

Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs

Filakovský, Marek; Nakajima, Tamio-Vesa; Opršal, Jakub; Tasinato, Gianluca; Wagner, Uli

DOI:

[10.4230/LIPIcs.STACS.2024.34](https://doi.org/10.4230/LIPIcs.STACS.2024.34)

License:

Creative Commons: Attribution (CC BY)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Filakovský, M, Nakajima, T-V, Opršal, J, Tasinato, G & Wagner, U 2024, Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs. in O Beyersdorff, MM Kanté, O Kupferman & D Lokshtanov (eds), *41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024)*, 34, Leibniz International Proceedings in Informatics (LIPIcs), vol. 289, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 34:1-34:19, 41st International Symposium on Theoretical Aspects of Computer Science, Clermont-Ferrand, France, 12/03/24. <https://doi.org/10.4230/LIPIcs.STACS.2024.34>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs

Marek Filakovský  

Masaryk University, Brno, Czech Republic

Tamio-Vesa Nakajima  

University of Oxford, UK

Jakub Opršal  

University of Birmingham, UK

Gianluca Tasinato  

ISTA, Klosterneuburg, Austria

Uli Wagner  

ISTA, Klosterneuburg, Austria

Abstract

A linearly ordered (LO) k -colouring of a hypergraph is a colouring of its vertices with colours $1, \dots, k$ such that each edge contains a unique maximal colour. Deciding whether an input hypergraph admits LO k -colouring with a fixed number of colours is NP-complete (and in the special case of graphs, LO colouring coincides with the usual graph colouring).

Here, we investigate the complexity of approximating the “linearly ordered chromatic number” of a hypergraph. We prove that the following promise problem is NP-complete: Given a 3-uniform hypergraph, distinguish between the case that it is LO 3-colourable, and the case that it is not even LO 4-colourable. We prove this result by a combination of algebraic, topological, and combinatorial methods, building on and extending a topological approach for studying approximate graph colouring introduced by Krokhin, Opršal, Wrochna, and Živný (2023).

2012 ACM Subject Classification Mathematics of computing \rightarrow Approximation algorithms; Mathematics of computing \rightarrow Algebraic topology; Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases constraint satisfaction problem, hypergraph colouring, promise problem, topological methods

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.34

Related Version *Full Version:* <https://arxiv.org/abs/2312.12981> [14]

Funding *Marek Filakovský:* This research was supported by Charles University (project PRIMUS/21/SCI/014), the Austrian Science Fund (FWF project P31312-N35), and MSCAfellow5_MUNI (CZ.02.01.01/00/22_010/0003229).

Tamio-Vesa Nakajima: This research was funded by UKRI EP/X024431/1 and by a Clarendon Fund Scholarship. All data is provided in full in the results section of this paper.

Jakub Opršal: This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 101034413.

Uli Wagner: This research was supported by the Austrian Science Fund (FWF project P31312-N35).



© Marek Filakovský, Tamio-Vesa Nakajima, Jakub Opršal, Gianluca Tasinato, and Uli Wagner;

licensed under Creative Commons License CC-BY 4.0

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).

Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov;

Article No. 34; pp. 34:1–34:19



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

Deciding whether a given finite graph is 3-colourable (or, more generally, k -colourable, for a fixed $k \geq 3$) was one of the first problems shown to be NP-complete by Karp [17]. Since then, the complexity of *approximating* the chromatic number of a graph has been studied extensively. The best-known polynomial-time algorithm approximates the chromatic number of an n -vertex graph within a factor of $O(n^{\frac{(\log \log n)^2}{(\log n)^3}})$ (Halldórsson [16]); conversely, it is known that the chromatic number cannot be approximated in polynomial time within a factor of $n^{1-\varepsilon}$, for any fixed $\varepsilon > 0$, unless $P = NP$ (Zuckerman [31]). However, this hardness result only applies to graphs whose chromatic number grows with the number of vertices, and the case of graphs with *bounded* chromatic number is much less well understood.

