonprojection $f_{W,i}$. Assign these two cases states 0 and 1. States do not increase under minors, and selections transport exactly whenever states agree. In a two-arrow minor chain two adjacent states agree; hence $\mathrm{BK}_{1,2}(\mathscr{Q}, \mathscr{P})$ holds.

Nevertheless there is no D-homomorphism $\mathscr{Q} \to \mathscr{P}$ for any finite D, even if compatibility is tested only on surjective minors. Given proposed lists C, set $m = D + 1$, $q = 4D + 1$, choose an m-ary projection e_j, and take $f = f_{W,i}$ on $[q]$, where $W = \{A \subseteq [q] : |A| \geq 2D + 1\}$ and $i \in [q]$ is arbitrary. Choose $b \notin C(e_j)$. Send $C(f)$ to b, send D unmarked coordinates bijectively to $[m] \setminus \{j\}$, and send all remaining coordinates to j. This map π is onto, and at least $2D + 1$ source coordinates map to j. Therefore $f^{\pi} = e_j$, whereas $\pi(C(f)) \cap C(e_j) = \varnothing$.

This separation is for the function minion $\mathscr{Q}$.

D.3 A separation at the Barto–Kozik condition

Example D.5 (A finite-template separation)**.** Let ST be the Boolean minion generated by

$$g_{2m+1}(x_1, \dots, x_{2m+1}) = \mathbf{1}\left[x_1 + \sum_{j=1}^{m} 2^j (x_{2j} - x_{2j+1}) > 0\right], \qquad m \geq 0.$$

Here $\mathbf{1}[P]$ is one if the statement P holds and zero otherwise. Banakh and Kozik realize ST as the polymorphism minion of a finite Boolean promise template by naming both unary singletons [10, Theorem 3.8]. Their proof gives injective layered choice for ST [10, Observation 3.9 and the proof of Theorem 3.2 in Section 5.2]. Their hardness proof [10, Theorem 3.1 and Appendix B] uses smooth layered Label Cover. Theorem 1.15 gives a PCP-free reduction through the injective variant of the Barto–Kozik condition in Section 1.6. For every d, r, their counterexample to a (d, r)-Barto–Kozik homomorphism $\mathrm{ST} \to \mathscr{P}$ uses only coordinate identifications, so it also excludes a surjective Barto–Kozik homomorphism [10, proof of Observation 3.7].

D.4 Consistency reductions

We recall the level construction in Theorem 1.3 and the two properties of consistency reductions used in the comparison, namely closure under composition and their separation from the chain condition.

Definition D.6 (Levels of a minion [8, Definition 3.2])**.** Fix a minion $\mathscr{R}$ and $k \geq 1$. Choose a finite nonempty family $\mathcal{E} = (E_x)_{x \in X}$ of finite nonempty sets. For $j \geq 0$ and $\bar{x} = (x_1, \dots, x_j) \in X^j$, put $E_{\bar{x}} = \prod_{i=1}^{j} E_{x_i}$; the empty product is a singleton. Choose elements $r_{\bar{x}} \in \mathscr{R}(E_{\bar{x}})$ for $\bar{x} \in X^k$. For $\sigma : [k] \to [k]$, let

$$p_{\bar{x},\sigma} : E_{\bar{x}} \to E_{\bar{x} \circ \sigma}, \qquad (a_1, \dots, a_k) \mapsto (a_{\sigma(1)}, \dots, a_{\sigma(k)}).$$

Require compatibility under reindexing,

$$r_{\bar{x} \circ \sigma} = r_{\bar{x}}^{p_{\bar{x},\sigma}} \qquad (\bar{x} \in X^k,\ \sigma : [k] \to [k]).$$

An I-ary element of the *k-th level* $\mathrm{lvl}_k(\mathscr{R})$ is a triple $(\mathcal{E}, r, w)$, where $w = (w_{\bar{y}})_{\bar{y} \in X^{k-1}}$ is obtained from some $z \in X$ and $\beta : E_z \to I$ by

$$w_{\bar{y}} = r_{(z,\bar{y})}^{\beta \times \mathrm{id}_{E_{\bar{y}}}} \in \mathscr{R}(I \times E_{\bar{y}}) \qquad (\bar{y} \in X^{k-1}).$$

The witnesses z, β are not part of the triple. For $\pi : I \to J$, define

$$(\mathcal{E}, r, w)^{\pi} = (\mathcal{E}, r, w'), \qquad w'_{\bar{y}} = w_{\bar{y}}^{\pi \times \mathrm{id}_{E_{\bar{y}}}}.$$

The fixed data $(\mathcal{E}, r)$ record a compatible table on products of k coordinate sets; taking minors changes only the distinguished coordinate set I. The witness for the π-minor is $z, \pi \circ \beta$, and the minor identities follow from those of $\mathscr{R}$.

Proposition D.7. *For fixed templates and k, the map T_k is computable in polynomial time and satisfies $\mathbf{X} \to \mathbf{A} \Longrightarrow T_k(\mathbf{X}) \to \mathbf{C}$. If some surviving family F_U is empty, then $T_k(\mathbf{X}) \not\to \mathbf{D}$.*

Proof. The restrictions of a global $\mathbf{A}$-solution survive pruning: they supply one another's required restrictions and extensions, so none can be the first deleted assignment. Evaluation at these restrictions gives compatible maps $\mathbf{C}^{F_U} \to \mathbf{C}$, and hence a map from the quotient. If F_U is empty, the unique vertex of $\mathbf{C}^{F_U}$ belongs diagonally to every relation. Its image in $\mathbf{D}$ would have the same property, making every instance map constantly to $\mathbf{D}$, contrary to the NO-instance assumption. There are polynomially many assignments in all families, and every pruning step deletes one. The tests, powers, and quotient can all be computed in polynomial time. □

Logical characterization. A *Datalog program* is a finite collection of positive relational rules, without function symbols, evaluated by repeatedly adding all consequences until a fixed point is reached. A *Datalog interpretation* uses designated relations defined by Datalog programs as domains of tuples and as relations of an output relational structure. It may have finitely many sorts; a *sort* is a separately tagged domain, and each argument of a relation has a prescribed sort. A union step merges selected sorts by disjoint union and merges the corresponding relations by union. A *Datalog-with-unions reduction* is a Datalog interpretation followed by such a union step.

Theorem D.8 (Dalmau–Opršal)**.** *For finite promise templates, a consistency reduction exists if and only if a Datalog-with-unions reduction exists. The existence of a consistency reduction is reflexive and transitive.*

The equivalence is [7, Theorem 3.12 and Corollary 3.13]. Transitivity follows from their composition theorem [7, Theorem 3.7], and reflexivity follows from the identity interpretation.

Lemma D.9 (Small orbit)**.** *Suppose there is an injective chain homomorphism $\mathscr{M} \to \mathscr{N}$ of height h and list bound d when tested on surjective minor chains. If a finite group G of coordinate permutations fixes $f \in \mathscr{M}(I)$, then G has an orbit in $\mathscr{N}(I)$ of size bounded in terms of h, d.*

Proof. Choose n large enough for Lemma 1.18 with $t = 2$. Form a chain of n copies of f, choosing each arrow independently and uniformly in G. Every arrow is a bijection and hence is injective on every set of accumulated marks. Write b for the number of pairs of h-sets $F < E$ among these positions. In every such chain some pair has matching entries, by the lemma applied to the lists in $\mathscr{N}(I)$ with the minor maps given by its arrows.