Given an input graph G that is promised to be 3-colourable, what is the complexity of finding a colouring of G with some larger number $k > 3$ of colours? Khanna, Linial, and Safra [19] showed that this problem is NP-hard for $k = 4$, and it is generally believed that the problem is NP-hard for any constant k . However, surprisingly little is known, and the only improvement and best result to date, hardness for $k = 5$, was obtained only relatively recently by Bulín, Krokhin, and Opršal [9]. On the other hand, the best polynomial-time algorithm, due to Kawarabayashi and Thorup [18], uses a number of colours (slightly less than $n^{1/5}$) that depends on the number n of vertices of the input graph.

More generally, it is a long-standing conjecture that finding a k -colouring of a c -colourable graph is NP-hard for all constants $k \geq c \geq 3$, but the complexity of this *approximate graph colouring* problem remains wide open. The results from [9] generalise to give hardness for $k = 2c - 1$ and all $c \geq 3$. For $c \geq 6$, this was improved by Wrochna and Živný [29], who showed that it is hard to colour c -colourable graphs with $k = \binom{c}{\lfloor c/2 \rfloor}$ colours. We remark that conditional hardness (assuming different variants of Khot's *Unique Games Conjecture*) for approximate graph coloring for all $k \geq c \geq 3$ was obtained by Dinur, Mossel, and Regev [12], Guruswami and Sandeep [15], and Braverman, Khot, Lifshitz, and Mulzer [6].

Given the slow progress on approximate graph colouring, we believe there is substantial value in developing and extending the available methods for studying this problem and related questions, and we hope that the present paper contributes to this effort. As our main result (Theorem 1.1 below), we establish NP-hardness of a relevant hypergraph colouring problem that falls into a general scope of *promise constraint satisfaction problems*; in the process, we considerably extend a topological approach and toolkit for studying approximate colouring that was introduced by Krokhin, Opršal, Wrochna, and Živný [21, 29, 23].

Graph colouring is a special case of the *constraint satisfaction problem (CSP)*, which has several different, but equivalent formulations. For us, the most relevant formulation is in terms of homomorphisms between relational structures. The starting point is the observation that finding a k -colouring of a graph G is the same as finding a *graph homomorphism* (an edge-preserving map) $G \rightarrow K_k$ where K_k is the complete graph with k vertices. The general formulation of the constraint satisfaction problem is then as follows (see Section 2.1 below for more details): Fix a relational structure \mathbf{A} (e.g., a graph, or a uniform hypergraph), which parametrises the problem. $\text{CSP}(\mathbf{A})$ is then the problem of deciding whether a given structure \mathbf{X} allows a homomorphism $\mathbf{X} \rightarrow \mathbf{A}$. One of the celebrated results in the complexity theory of CSPs is the Dichotomy Theorem of Bulatov [8] and Zhuk [30], which asserts that for every finite relational structure \mathbf{A} , $\text{CSP}(\mathbf{A})$ is either NP-complete, or solvable in polynomial time.

The framework of CSPs can be extended to *promise constraint satisfaction problems (PCSPs)*, which include approximate graph colouring. PCSPs were first introduced by Austrin, Guruswami, and Håstad [1], and the general theory of these problems was further

developed by Brakensiek and Guruswami [5], and by Barto, Bulín, Krokhin, and Opršal [3]. Formally, a PCSP is parametrised by two relational structures \mathbf{A} and \mathbf{B} such that there exists a homomorphism $\mathbf{A} \rightarrow \mathbf{B}$. Given an input structure \mathbf{X} , the goal is then to distinguish between the case that there is a homomorphism $\mathbf{X} \rightarrow \mathbf{A}$, and the case that there does not even exist a homomorphism $\mathbf{X} \rightarrow \mathbf{B}$ (these cases are distinct but not necessarily complementary, and no output is required in case neither holds); we denote this decision problem by $\text{PCSP}(\mathbf{A}, \mathbf{B})$. For example, $\text{PCSP}(K_3, K_k)$ is the problem of distinguishing, given an input graph G , between the case that G is 3-colourable and the case that G is not k -colourable. This is the *decision version* of the approximate graph colouring problem whose *search version* we introduced above. We remark that the decision problem reduces to the search version, hence hardness of the former implies hardness of the latter.

PCSPs encapsulate a wide variety of problems, including versions of hypergraph colouring studied by Dinur, Regev, and Smyth [13] and Brakensiek and Guruswami [4]. A variant of hypergraph colouring that is closely connected to approximate graph colouring and generalises (monotone¹) *1-in-3SAT* is *linearly ordered (LO) hypergraph colouring*. A linearly ordered k -colouring of a hypergraph H is an assignment of the colours $[k] = \{1, \dots, k\}$ to the vertices of H such that, for every hyperedge, the maximal colour assigned to elements of that hyperedge occurs exactly once. Note that for graphs, linearly ordered colouring is the same as graph colouring. Moreover, LO 2-colouring of 3-uniform hypergraphs corresponds to (monotone) 1-in-3SAT (by viewing the edges of the hypergraph as clauses). In the present paper, we focus on 3-uniform hypergraphs; whether such a graph has an LO k -colouring can be expressed as $\text{CSP}(\mathbf{LO}_k)$ for a specific relational structure \mathbf{LO}_k with one ternary relation (see Section 2.1); in particular, 1-in-3SAT corresponds to $\text{CSP}(\mathbf{LO}_2)$.

The promise version of LO hypergraph colouring was introduced by Barto, Battistelli, and Berg [2], who studied the *promise 1-in-3SAT* problem. More precisely, let \mathbf{B} be a fixed ternary structure such that there is a homomorphism $\mathbf{LO}_2 \rightarrow \mathbf{B}$. Then $\text{PCSP}(\mathbf{LO}_2, \mathbf{B})$ is the following decision problem: Given an instance \mathbf{X} of 1-in-3SAT, distinguish between the case that \mathbf{X} is satisfiable, and the case that there is no homomorphism $\mathbf{X} \rightarrow \mathbf{B}$. For structures \mathbf{B} with three elements, Barto et al. [2] obtained an almost complete dichotomy; the only remaining unresolved case is $\mathbf{B} = \mathbf{LO}_3$, i.e., the complexity of $\text{PCSP}(\mathbf{LO}_2, \mathbf{LO}_3)$. They conjectured that this problem is NP-hard, and more generally that $\text{PCSP}(\mathbf{LO}_c, \mathbf{LO}_k)$ is NP-hard for all $k \geq c \geq 2$ [2, Conjecture 27]. Subsequently, the following conjecture emerged and circulated as folklore (first formally stated by Nakajima and Živný [27]): $\text{PCSP}(\mathbf{LO}_2, \mathbf{B})$ is either solved by the *affine integer programming relaxation*, or NP-hard (see Ciardo, Kozik, Krokhin, Nakajima, and Živný [10] for recent progress in this direction).

Promise LO hypergraph colouring was further studied by Nakajima and Živný [26], who found close connections between promise LO hypergraph colouring and approximate graph colouring. In particular, they provide a polynomial time algorithm for LO-colouring 2-colourable 3-uniform hypergraphs with a superconstant number of colours, by adapting methods used for similar algorithms for approximate graph colouring, e.g., [18]. In the other direction, NP-hardness of $\text{PCSP}(\mathbf{LO}_k, \mathbf{LO}_c)$ for $4 \leq k \leq c$ follows relatively easily from NP-hardness of the approximate graph colouring $\text{PCSP}(K_{k-1}, K_{c-1})$, as was observed by Nakajima and Živný and by Austrin (personal communications).²

¹ In the present paper, we will only consider the monotone version of 1-in-3SAT, i.e., the case where clauses contain no negated variable, and we will often omit the adjective “monotone” in what follows.

² To see why, observe that $(\mathbf{LO}_k, \mathbf{LO}_c)$ promise primitive-positive defines (K_{k-1}, K_{c-1}) ; in particular, we can define $x \neq y$ by $\exists z \cdot R(z, z, x) \wedge R(z, z, y) \wedge R(x, y, z)$. We then see that if R is interpreted in \mathbf{LO}_k , then the required z exists if and only if $x \neq y$, as required.

Our main result is the following, which cannot be obtained using these arguments.

► **Theorem 1.1.** $\text{PCSP}(\mathbf{LO}_3, \mathbf{LO}_4)$ is NP-complete.