Fix a pair $F < E$ and fix every arrow except the one immediately after $\max F$. Both lists are then fixed. The composite from $\max F$ to $\max E$ is still uniform in G, so the image of each of the at most d entries on F is uniform on its orbit. If every orbit had size greater than bd^2, the probability of hitting one of the at most d entries on E would be less than $1/b$, even after summing over the entries on F. Summing over the b pairs would give probability less than one, a contradiction. □

Example D.10 (Consistency does not imply the chain condition)**.** Let

$$\mathbf{H} = (\{0,1\}; R_\wedge, \{0\}, \{1\}), \qquad R_\wedge = \{(x,y,z) : x \wedge y = z\}, \qquad \mathscr{H} = \mathrm{Pol}(\mathbf{H}).$$

There is a 3-consistency reduction from $\mathrm{CSP}(K_2)$ to $\mathrm{CSP}(\mathbf{H})$. However, there is no chain homomorphism $\mathscr{H} \to \mathrm{Pol}(K_2)$. This remains true even if only surjective injective chains are tested, for any choice of bounded marks. This separation extends the symmetry argument of [7, Example C.9].

Proof. The polymorphisms of $\mathbf{H}$ are exactly the nonempty conjunctions. Indeed, preservation of meet makes the inverse image of one an upward-closed, meet-closed family of Boolean tuples. Preservation of the two unary relations makes it nonempty and excludes the zero tuple. On a finite Boolean cube it is therefore a principal filter generated by a nonzero tuple, which gives precisely a nonempty conjunction.

Suppose the 3-consistency procedure for two-coloring a graph has no empty family. Every surviving assignment to the endpoints of a walk, including one whose endpoints coincide, has the parity prescribed by its length. This follows by induction: extend it to the endpoints and the penultimate vertex, then restrict to the shorter walk and to the last edge. An odd cycle is therefore impossible, while a loop is detected by a local constraint. Thus a non-bipartite graph produces an empty family, whose zero power has both unary relations on its unique vertex and cannot map to $\mathbf{H}$. If the input graph $\mathbf{X}$ is bipartite, Proposition D.7 gives $T_3(\mathbf{X}) \to \mathbf{H}$. This proves the reduction.

For the algebraic obstruction, take $I = \mathbb{F}_2^t$ and let G be the group of translations of I. The conjunction of all coordinates is fixed by G. The operations of $\mathrm{Pol}(K_2)(I)$ are exactly the *self-dual Boolean functions*, namely those satisfying $g(\neg x) = 1 - g(x)$. Every nonidentity translation τ is a fixed-point-free involution. Choose $x \in \{0,1\}^I$ with opposite bits at the two coordinates of each τ-pair. For every self-dual g,

$$g^\tau(x) = g(\neg x) = 1 - g(x) \neq g(x).$$

Consequently every G-orbit in $\mathrm{Pol}(K_2)(I)$ has size 2^t. An injective chain homomorphism on these surjective chains would force an orbit of bounded size by Lemma D.9. Taking t large enough gives a contradiction. □

D.5 Logspace construction and decoding

Proposition D.11 (LOGSPACE implementation)**.** *For fixed finite templates and k, the unpruned local-power construction $\widetilde{T}_k$ of Definition 1.14 is computable in LOGSPACE. For the coloring application, replacing a quotient with a loop by K_{n+1} preserves this bound.*

Proof. Let N be the input length. Since the templates and k are fixed, each U can be stored in $O(\log N)$ space, and $|F_U| \le |A|^k$ is a constant. Enumerate the assignments on U and check each against the input constraints by rescanning the input. This enumerates F_U, detects empty families, and computes each restriction map in logspace. Each power $\mathbf{C}^{F_U}$ has constant size, so the disjoint union, its relations, and the generating identifications can also be enumerated in logspace.

Regard the generating identifications as edges of an undirected graph on this disjoint union. The graph has polynomially many vertices and logspace-computable adjacency. Two vertices are identified precisely when they are connected. Reingold's theorem [9, Theorem 4.1] decides this in logspace. Represent each component by its least vertex in the enumeration; scanning candidate vertices and testing connectivity computes that representative in logspace. We can therefore output the quotient vertices and relation tuples in logspace, suppressing duplicate tuples by enumeration if required. The output has polynomial size. This proves the asserted bound on the instance map.

For the coloring application, a quotient has a loop precisely when some edge of the disjoint union has both endpoints in one equality component. Enumerate the edges and use the same connectivity test to detect this in logspace, before producing the output. The fixed NO instance K_{n+1} can then be output whenever required. □

Lemma D.12 (Finite matrix alternative)**.** *Let M be a finite zero–one matrix with nonempty sets of rows and columns, and $0 < \delta < 1$. One can find either a distribution on the rows giving the ones in every column mass less than δ, or a distribution on the columns giving the ones in every row mass at least $\delta/4$. For fixed rational δ this takes polynomial time.*

Proof. Give each of the N rows weight one. In a fixed order, take the first column whose ones have at least a δ fraction of the total weight. Record it and double the weights of the rows where it has a zero. If no column qualifies, normalize the row weights and stop. Otherwise run for $H = \lceil 4/\delta \rceil \lceil \log_2(2N) \rceil$ rounds. The total final weight is at most $N(2-\delta)^H$. A row having ones in a of the recorded columns has weight 2^{H-a}. Therefore

$$a \ln 2 \;\ge\; -H \ln(1-\delta/2) - \ln N \;\ge\; H\delta/2 - \ln N \;\ge\; H\delta/4.$$

The uniform distribution on the recorded columns, with repetitions, has the required mass in every row. The weights have $O(H)$ bits. □

The ordinary two-level theorem.

Proof of Theorem 1.10. We use the argument of Barto and Kozik [6, Section 3.1]. Allow separate bounds u, ℓ on the upper lists J and lower lists K. Fix m and induct lexicographically on (ℓ, u), always taking $n = p$. For $u = \ell = 1$ take $p = 2m$, $q = m$. Any two lower scopes lie in a common upper scope, so their unique entries agree on the intersection. They glue to an m-solution of K.

Otherwise take arities (a, b) for upper and lower bounds $(\ell, 1)$, and (s, t) for $(u-1, \ell)$ if $u > 1$. Set $s = t = 1$ if $u = 1$. The first call has rank $(1, \ell)$, which is smaller even when $\ell = 1$ because then $u > 1$. The second call decreases u. Put

$$q = a + \binom{a}{b} t, \qquad M = \binom{s}{t} q, \qquad p = 1 + \max\{s, q + \ell M\}.$$

Suppose every b-set W has a tuple z_W and a t-set Y_W such that $z_W \in K(Z)|_W$ whenever Z is a q-set containing $W \cup Y_W$. For each a-set X, choose a q-set Z_X containing X and all Y_W for $W \in \mathcal{P}_b(X)$. The upper lists $K(Z_X)|_X$ and lower singletons $\{z_W\}$ have ordinary overlap. All their entries are restrictions of entries of K, so induction gives an m-solution of K.

Otherwise fix a b-set W where every tuple can be avoided above every prescribed t-set. Choose a q-set $Z_0 \supseteq W$. Some $z \in K(Z_0)|_W$ has the following property. For every set T of at most M variables, some p-set $U \supseteq W \cup T$ has $z \in J(U)|_W$. Indeed, if every candidate z had a blocking set T_z, a p-set containing Z_0 and all the T_z would contradict ordinary overlap at Z_0.

For each t-set Y choose a q-set $Z_Y \supseteq W \cup Y$ with $z \notin K(Z_Y)|_W$. For each s-set X, the union of the Z_Y over $Y \in \mathcal{P}_t(X)$ contains X and has at most M variables. Choose a p-set U_X containing that union with $z \in J(U_X)|_W$. Delete from $J(U_X)$ all entries restricting to z on W, then restrict the survivors to X. For every $Y \in \mathcal{P}_t(X)$, ordinary overlap gives an upper entry matching an entry of $K(Z_Y)$ on Z_Y. It survives because $K(Z_Y)|_W$ avoids z. Thus the new upper lists are nonempty, have size at most $u-1$, and have ordinary overlap with $K(Z_Y)|_Y$. For $u = 1$ a survivor is impossible. Otherwise induction gives an m-solution that lifts to the original system supplying its entries.