Apart from the intrinsic interest of LO hypergraph colouring, we believe that the main contribution of this paper is on a technical level, by extending the topological approach of [23] and bringing to bear more advanced methods from algebraic topology, in particular *equivariant obstruction theory*. To our knowledge, this paper is the first that uses these methods in the PCSP context; we view this as a “proof of concept” and believe these tools will be useful to make further progress on approximate graph colouring and related problems.

The proof of Theorem 1.1 has two main parts. For a natural number n , let $(\mathbf{LO}_3)^n$ be the n -fold power of the relational structure \mathbf{LO}_3 (see Section 2.2). In the first part of the proof, we use topological methods to show (Lemma 3.2 below) that with every homomorphism $f: (\mathbf{LO}_3)^n \rightarrow \mathbf{LO}_4$, we can associate an *affine map* $\chi(f): \mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$ (i.e., a map of the form $(x_1, \dots, x_n) \mapsto \sum_{i=1}^n \alpha_i x_i$, for some $\alpha_i \in \mathbb{Z}_3$ and $\sum_{i=1}^n \alpha_i \equiv 1 \pmod{3}$); moreover, the assignment $f \mapsto \chi(f)$ preserves natural *minor relations* that arise from maps $\pi: [n] \rightarrow [m]$, i.e., χ is a *minion homomorphism* (see Section 2.2 for the precise definitions).

In the second part of the proof, we show by combinatorial arguments that the maps $\chi(f): \mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$ form a very restricted subclass of affine maps: They are projections $\mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$, $(x_1, \dots, x_n) \mapsto x_i$ (Corollary 3.4). Theorem 1.1 then follows from a hardness criterion (Theorem 2.6) obtained as part of a general algebraic theory of PCSPs [3].

In a nutshell, topology enters in the first part of the proof as follows. First, with every homomorphism $f: (\mathbf{LO}_3)^n \rightarrow \mathbf{LO}_4$ we associate a continuous map $f_*: T^n \rightarrow P^2$, where T^n is the n -dimensional torus (the n -fold power of the circle S^1) and P^2 is a suitable target space that will be described in more detail later; moreover, the cyclic group \mathbb{Z}_3 naturally acts on both T^n and P^2 , and the map f_* preserves these symmetries (it is *equivariant*). This first step uses *homomorphism complexes* (a well-known construction in topological combinatorics that goes back to the work of Lovász [24], see Section 2.3). Second, we show that equivariant continuous maps $T^n \rightarrow P^2$, when considered up to a natural equivalence relation of symmetry-preserving continuous deformation (*equivariant homotopy*), are in bijection with affine maps $\mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$. This second step uses *equivariant obstruction theory*.³

We remark that, with some additional work, our method could be extended to prove NP-hardness of $\text{PCSP}(\mathbf{LO}_k, \mathbf{LO}_{2k-2})$ (but as remarked above, this already follows from known hardness results for approximate graph colouring).

2 Preliminaries

We use the notation $[n]$ for the n -element set $\{1, \dots, n\}$, and identify tuples $a \in A^n$ with functions $a: [n] \rightarrow A$, and we use the notation a_i for the i th entry of a tuple. We denote by 1_X the identity function on a set X .

2.1 Promise CSPs

We start by recalling some fundamental notions from the theory of promise constraint satisfaction problems, following the presentation of [3] and [22].

³ By contrast, the topological argument in [23] required understanding maps from T^n to the circle S^1 that preserve natural \mathbb{Z}_2 -symmetries on both spaces, again up to equivariant homotopy; such maps can be classified by more elementary arguments using the *fundamental group* because S^1 is 1-dimensional.

A *relational structure* is a tuple $\mathbf{A} = (A; R_1^{\mathbf{A}}, \dots, R_k^{\mathbf{A}})$, where A is a set, and $R_i^{\mathbf{A}} \subseteq A^{\text{ar}(R_i)}$ is a relation of arity $\text{ar}(R_i)$. The *signature* of \mathbf{A} is the tuple $(\text{ar}(R_1), \dots, \text{ar}(R_k))$. For two relational structures $\mathbf{A} = (A; R_1^{\mathbf{A}}, \dots, R_k^{\mathbf{A}})$ and $\mathbf{B} = (B; R_1^{\mathbf{B}}, \dots, R_k^{\mathbf{B}})$ with the same signature, a *homomorphism* from \mathbf{A} to \mathbf{B} , denoted $h: \mathbf{A} \rightarrow \mathbf{B}$, is a function $h: A \rightarrow B$ that preserves all relations, i.e., such that $h(a) \in R_i^{\mathbf{B}}$ for each $i \in \{1, \dots, k\}$ and $a \in R_i^{\mathbf{A}}$ where $h(a)$ denotes the componentwise application of h on the elements of a . To express the existence of such a homomorphism, we will also use the notation $\mathbf{A} \rightarrow \mathbf{B}$. The set of all homomorphisms from \mathbf{A} to \mathbf{B} is denoted by $\text{hom}(\mathbf{A}, \mathbf{B})$.