On finite explicit input, break ties by input order. A candidate forced tuple on W can be enumerated among restrictions of lower entries on scopes containing W. Every other test scans fixed-size sets and list entries. Store variable indices and pointers to value names, comparing names by rescanning the input. Regenerate the derived lists on demand. The induction depth and all arities are fixed, so outputting the chosen system and the assignment one variable at a time takes logarithmic work space [6, Appendix B]. □

The empty-mark case of Theorem 1.16. Under ordinary chain consistency the two children at a split can be joined by deterministic choices. We give the size bounds and choices here; they preserve the ordinary theorem for arbitrary value sets and the logspace bound for finite explicit input. For $h = 1$ use Theorem 1.10. For $h \geq 2$ take the full binary tree of the height excluded by the matching argument in Section 1.7.2. Prescribe sizes recursively at a node with minimum size k and a required set of at most ℓ variables. Start with $k = a$ and $\ell = 0$. At a leaf use h consecutive sizes with smallest size $\max(k, \ell)$. If each child has m positions, use Theorem 1.10 with list bound $d\binom{m}{h}$ to obtain sizes $p > q$ and a threshold n. Prescribe the lower child with minimum size k and budget $\ell + q$. If its largest size is k_D, prescribe the upper child with minimum size $k_D + 1$ and budget $p + \binom{p}{q}k_D$. Take the threshold for $|V|$ to be at least the largest prescribed size and every n so obtained.

Choose a node of least height for which some set Z within its budget cannot be retained throughout a block. Such a node exists since no root block exists, and every leaf can retain its required set. For every q-set X, its lower child admits a block D_X retaining $X \cup Z$. For every p-set T, its upper child admits a block U_T retaining $T \cup \bigcup_{X \in \mathcal{P}_q(T)} \max D_X$. At $X \subseteq T$ the concatenation $U_T D_X$ retains Z and has the prescribed sizes. Some upper face therefore matches a lower face. Restrict all upper face entries to T to form $J(T)$, and all lower face entries to X to form $K(X)$. These lists meet the hypotheses of Theorem 1.10. All these choices also work on an infinite variable set.

For finite explicit input the tree and all size bounds depend only on a, h, d. Enumerate nodes by increasing height, from left to right at each height. For each node enumerate the permitted sets Z, and for each Z enumerate the decreasing chains of the prescribed sizes whose smallest member contains Z. Use lexicographic order on sets and chains, using the input indices of the variables. To test whether a candidate is a block, inspect each of its internal splits. For each pair of h-faces on opposite sides, scan their lists for entries agreeing on the smaller domain. A chain is a block exactly

when no such pair matches at any split. This finds the first node and Z for which no candidate is a block.

For this node and Z, use the first qualifying lower block D_X for each X. On a query for U_T, recompute the D_X for $X \in \mathcal{P}_q(T)$ to form the required padding set, then take the first qualifying upper block. Generate $J(T)$ and $K(X)$ by restricting their face-list entries, suppressing duplicates by rescanning earlier entries. Each chain and padding set contains a bounded number of variables. Storing their input indices and pointers to value names, and comparing names by rescanning, uses $O(\log N)$ space. The two families of lists have polynomial total size. Simulate the decoder of Theorem 1.10 on these lists, regenerating its input symbols on demand, to output a solution in $O(\log N)$ space.

Proof of the LOGSPACE decoding claim in Theorem 1.15. For the ordinary chain homomorphism, the input consists of $\mathbf{X}$ and an explicitly represented homomorphism $q : \widetilde{T}_k(\mathbf{X}) \to \mathbf{D}$. For each bounded set U, enumerate F_U in a fixed order. Each table f_U is obtained by evaluating q on the corresponding quotient representatives, which are computable in LOGSPACE by Proposition D.11. The finite templates and k bound all these table sizes by a constant. The values of Ξ on the resulting bounded minor chains can therefore be fixed in the algorithm. Evaluating b_g as in the main proof gives LOGSPACE query access to each $L(C)$. Simulate the reconstruction algorithm just proved on the polynomial-size explicit encoding of these lists, regenerating requested symbols in LOGSPACE. This gives the required search reduction with a LOGSPACE search decoder. If $|X| < N_0 \le k$, enumerate the bounded set F_X instead and output the first member composed with a fixed promise homomorphism $\mathbf{A} \to \mathbf{B}$.

E Surjective variants

In this section we prove the implications in the right panel of the diagram in Section 1. Consistency pruning makes the restriction maps between surviving local solutions surjective.

E.1 Surjective-minor homomorphisms and arc-consistency construction

Definition E.1 (Surjective-minor homomorphism)**.** A *surjective-minor homomorphism* $\xi : \mathscr{M} \to \mathscr{N}$ is an arity-preserving map satisfying

$$\xi(f^\pi) = \xi(f)^\pi \qquad \text{whenever } \pi : I \twoheadrightarrow J.$$

The adjective refers to the coordinate map π; no surjectivity of ξ is required.

To express surjective preservation as ordinary minor preservation, record a set of coordinates on which each operation is represented. An arbitrary coordinate map is surjective onto the image of this recorded set.

Definition E.2 (The minion $\omega(\mathscr{M})$)**.** For a minion $\mathscr{M}$, define

$$\omega(\mathscr{M})(I) = \{(S, f) : \varnothing \neq S \subseteq I,\ f \in \mathscr{M}(S)\}, \qquad (S, f)^\pi = (\pi(S), f^{\pi|_S}). \tag{16}$$

In this formula $\pi|_S : S \to \pi(S)$ is surjective. The set S is the *support* of the element (S, f). The support is specified as part of the element, even when f has dummy coordinates.

The minor identities for $\omega(\mathscr{M})$ follow immediately from those for $\mathscr{M}$. This is the construction of [7, Definition 5.4].

The corresponding arc-consistency construction retains possible values at each variable and possible tuples at each constraint, then glues the associated powers. Given an input $\mathbf{X}$, start with $F_v = A$ for every variable v. For each constraint $c = (R, v_1, \dots, v_m)$ put

$$E_c = R^{\mathbf{A}} \cap (F_{v_1} \times \cdots \times F_{v_m}).$$

Repeatedly replace each F_{v_i} by its intersection with the i-th coordinate projection of E_c, and recompute the sets E_c, until no set changes. In the resulting relational structure, take a copy of $\mathbf{C}^{F_v}$ for each v and a copy of $\mathbf{C}^{E_c}$ for each constraint c. For every coordinate projection $\rho_{c,i} : E_c \to F_{v_i}$, identify $x \in C^{F_{v_i}}$ with $x \circ \rho_{c,i} \in C^{E_c}$. Relations in the quotient are the images of the relations in these powers. This is the arc-consistency construction of [7, Definition 5.1].

Theorem E.3 (Dalmau–Opršal)**.** *The arc-consistency construction is a reduction in* (1) *if and only if there is a surjective-minor homomorphism* $\mathscr{M} \to \mathscr{N}$*.*

Dalmau–Opršal [7, Theorem 5.5] characterize the construction by $\mathrm{Hom}(\omega(\mathscr{M}), \mathscr{N})$. Proposition E.7 below translates their condition.

Example E.4 (Preserving surjective minors)**.** Let $\mathscr{T}$ be the majority minion of Example D.2. The map $f \mapsto (I, f)$ from $\mathscr{T}(I)$ to $\omega(\mathscr{T})(I)$ preserves surjective minors. Suppose there were an ordinary minion homomorphism $\mathscr{T} \to \omega(\mathscr{T})$. Every operation is a minor of ternary majority, so every image would have support of size at most three. Taking these supports as coordinate lists would give an exact 3-set-valued homomorphism to $\mathscr{P}$, since a minor sends a support S to $\pi(S)$. This contradicts Example D.2.

E.2 Surjective Barto–Kozik homomorphisms

Definition E.5. A *surjective* (d, r)*-Barto–Kozik homomorphism* has the lists of Definition 1.7 and requires (4) for every surjective minor chain of r arrows. Write $\mathrm{BK}^{\mathrm{surj}}_{d,r}(\mathscr{M}, \mathscr{N})$ for existence of such a map, and $\mathrm{BK}^{\mathrm{surj}}(\mathscr{M}, \mathscr{N})$ if some finite d, r suffice.

Every Barto–Kozik homomorphism restricts to surjective chains. With $d = r = 1$ the condition is exactly preservation of surjective minors, since intersecting singletons are equal. Example D.4 gives a surjective $(1, 2)$-Barto–Kozik homomorphism to $\mathscr{P}$ and rules out every one-arrow bound, already using surjections. For finite templates $\mathrm{BK}^{\mathrm{surj}}(\mathscr{M}, \mathscr{N})$ implies a consistency reduction by [7, Corollary C.6] and Proposition E.7.

E.3 Surjective chain homomorphisms

Definition E.6 (Surjective chain homomorphism)**.** A *surjective* (h, d)*-chain homomorphism* assigns lists as in Definition 1.13 to surjective minor chains with h operations. The *surjective chain condition* requires the transported face lists to intersect on every surjective minor chain with $h + 1$ operations. Write $\mathrm{Ch}^{\mathrm{surj}}_{h,d}(\mathscr{M}, \mathscr{N})$ for its existence, and $\mathrm{Ch}^{\mathrm{surj}}(\mathscr{M}, \mathscr{N})$ if some finite h, d suffice.

The injective variant uses one choice of bounded coordinate marks and tests the intersection on every chain that is both surjective and injective. For a single-output rule, exact lists, intersecting lists, or a Barto–Kozik homomorphism the same two restrictions apply to the respective defining condition.

Every chain homomorphism restricts to surjective minor chains. One support construction describes all three conditions.

Proposition E.7 (Support translation)**.** *For each of the five rows of the diagram in Section 1, with or without the injectivity requirement, the surjective version from $\mathscr{M}$ to $\mathscr{N}$ exists if and only if the unrestricted version from $\omega(\mathscr{M})$ to $\mathscr{N}$ exists, with the same list, height and mark bounds, where applicable.*

Proof. For a surjection $\pi : I \to J$ we have $(I, f)^\pi = (J, f^\pi)$. Thus a rule on $\omega(\mathscr{M})$ gives the surjective rule on $\mathscr{M}$ by using full supports.

Conversely, replace each input $(S, f) \in \omega(\mathscr{M})(I)$ by $f \in \mathscr{M}(S)$. On a chain restrict each arrow π to $\pi|_S : S \to \pi(S)$, which is surjective. Apply the given rule to these inputs. If the last support is $S \subseteq I$, take the minor of each output along the inclusion $\iota_S : S \hookrightarrow I$. Since $\pi \circ \iota_S = \iota_{\pi(S)} \circ (\pi|_S)$, extending by inclusion commutes with transport. It preserves the defining equality for single outputs, the witness for a compatible pair of operation lists, and the common entry for all face lists of a chain. It also carries equality of exact lists to equality of their images. For the injective variants, give (S, f) the marks of f viewed as elements of I via $S \hookrightarrow I$. The accumulated marks at each position are exactly the images of those on the restricted chain. A chain satisfying the test in I therefore satisfies it on the supports. In the converse direction take the marks at full supports. Inclusions preserve the mark bounds, and list sizes cannot increase. □

Apply Proposition F.1 to $\omega(\mathscr{M})$ and use the support translation. A surjective (d, r)-Barto–Kozik homomorphism gives a surjective (r, rd)-chain homomorphism. Height one gives the one-arrow condition, and singleton lists at any height give precisely the surjective-minor homomorphisms.

Theorem E.8 (Surjective chain-homomorphism reduction)**.** *For the finite templates in* (1)*, if there is a homomorphism on injective chains from $\mathscr{M}$ to $\mathscr{N}$ with the test restricted to surjective minor chains, then a consistency reduction exists with a polynomial-time search decoder.*

Proof. Use the pruned construction T_k from Definition 1.2, with the size chosen in Theorem 1.15. Proposition D.7 gives completeness, handles empty families, and bounds the construction time. The surviving families satisfy (2), so their restriction maps are surjective. Apply the soundness argument of Theorem 1.15 to these families: all induced minor chains are now surjective, so the hypothesis gives the required face intersections on the injective chains. The marks are still local assignments on the current scope. Theorem 1.16 therefore reconstructs a solution in polynomial time, including the time to obtain the pruned families. □

The converse fails for finite templates by Example D.10.

F Comparison and composition of chain homomorphisms

Proposition F.1 (From Barto–Kozik to chain homomorphisms)**.** *For positive parameters,*

$$\begin{aligned}\mathrm{BK}_{d,r}(\mathscr{M},\mathscr{N}) &\Longrightarrow \mathrm{Ch}_{r,rd}(\mathscr{M},\mathscr{N}),\\ \mathrm{Ch}_{1,d}(\mathscr{M},\mathscr{N}) &\Longleftrightarrow \mathrm{BK}_{d,1}(\mathscr{M},\mathscr{N}),\\ \mathrm{Ch}_{h,1}(\mathscr{M},\mathscr{N}) &\Longleftrightarrow \mathrm{Hom}(\mathscr{M},\mathscr{N}).\end{aligned}$$

The same three implications hold for the injective variants, with the same bound on coordinate marks. Precomposition and postcomposition by ordinary minion homomorphisms preserve the parameters of a chain homomorphism.

Proof. Given BK vertex lists, put on a minor chain with r operations the union of its r lists, transported to its end. On a minor chain with $(r+1)$ operations, the compatible pair has a common terminal image. Every r-face retains at least one member of that pair, so its union contains this common image. The bound is rd. Height one is exactly the two-list intersection condition on an arrow.

For singleton chain lists, write $v(F)$ for the unique value on a minor chain F, and define $\xi(f)$ using the identity minor chain with h operations at f. Given $f \xrightarrow{\pi} g$, let F_j contain $h-j$ copies of f followed by j copies of g, for $0 \le j \le h$. Use identity maps within each run and π between the runs. For $1 \le j \le h$, compare the faces of the minor chain with $(h+1)$ operations with $h+1-j$ copies of f followed by j copies of g. Deleting an f gives F_j; deleting a g gives F_{j-1}. The equality of the transported singleton entries gives $v(F_0)^\pi = v(F_1)$ when $j=1$, and $v(F_{j-1}) = v(F_j)$ when $j>1$, since both endpoints are then g. Thus $\xi(f)^\pi = \xi(g)$. Conversely, a minion homomorphism ξ gives the singleton $\{\xi(g)\}$ on each chain ending at g. Precomposition preserves the coordinate maps of a chain, and postcomposition sends a common entry of the face lists to a common entry of their images. For the injective variants use the given marks throughout. A subchain of an injective chain is injective. In the singleton argument, inserting identity arrows and repeating either endpoint of an injective arrow preserves injectivity on the accumulated marks. □

Theorem F.2 (Composition of chain homomorphisms)**.** *If $\Xi : \mathscr{M} \to \mathscr{N}$ and $\eta : \mathscr{N} \to \mathscr{Q}$ are chain homomorphisms with parameters (h,d) and (k,e), then there is a chain homomorphism $\mathscr{M} \to \mathscr{Q}$ whose parameters depend only on h,d,k,e.*

Proof. Choose $H \ge h+k-1$ so that Lemma 1.18 applies to $H+1$ positions with list bound d, face size h, and $t = k+1$. To form the list $\Theta(F)$ on an H-position chain F, choose an h-subchain B and $k-1$ later positions of F. For each $g \in \Xi(B)$, follow g by its minors at those positions. Apply η to this k-position chain and transport its output list to the end of F. Take the union over all choices. These lists are nonempty, and there are at most $de\binom{H}{h+k-1}$ entries in the union.