Our focus is on structures with a single ternary relation R , i.e., pairs $(A; R^{\mathbf{A}})$ with $R^{\mathbf{A}} \subseteq A^3$. Moreover, most structures in this paper have a *symmetric* relation, i.e., the relation $R^{\mathbf{A}}$ is invariant under permuting coordinates. Such structures can be also viewed as 3-uniform hypergraphs, keeping in mind that edges of the form (a, a, b) are allowed.

► **Definition 2.1** (Promise CSP). *Fix two relational structures such that $\mathbf{A} \rightarrow \mathbf{B}$. The promise CSP with template \mathbf{A}, \mathbf{B} , denoted by $\text{PCSP}(\mathbf{A}, \mathbf{B})$, is a computational problem that has two versions:*

- *In the search version of the problem, we are given a relational structure \mathbf{X} with the same signature as \mathbf{A} and \mathbf{B} , we are promised that $\mathbf{X} \rightarrow \mathbf{A}$, and we are tasked with finding a homomorphism $h: \mathbf{X} \rightarrow \mathbf{B}$.*
- *In the decision version of the problem, we are given a relational structure \mathbf{X} , and we must answer Yes if $\mathbf{X} \rightarrow \mathbf{A}$, and No if $\mathbf{X} \not\rightarrow \mathbf{B}$. (These cases are mutually exclusive since $\mathbf{A} \rightarrow \mathbf{B}$ and homomorphisms compose.)*

The decision version reduces to the search version; thus for proving the hardness of both versions of problems, it is sufficient to prove the hardness of the decision version of the problem, and in order to prove tractability of both versions, it is enough to provide an efficient algorithm for the search version.

To complete this section, we define the relational structure \mathbf{LO}_k , $k \in \mathbb{N}$, that appears in our main result. The domain of \mathbf{LO}_k is $\{1, \dots, k\}$, and \mathbf{LO}_k has one ternary relation, containing precisely those triples (a, b, c) which contain a *unique maximum*. In other words, $(a, b, c) \in R^{\mathbf{LO}_k}$ if and only if $a = b < c$, $a = c < b$, $b = c < a$, or all three elements a, b, c are distinct. For example, $(1, 1, 2)$ or $(1, 2, 3)$ are triples of the relation of \mathbf{LO}_3 , but not $(2, 2, 1)$.

2.2 Polymorphisms and a hardness condition

Our proof of Theorem 1.1 uses a hardness criterion (Theorem 2.6 below) obtained as part of a general algebraic theory of PCSPs developed in [3], which we will briefly review.

► **Definition 2.2.** *Given a structure \mathbf{A} , we define its n -fold power to be the structure \mathbf{A}^n with the domain A^n and*

$$R_i^{\mathbf{A}^n} = \{(a_1, \dots, a_{\text{ar}(R_i)}) \mid (a_1(i), \dots, a_{\text{ar}(R_i)}(i)) \in R_i^{\mathbf{A}} \text{ for all } i \in [n]\}$$

for each i .

An n -ary polymorphism from a structure \mathbf{A} to a structure \mathbf{B} is a homomorphism from \mathbf{A}^n to \mathbf{B} . We denote the set of all polymorphisms from \mathbf{A} to \mathbf{B} by $\text{pol}(\mathbf{A}, \mathbf{B})$, and the set of all n -ary polymorphisms by $\text{pol}^{(n)}(\mathbf{A}, \mathbf{B})$.⁴

⁴ Untraditionally, we use lowercase notation for polymorphisms to highlight that we are not considering any topology on them contrary to the homomorphism complexes introduced below.