Now let C have $H+1$ operations on coordinate sets I_i. Apply Lemma 1.18 to $S_i = \mathscr{N}(I_i)$ with the minor maps $u \mapsto u^{\pi_{ij}}$ and lists Ξ. This gives disjoint h-subchains $B_0 < \cdots < B_k$ and entries $g_j \in \Xi(B_j)$ forming a minor chain. Its k-faces have η-lists with a common image at the end of B_k. Transport that image to the end of C and call it z.

Each H-face F of C misses a position of at most one B_j. Discard that B_j, or any one if all lie in F. The first remaining subchain and the endpoints of the other $k-1$ are among the choices defining $\Theta(F)$. They give the corresponding k-face of the chain of the g_j. Its list has an entry with image z at the end of C. Thus z belongs to every transported face list. □

The composition proof also applies when the first chain homomorphism is injective and the second is an ordinary chain homomorphism. Keep the marks of the first map. On an injective

chain every subchain used to obtain the g_j is injective, and the second map can be evaluated on the resulting chain. Precomposition of an injective chain homomorphism by a minion homomorphism α uses the marks $S(\alpha(f))$ on f.

G Proofs for unate Boolean polymorphisms

In this section we prove the direction certificate bound, the rotation obstruction for increasing operations, and the passage from that obstruction to an injective chain homomorphism.

G.1 Cancellation and block symmetry

In this section we prove Proposition 2.7. For a given unate operation g, label each increasing coordinate positive and each decreasing coordinate negative, assigning either sign to a dummy coordinate. Write h for the increasing Boolean function obtained by complementing the negative inputs of g. Only g is used to take minors. For a coordinate x, a set B of other coordinates is a *substitute* for x if

$$h(Z \cup \{x\}) \leq h(Z \cup B) \qquad (Z \cap (B \cup \{x\}) = \varnothing).$$

Write $x \preceq B$ for this inequality. A collection of many disjoint substitutes lets us change all coordinates read by a predictor while controlling its output.

Lemma G.1. *For each $k \geq 1$ there is c such that every f in a minion of unate Boolean operations has a direction certificate of size at most c, or has a minor g in which every coordinate has k pairwise disjoint substitutes of size at most two.*

Proof. For any set S of coordinates of f, identifying all inputs in S gives a unate minor. Thus the differences $f(Z \cup S) - f(Z)$ for $Z \cap S = \varnothing$ are either all nonnegative or all nonpositive. Let $\sigma(S) \in \{-1, 0, 1\}$ be the sign of their average over a uniform Z outside S. In particular $\sigma(S) = 0$ means every difference is zero. Adding an increasing coordinate to S can only raise the average, since this fixes that input to one in the first term and to zero in the second. Removing a decreasing coordinate has the same effect. Consequently if P, N are the positive and negative coordinate sets, any S with $\sigma(S) \geq 0$ gives the increasing certificate $(S \cap P, N \setminus S)$. Dually $\sigma(S) \leq 0$ gives the decreasing certificate $(S \cap N, P \setminus S)$. Disjoint unions of increasing sets are increasing, since their inputs can be changed in successive steps.

Put $s = k+1$ and suppose that both directions lack certificates of size at most $100s$. In particular both P and N have more than $100s$ coordinates, by applying the preceding observation to the empty set. Choose $T_- \subseteq N$ and $T_+ \subseteq P$ maximizing their combined size subject to the existence of s distinct $u \in P \setminus T_+$ with $\sigma(T_- \cup \{u\}) \geq 0$ and s distinct $v \in N \setminus T_-$ with $\sigma(T_+ \cup \{v\}) \leq 0$. Empty sets are feasible. Retain one set U of the s positive witnesses and one set V of the s negative witnesses. At least $100s - 1$ negatives remain outside T_-, since $T_- \cup \{u\}$ gives an increasing certificate. The dual bound holds outside T_+.

Choose $8s$ negative test coordinates outside $T_- \cup V$, and $8s$ positive test coordinates outside $T_+ \cup U$. For a negative test j, fewer than s coordinates $i \in P \setminus T_+$ satisfy $\sigma(T_- \cup \{i, j\}) > 0$. Otherwise j could be added to T_- while retaining V and s positive witnesses. Call i exceptional if it satisfies this strict inequality for at least $4s$ negative tests. Counting the pairs (i, j) gives fewer than $2s$ exceptional positives. Define the exceptional negatives by the reverse inequalities. There are fewer than $2s$ of these as well.

For every exceptional positive i, take a block B_i consisting of i and negative coordinates from outside T_-, V, the negative tests and the exceptional negatives. Choose these blocks disjoint, maximizing their total size subject to $\sigma(T_- \cup B_i) \geq 0$ for each i. The singletons are feasible, since removing j from a positive test comparison can only increase its sign. If any B_i were chosen, at least s eligible negatives j remain unused. Indeed, each B_i is increasing after removing T_- from $T_- \cup B_i$. Thus the set $T_- \cup \bigcup_i B_i$ is the union of $T_- \cup B_i$ for one i with the other disjoint increasing blocks. It contains fewer than $2s$ positives. If fewer than s eligible negatives remained, it would

miss at most $s + 8s + 2s + s$ negatives, and give a certificate smaller than $100s$. By maximality every unused eligible j satisfies

$$\sigma(T_- \cup B_i \cup \{j\}) \le 0 \qquad \text{for every exceptional positive } i.$$

Perform the dual construction for exceptional negatives. The two constructions use disjoint coordinates, since exceptional coordinates of either sign can enter a block only as its distinguished coordinate.

Identify T_-, T_+ and each of the enlarged blocks separately, leaving the remaining inputs single, to obtain g. Empty reference sets are represented by distinct dummy inputs. Denote the reference inputs by r_-, r_+. Give each enlarged input the sign of its distinguished coordinate. This is an admissible orientation because the block itself has that weak direction. Weak direction inequalities for unions of blocks survive every one of these identifications.

Consider a positive input $x \neq r_+$. If x is an enlarged block, the preceding maximality gives at least s distinct untouched negatives j such that $\{x, r_-, j\}$ is decreasing in g. If x is a singleton, at least $4s$ negative tests gave this comparison, and fewer than $2s$ were enlarged as exceptional negatives. Each of the s witnesses in U still gives a positive input u for which $\{u, r_-\}$ is increasing, even if the witness was enlarged. In the oriented cube these inequalities are

$$x \preceq \{r_-, j\}, \qquad r_- \preceq \{u\}.$$

They imply $x \preceq \{u, j\}$ whenever $u \neq x$. If the outside context contains r_-, first replace r_- by u, then replace x by $\{r_-, j\}$. Otherwise perform the replacements in reverse order. Pairing distinct j with distinct $u \neq x$ gives k disjoint substitutes. The reference r_- has its s singleton substitutes. Reversing the signs treats the negative inputs and r_+. □

Proof of Proposition 2.7. Fix t. We show that for k sufficiently large in terms of t, the second conclusion of Lemma G.1 gives a minor that is block-symmetric on blocks of sizes $t, t+1$.

Given $\alpha \in [1/4, 3/4]$, send each positive coordinate of g independently to the first block with probability α, and otherwise to the second. For a negative coordinate use probability $1 - \alpha$ for the first block. In the chosen block pick a uniform input position. This defines an ordinary random minor of g. At an input with densities of ones a, b on the two blocks, the corresponding independent inputs to the oriented function h have probabilities

$$b + \alpha(a - b) \quad \text{and} \quad 1 - a + \alpha(a - b)$$

of being one on the positive and negative coordinates, respectively. Both probabilities have slope $a - b$. Thus the mean $F_{a,b}(\alpha)$ is monotone in α for each density pair. For a nonconstant input, coprimality of $t, t+1$ gives $|a - b| \ge 1/(t(t+1))$.