Concretely, in the special case of structures with a ternary relation, a polymorphism is a mapping $f: A^n \rightarrow B$ such that, for all triples $(u_1, v_1, w_1), \dots, (u_n, v_n, w_n) \in R^{\mathbf{A}}$, we have

$$(f(u_1, \dots, u_n), f(v_1, \dots, v_n), f(w_1, \dots, w_n)) \in R^{\mathbf{B}}.$$

Polymorphisms are enough to describe the complexity of a PCSP up to certain log-space reductions. Loosely speaking, the more complex the polymorphisms are, the easier the problem is. We will use a hardness criterion that essentially states that the problem is hard if the polymorphisms have no interesting structure. To define what do we mean by interesting structure, we have to define the notions of *minor*, *minion* and *minion homomorphism*.

► **Definition 2.3.** Fix two sets A and B , and let $f: A^n \rightarrow B$, $\pi: [n] \rightarrow [m]$ be functions. The π -minor of f is the function $g: A^m \rightarrow B$ defined by $g(x) = f(x \circ \pi)$, i.e., such that

$$g(x_1, \dots, x_m) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$$

for all $x_1, \dots, x_m \in A$. We denote the π -minor of f by f^π .

Abstracting from the fact that the polymorphism of any template are closed under taking minors leads to the following notion of (*abstract*) *minions*:⁵

► **Definition 2.4.** An (abstract) minion \mathcal{M} is a collection of sets $\mathcal{M}^{(n)}$, where $n > 0$ is an integer, and mappings $\pi^{\mathcal{M}}: \mathcal{M}^{(n)} \rightarrow \mathcal{M}^{(m)}$ where $\pi: [n] \rightarrow [m]$ such that $\pi^{\mathcal{M}} \circ \sigma^{\mathcal{M}} = (\pi \circ \sigma)^{\mathcal{M}}$ for each π and σ , and $(1_{[n]})^{\mathcal{M}} = 1_{\mathcal{M}^{(n)}}$.⁶

The polymorphisms of a template \mathbf{A}, \mathbf{B} form a minion \mathcal{M} defined by $\mathcal{M}^{(n)} = \text{pol}^{(n)}(\mathbf{A}, \mathbf{B})$, and $\pi^{\mathcal{M}}(f) = f^\pi$. With a slight abuse of notation, we will use the symbol $\text{pol}(\mathbf{A}, \mathbf{B})$ for this minion. Conversely, if \mathcal{M} is an abstract minion, we will call $\pi^{\mathcal{M}}(f)$ the π -minor of f , and write f^π instead of $\pi^{\mathcal{M}}(f)$.

An important example is the minion of projections denoted by \mathcal{P} . Abstractly, it can be defined by $\mathcal{P}^{(n)} = [n]$ and $\pi^{\mathcal{P}} = \pi$. Equivalently, and perhaps more concretely, \mathcal{P} can also be described as follows: Given a finite set A with at least two elements and integers $i \leq n$, the i -th n -ary *projection* on A is the function $p_i: A^n \rightarrow A$ defined by $p_i(x_1, \dots, x_n) = x_i$. The set of coordinate projections is closed under minors as described above and forms a minion isomorphic to \mathcal{P} . In particular, \mathcal{P} is also isomorphic to the polymorphism minion $\text{pol}(\mathbf{LO}_2, \mathbf{LO}_2)$.

► **Definition 2.5.** A minion homomorphism from a minion \mathcal{M} to a minion \mathcal{N} is a collection of mappings $\xi_n: \mathcal{M}^{(n)} \rightarrow \mathcal{N}^{(n)}$ that preserve taking minors, i.e., such that for each $\pi: [n] \rightarrow [m]$, $\xi_m \circ \pi^{\mathcal{M}} = \pi^{\mathcal{N}} \circ \xi_n$. We denote such a homomorphism simply by $\xi: \mathcal{M} \rightarrow \mathcal{N}$, and write $\xi(f)$ instead of $\xi_n(f)$ when the index is clear from the context.⁷

Using the minion \mathcal{P} , we can now formulate the following hardness criterion (which follows from [3, Theorem 3.1 and Example 2.17] and can also be derived from [3, Corollary 5.2]; see also Section 5.1 of that paper for more details).