Choose $\varepsilon > 0$ so that $3\varepsilon 2^{2t+1} < 1$. Divide $[1/4, 3/4]$ into more than $(t+1)(t+2)/\varepsilon$ equal intervals. The sum of the oscillations over all intervals is at most one for each density pair. Thus some interval has oscillation at most ε for every pair. Use its midpoint for the random minor. For a fixed nonconstant input, couple the endpoint and midpoint inputs to h by thresholding independent uniform ranks. In increasing order of the biases write the resulting sets as $L \subseteq X \subseteq H$. Each coordinate belongs to $X \setminus L$ with probability at least $\delta > 0$, and to $H \setminus X$ with probability at least δ, where δ depends only on t and the interval length.

Use ranks uniform on $[M]$, taking a common denominator for all biases at endpoints and midpoints of the original partition, with $1/M \le \delta$. At a witness for a nearby color, each midpoint rank moves by at most one step, so the resulting set lies between L and H. By Lemma A.14 there is a predictor reading at most $b = b(M, \varepsilon)$ ranks, on a set J, whose answer lies between $h(L)$ and

$h(H)$ except with probability ε. Independently refresh the ranks on J to obtain X'. The predictor and $h(X')$ are independent.

To bracket $h(X')$ by the same endpoints, choose l disjoint sets outside J, each containing a separate substitute for every $j \in J$. Use disjoint substitutes across all these choices. Greedy selection works when $k > b + 2bl$, since a previously occupied coordinate meets at most one candidate in the disjoint family for the next j. Each chosen set has size at most $2b$. The probability that one specified gap contains none of the l sets is at most $(1 - \delta^{2b})^l$, by independence across the sets. Choose l so twice this probability is at most ε. All these bounds are independent of g and its arity. For J empty the bracketing is immediate.

If each gap contains a chosen set, then $h(L) \le h(X') \le h(H)$ for every filling of J. Indeed every occupied member of J in L can be replaced by its fresh substitute in $X \setminus L$ to yield a subset of $X \setminus J$. The same replacement in X' using the upper gap gives a subset of $H \setminus J$. Finally $\Pr[h(L) \neq h(H)] \le \varepsilon$ by the choice of interval. The independent predictor and $h(X')$ therefore agree with probability at least $1 - 3\varepsilon$. Since $h(X')$ has mean $F_{a,b}$ at the midpoint, that mean lies within 3ε of an endpoint of $[0, 1]$.

Round each midpoint mean to that endpoint. The two constant inputs already have deterministic outputs. On every input the random minor disagrees with the rounded table with probability at most 3ε. A union bound over all 2^{2t+1} inputs gives one minor agreeing everywhere with this table of the two densities. □

G.2 Rotation for increasing operations

In this section we prove Lemma 2.8. Suppose that $g(S) \le g((S + d) \cup (S + d + 1))$ for every S. We evaluate g on random unions of cyclic intervals. First we find a bias at which a bounded number of intervals predicts the output, while both outputs occur with positive probability. The assumed inequality will allow us to rotate the predictor away from all its coordinates.

Fix a prime $p \ge 3$. Partition the cycle of length $n = pd + 1$ at cuts $id + t_i$, $0 \le i < p$, where the t_i are independent and uniform on $\{0, \dots, W - 1\}$ and $d = W + 2$. Give interval i, from cut i up to cut $i + 1$ in the cyclic order, an independent bit x_i of mean u. Let H be the value of g on the union of the intervals with bit one, and $Q(u) = \mathbb{E}H$. The function Q is nondecreasing.

Put $a = 2^{-m-2}$. For each $s \in \{m, m + 1\}$ there is an interior j/s with $a < Q(j/s) < 1 - a$. Otherwise round all the values $Q(j/s)$ to the nearer endpoint of $[0, 1]$. They specify a symmetric threshold by monotonicity. Indeed, choose the cuts once and send each interval independently to a uniform label in $[s]$. On every input of weight j the resulting minor has mean $Q(j/s)$, so a union bound gives positive probability of agreeing everywhere with the rounded threshold. Constants are also symmetric thresholds, so $g(\varnothing) = 0$ and $g(\mathbb{Z}_n) = 1$ under the hypothesis. The two interior grid points differ by at least $1/(m(m + 1))$. On the interval between them $Q \in [a, 1 - a]$ and $\min(u, 1 - u) \ge 1/(m + 1)$. This is the random identification argument of Brakensiek, Guruswami and Sandeep [17, Lemma 3.4].

To use the nearby-color predictor, we represent H by independent parameters in which it is increasing. At boundary i, the $W - 1$ sites $id, \dots, id + W - 2$ may be assigned to either neighboring interval. When the adjacent bits are 10 take a prefix of these sites, and when they are 01 take a suffix, in each case of length L_i uniform on $\{0, \dots, W - 1\}$. At 00 take none and at 11 take all. The other sites are fixed by their interval bits. With independent L_i this gives the original law conditional on the bits, since at a 10 boundary $t_i = L_i$ and at a 01 boundary $t_i = W - 1 - L_i$. At equal bits we may take $t_i = L_i$. The selected set now increases in each x_i and each L_i.

Fix a large integer M in terms of m only and let W be a power of M. Encode x_i by thresholding a uniform rank on $[M]$ at Mu, for $u \in \{0, 1/M, \dots, 1\}$. Encode L_i in base M by fresh uniform digits.

Moving every rank and digit by at most one allows bit bias changes of $1/M$ in either direction. The perturbed length lies between $L_i - D$ and $L_i + D$ clipped to $[0, W-1]$, where $D = (W-1)/(M-1)$. Each clipped length law has total variation distance at most $\theta = D/W \le 1/(M-1)$ from the uniform length law.

We bound the effect of changing all these length laws. Freeze all bits and lengths except the two adjacent bits at one boundary and its length. Write H_{ij} for the output with adjacent bits i, j. The values H_{00}, H_{11} are independent of this length. At fixed bit probabilities, changing the length law changes the conditional mean by at most $\theta(H_{11} - H_{00})$. Raising both bit probabilities by θ/η increases the mean by at least this amount, as long as they stay in $[\eta, 1-\eta]$. Indeed, for each fixed length first raise the left probability and then the right. The pivotal edges of $00, 10, 11$ contribute at least

$$(\theta/\eta)\,\eta\big[(H_{10} - H_{00}) + (H_{11} - H_{10})\big].$$

The same comparison in reverse bounds a downward change. Restoring one length law at a time and compensating its two neighbors uses each bit twice. Thus the means of the lower and upper parameter bounds for all neighboring inputs lie between $Q(u-\gamma)$ and $Q(u+\gamma)$, where $\gamma = 1/M + 2/(\eta(M-1))$ and $\eta = 1/(2(m+1))$, whenever $[u-\gamma, u+\gamma] \subseteq [\eta, 1-\eta]$.

For M sufficiently large the balanced interval found above contains more than $32/a$ disjoint intervals $[u-\gamma, u+\gamma]$ centered at grid points u. Since Q increases by at most one, one of these intervals has variation at most $a/32$. At its midpoint the upper and lower parameter bounds give the same output except with probability at most $a/32$, by monotonicity. Apply Lemma A.14 with error $a/32$ to the Boolean coloring of the uniform ranks and digits. Its predictor J thus agrees with H except with probability $a/16$. It reads digits or ranks from at most $b = b(m)$ interval indices, independently of W and p.

Let R rotate the bits and lengths forward by one interval. For the chosen set S, the translated cut positions in $S + d$ are the rotated cut positions except that cut zero lies one site earlier, since $n = pd + 1$. Passing from $S + d$ to $(S + d) \cup (S + d + 1)$ extends each run of ones by one site at its right end. Thus the latter union differs from the rotated union by at most $p + 1$ single-site cut moves. Conditional on the bits the raw cuts are independent and uniform, and the required shifts depend only on the bits. For fixed other cuts, moving one cut along its window changes the output at most once by monotonicity of g. Each unit move therefore changes the output with probability at most $1/W$, including when the window has already been shifted one site. The spacing $d = W + 2$ keeps adjacent windows disjoint during the comparison. By our assumed inequality, $\Pr[H = 1, H \circ R = 0] \le (p+1)/W$. Since both have the same law,

$$\Pr[H \ne H \circ R^j] \le 2j(p+1)/W \qquad (0 \le j < p).$$

Choose $p > b^2 + 2$ prime and then W a large enough power of M to make $2p(p+1)/W \le a/8$. Some rotation by j moves all indices of the predictor off themselves, since there are at most b^2 differences between them. Then $J, J \circ R^j$ are independent with common mean in $[a - a/16, 1 - a + a/16]$, so their disagreement probability exceeds $a/2$. The two prediction errors and the rotation estimate bound it by $a/4$. This contradiction proves Lemma 2.8.

G.3 Direction certificates on an injective chain

In this section we use direction certificates to apply the tight-cycle proof to injective chains.

Lemma G.2. *Let p be prime and $L > 1$ satisfy $L \equiv 1 \pmod p$. Let E be the set of tuples of $\mathrm{CYCLE}_p[L]$, and let $\rho_a : E \to \mathbb{Z}_L$ read entry a, for $0 \le a < p$. Suppose every operation of a Boolean function minion $\mathscr{M}$ has an increasing direction certificate of size at most b. If no increasing $F \in \mathscr{M}(E)$ satisfies*

$$F^{\rho_0} = \cdots = F^{\rho_{p-1}},$$

then $\mathrm{Ch}^{\mathrm{inj}}(\mathscr{M}, \mathscr{P})$.

Proof. Use the clock of Appendix A, with block labeling $\lambda : \mathbb{Z}_N \to \mathbb{Z}_L$ from Lemma A.12. For $f \in \mathscr{M}(I)$ color $x \in \mathbb{Z}_N^I$ by $c_f(x) = f^{\lambda \circ x}$. These colorings commute with minors. The palette has size $q = |\mathscr{M}(\mathbb{Z}_L)|$. Put $s = q(p-1)+1$ and $R = 2ps^2$, the distance bound in the intersection argument of Section A.2.3. Choose N divisible by pL with $N/(pL) > 4R + 4$.

For every $e \in E$ there is an interval J_e of starting positions $z \in \mathbb{Z}_N$ for which the p readings $\lambda(z - aN/p)$ give e, even when each position is independently moved by at most R steps. Indeed, write $L = p\ell + 1$. In units of clock blocks the turn is $-\ell - 1/p$. The starting block specifies the initial label, and its p equal subintervals specify which step crosses an extra block boundary. The extra crossing gives the increment $1 - p$ in the labels, and all other increments are 1 modulo L. This specifies all tuples of E. Take J_e to be the middle half of the corresponding subinterval, rounding inwards to clock positions. Put $\gamma = \min_e |J_e|/N > 0$.

Fix r large enough that

$$N^s p^s |E| (1 - \gamma^b)^r < 1/4.$$

For an increasing operation take an empty mark. For each other operation, take a maximal family of bounded increasing certificates with pairwise disjoint supports, stopping at r certificates. Here and below the bound on each certificate is b. Mark the union of the chosen supports. These marks have size at most rb.

First consider an injective chain of s operations, each increasing or having r disjoint certificates. For its faces use the coordinate lists of Lemma A.15 for the colorings c_f, with $2\varepsilon s N^s p^s < 1/4$. The face test applies to arbitrary finite colorings. We verify the intersection of these transported lists by the same cyclic covering. Suppose their intersection is empty. Label the terminal coordinates by missing faces and use the random offsets and the prescribed row-one input $x = x_1(t, \alpha)$ of (12), for all depth profiles and arm completions on the star product with at least one full row.

For a fixed profile and completion, the coordinates of x on the first mark are independent and uniform. Indeed, condition on all offsets except those of one full row k. Each marked coordinate v then receives the independent uniform offset at $\pi_{1k}(v)$, and injectivity makes these indices distinct. If the first operation has r chosen certificates, fix $e \in E$ and choose $e' \neq e$. Each certificate (A, D) has probability at least γ^b of putting all x_v, $v \in A$, in J_e and all x_v, $v \in D$, in $J_{e'}$. The r events are independent. The choice of r therefore gives probability greater than $3/4$ that every e receives such a certificate at every tested input. For an increasing first operation its empty certificate already has this property.

On this event choose arbitrary inputs y^a at distance at most R from $x - aN/p$. The columns $\theta(v) = (\lambda(y_v^a))_{0 \le a < p}$ belong to E by the clock sandwich, since $2R \le N/(pL)$. Each fiber $\theta^{-1}(e)$ contains the required coordinates of a certificate and avoids its forbidden coordinates. Hence $F = f_1^\theta$ is increasing in every coordinate. The phase colors are F^{ρ_a}, so the hypothesis excludes a common color in the p boxes of radius R.

On the other hand, with probability greater than $3/4$ every turn at a zero-depth row changes each color distance by at most $2(s-1)$. Indeed, deletion of that row gives exactly the face-test distribution: the first remaining full row supplies the uniform tuple, and the turned block misses its face list. Sum the face-test error over the at most $sN^s p^s$ choices of row, depths and arms, as in

Section A.2.3. Choose offsets for which both events hold. By Lemma A.16, some cell has the same color on all rotations. Changing the arms at its centers and pulling witnesses back to row one uses at most $2(s-1)^2(p-1)+s \le R$ steps, by the intersection argument in that section. This gives the excluded common color. Thus the face lists intersect.

It remains to allow operations with fewer than r disjoint certificates. Use $s+1$ rows. On each s-row face list the transported marks of its rows with such a maximal family, excluding increasing rows. Also list the transported face-test lists of all $(s-1)$-row subchains consisting of the other rows. Add an arbitrary coordinate if the list is empty. The sizes are uniformly bounded. If two rows $i<j$ with small maximal families occur, a certificate at row i transports to a bounded increasing certificate at j: injectivity on its support keeps the two parts disjoint, and pulling a block back proves the direction inequality. This support meets the mark at j by maximality. A common coordinate transported to the last row thus appears on every face, since each face retains one of these two rows. Otherwise choose s rows of the type already treated. Their subchain is injective, and its face lists have a common transported coordinate. Every face of the full chain contains at least one of those shorter faces, so the same coordinate works there. This gives the required injective chain homomorphism. □

References

[1] OpenAI. *Hardness of finding large independent sets in three-colorable graphs*. Preprint dated 24 September 2026, Theorem 1.1 and Lemmas 4.1–4.2. Public repository version.

[2] Y. Fei, D. Minzer, and S. Wang. *On the hardness of* 4*-to-*1 *games with perfect completeness*. ECCC TR26-179, first posted 14 September 2026. Corollaries 1.7 and 1.10 in arXiv:2610.12378v2, 22 October 2026.

[3] L. Barto, J. Bulín, A. Krokhin, and J. Opršal. *Algebraic approach to promise constraint satisfaction*. Journal of the ACM 68(4), Article 28, 2021. Published article. Full version.

[4] A. Krokhin, J. Opršal, M. Wrochna, and S. Živný. *Topology and adjunction in promise constraint satisfaction*. SIAM Journal on Computing 52(1) (2023), 38–79. Published article. Author version.

[5] L. Barto and M. Kozik. *Combinatorial Gap Theorem and Reductions between Promise CSPs*. SODA 2022, pp. 1204–1220. Published article. Theorem numbers refer to arXiv:2107.09423v1.

[6] L. Barto and M. Kozik. *A baby PCP theorem and new reductions for promise CSPs*. Unpublished manuscript, undated. Theorem 3.1, Section 3.1 and Appendices A–B.

[7] V. Dalmau and J. Opršal. *Local consistency as a reduction between constraint satisfaction problems*. LICS 2024, Article 29, pp. 1–15. Published article. Numbering follows the full version.

[8] L. Barto, M. Hadek, and D. Zhuk. *Toward a Uniform Algorithm and Uniform Reduction for Constraint Problems*. 2026 preprint, arXiv:2604.06335v1. Definition 3.2 and Theorems 6.7 and 6.14. `https://arxiv.org/html/2604.06335v1`.

[9] O. Reingold. *Undirected connectivity in log-space*. Journal of the ACM 55(4), Article 17, 2008. Published article. Author version.

[10] D. Banakh and M. Kozik. *Injective hardness condition for PCSPs*. LICS 2024, Article 8, pp. 1–10. Published article. Numbering follows the full version.

[11] D. Banakh and K. Michno. *Boolean PCSPs through the lens of Fourier Analysis*. Accepted for SODA 2027. Full version, 29 October 2026.

[12] M. Wrochna. *A note on hardness of promise hypergraph colouring*. 2022. `https://arxiv.org/abs/2205.14719v1`.

[13] *A nonlinear class with a combinatorial boundary homomorphism*. Unpublished project note, October 2026. `research/coloring-combinatorial/quorum-proof.txt`.

[14] F. P. Ramsey. *On a problem of formal logic*. Proceedings of the London Mathematical Society (2) 30 (1930), 264–286. `https://doi.org/10.1112/plms/s2-30.1.264`.

[15] C. C. Harner and R. C. Entringer. *Arc colorings of digraphs*. Journal of Combinatorial Theory, Series B 13 (1972), 219–225. `https://doi.org/10.1016/0095-8956(72)90057-3`.

[16] L. Barto, M. Kozik, and T. Niven. *The CSP Dichotomy Holds for Digraphs with No Sources and No Sinks (A Positive Answer to a Conjecture of Bang-Jensen and Hell)*, SIAM Journal on Computing 38(5) (2009), pp. 1782–1802. Published article. Author version.

[17] J. Brakensiek, V. Guruswami, and S. Sandeep. *Conditional Dichotomy of Boolean Ordered Promise CSPs*. TheoretiCS 2 (2023), Article 2, 1–22. `https://doi.org/10.46298/theoretics.23.2`.

[18] J. Brakensiek, V. Guruswami, M. Wrochna, and S. Živný. *The Power of the Combined Basic Linear Programming and Affine Relaxation for Promise Constraint Satisfaction Problems*. SIAM Journal on Computing 49(6) (2020), 1232–1248. Published article. Author version.

[19] D. Banakh, A. Barsukov and T.-V. Nakajima, *Towards Infinite PCSP: A Dichotomy for Monochromatic Cliques*, LICS 2026, LIPIcs 380, 13:1–13:28. `https://doi.org/10.4230/LIPIcs.LICS.2026.13`. Numbered references use the full version of 18 October 2026, `https://arxiv.org/abs/2605.09815v2`.

[20] A. Krokhin and D. Vagnozzi, *Approximating 1-in-3 SAT by Linearly Ordered Hypergraph 3-Colouring Is NP-Hard*, ICALP 2026, LIPIcs 374, 184:1–184:18. `https://doi.org/10.4230/LIPIcs.ICALP.2026.184`.

[21] V. Guruswami and S. Sandeep, *Rainbow Coloring Hardness via Low Sensitivity Polymorphisms*, SIAM Journal on Discrete Mathematics 34(1) (2020), 520–537. `https://doi.org/10.1137/19M127731X`. Numbered references use the APPROX/RANDOM 2019 version, LIPIcs 145, 15:1–15:17. `https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2019.15`.

[22] J. Brakensiek and V. Guruswami, *New Hardness Results for Graph and Hypergraph Colorings*, CCC 2016, LIPIcs 50, 14:1–14:27. https://doi.org/10.4230/LIPIcs.CCC.2016.14.

[23] J. Brakensiek and V. Guruswami, *Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy*, SIAM Journal on Computing 50(6) (2021), 1663–1700. https://doi.org/10.1137/19M128212X.

[24] Y. Zhou, *Circle configurations and the hardness of promise graph and hypergraph coloring*, Hexagon 2611.00619v1, 19 November 2026. https://hexagonmath.org/2611.00619v1.

[25] P. Austrin, A. Bhangale and A. Potukuchi, *Improved Inapproximability of Rainbow Coloring*, SODA 2020, 1479–1495. https://doi.org/10.1137/1.9781611975994.90. Numbered references use https://arxiv.org/abs/1810.02784.

[26] V. V. S. P. Bhattiprolu, V. Guruswami and E. Lee, *Approximate Hypergraph Coloring under Low-discrepancy and Related Promises*, APPROX/RANDOM 2015, LIPIcs 40, 152–174. https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2015.152.

[27] A. Krokhin and J. Opršal. *An Invitation to the Promise Constraint Satisfaction Problem.* ACM SIGLOG News 9(3) (2022), pp. 30–59. Published survey. Author version.

[28] J. Brakensiek and V. Guruswami, *The Quest for Strong Inapproximability Results with Perfect Completeness*, APPROX/RANDOM 2017, LIPIcs 81, 4:1–4:20. https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2017.4.

[29] P. Austrin, J. Håstad and B. Martinsson, *On the Usefulness of Promises*, SODA 2026, 6332–6378. https://doi.org/10.1137/1.9781611978971.228. Section references use https://arxiv.org/abs/2511.21450v1.

[30] M. Filakovský, T.-V. Nakajima, J. Opršal, G. Tasinato and U. Wagner, *Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs*, ACM Transactions on Computation Theory 18(2) (2026), Article 10. https://doi.org/10.1145/3779121. Page and footnote references use STACS 2024, LIPIcs 289, Article 34. https://doi.org/10.4230/LIPIcs.STACS.2024.34.

[31] T.-V. Nakajima, Z. Verwimp, M. Wrochna and S. Živný, *Complexity of Approximate Conflict-Free, Linearly-Ordered, and Nonmonochromatic Hypergraph Colourings*, ACM Transactions on Computation Theory (2026). https://doi.org/10.1145/3815179. Theorem numbers refer to https://arxiv.org/abs/2501.12062v3.

[32] V. Guruswami and E. Lee, *Strong Inapproximability Results on Balanced Rainbow-Colorable Hypergraphs*, Combinatorica 38(3) (2018), 547–599. https://doi.org/10.1007/s00493-016-3383-0. Numbered references use https://eccc.weizmann.ac.il/report/2014/043/download/.

[33] P. Austrin, V. Guruswami and J. Håstad, $(2+\varepsilon)$-*SAT Is NP-hard*, SIAM Journal on Computing 46(5) (2017), 1554–1573. https://doi.org/10.1137/15M1006507.

[34] I. Dinur, O. Regev and C. Smyth, *The Hardness of 3-Uniform Hypergraph Coloring*, Combinatorica 25(5) (2005), 519–535. https://doi.org/10.1007/s00493-005-0032-4.

[35] T. Austin. *On the failure of concentration for the ℓ_∞-ball*. Israel Journal of Mathematics **211** (2016), 221–238. Author version, Theorem 1.1.

[36] L. Russo, *An approximate zero-one law*, Z. Wahrscheinlichkeitstheorie verw. Gebiete **61** (1982), 129–139.

[37] B. Hanke, R. Sanyal, C. Schultz, and G. M. Ziegler. *Combinatorial Stokes formulas via minimal resolutions*. Journal of Combinatorial Theory, Series A **116** (2009), 404–420. Corollary 5.10 and Theorem 5.14.

[38] T. Prescott and F. E. Su. *A constructive proof of Ky Fan's generalization of Tucker's lemma*. Journal of Combinatorial Theory, Series A **111** (2005), 257–265. Theorem 2.

[39] A. Dold. *Simple proofs of some Borsuk–Ulam results*. In *Proceedings of the Northwestern Homotopy Theory Conference*, Contemporary Mathematics **19** (1983), 65–69.