► **Theorem 2.6** ([3, corollary of Theorem 3.1]). For every promise template \mathbf{A}, \mathbf{B} such that there is a minion homomorphism $\xi: \text{pol}(\mathbf{A}, \mathbf{B}) \rightarrow \mathcal{P}$, $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is NP-complete.

⁵ Abstract minions as defined here are a generalization of so-called *function minions* defined in [3]; the relation between a function minion and an abstract minion is analogous to the distinction between a permutation group and a group.

⁶ In the language of category theory, a minion is defined as a functor from the category of finite sets to the category of sets, which satisfies a non-triviality condition: $\mathcal{M}(X) = \emptyset$ if and only $X = \emptyset$. The definition given abuses the fact that the sets $[n]$ form a skeleton of the category of finite sets.

⁷ A minion homomorphism is a natural transformation between the two functors.

2.3 Homomorphism complexes

We will need a number of notions from topological combinatorics, which we will review briefly now. We refer the reader to [25] for a detailed and accessible introduction (see also [23], in particular for further background on homomorphism complexes).

A *(finite, abstract) simplicial complex* K is finite system of sets that is downward closed under inclusion, i.e., $F \subseteq G \in K$ implies $F \in K$. The (finite) set $V = \bigcup K$ is called the set of *vertices* of K , and the sets in K are called *simplices* or *faces* of the simplicial complex. A *simplicial map* $f: K \rightarrow L$ between simplicial complexes is a map between the vertex sets that preserves simplices, i.e., $f(F) \in L$ for all $F \in K$.

An important way of constructing simplicial complexes is the following: Let P be a partially ordered set (poset). A *chain* in P is a subset $\{p_0, \dots, p_k\} \subseteq P$ such that $p_0 < p_1 < \dots < p_k$. The set of all chains in P is a simplicial complex, called the *order complex* of P . Note that an order-preserving map between posets naturally induces a simplicial map between the corresponding order complexes.

With every simplicial complex K , one can associate a topological space $|K|$, called the *underlying space* or *geometric realization* of K , as follows: Identify the vertex set of K with a set of points *in general position* in a sufficiently high-dimensional Euclidean space (here, general position means that the points in $F \cup G$ are affinely independent for all $F, G \in K$). Then, in particular, the convex hull $\text{conv}(F)$ is a geometric simplex for every $F \in K$, and the geometric realization can be defined as the union $|K| = \bigcup_{F \in K} \text{conv}(F)$ of these geometric simplices (see, e.g., [25, Lemma 1.6.2]). We also say that the simplicial complex K is a triangulation of the space $|K|$. Every simplicial map $f: K \rightarrow L$ between abstract simplicial complexes induces a continuous map $|f|: |K| \rightarrow |L|$ between their geometric realizations. In what follows, we will often blur the distinction between a simplicial complex and its geometric realization (especially when considering properties that do not depend on a particular triangulation).

Following [25, Section 5.9], we define homomorphism complexes as order complexes of the poset of *multihomomorphisms* from one structure to another.⁸

► **Definition 2.7.** *Suppose \mathbf{A}, \mathbf{B} are relational structures. A multihomomorphism from \mathbf{A} to \mathbf{B} is a function $f: A \rightarrow 2^B \setminus \{\emptyset\}$ such that, for each relational symbol R and all tuples $(u_1, \dots, u_k) \in R^{\mathbf{A}}$, we have that*

$$f(u_1) \times \dots \times f(u_k) \subseteq R^{\mathbf{B}}.$$

We denote the set of all such multihomomorphisms by $\text{mhom}(\mathbf{A}, \mathbf{B})$.

Multihomomorphisms are partially ordered by component-wise comparison, i.e., $f \leq g$ if $f(u) \subseteq g(u)$ for all $u \in A$. They can also be composed in a natural way, i.e., if $f \in \text{mhom}(\mathbf{A}, \mathbf{B})$ and $g \in \text{mhom}(\mathbf{B}, \mathbf{C})$, then $(g \circ f)(a) = \bigcup_{b \in f(a)} g(b)$ is a multihomomorphism from \mathbf{A} to \mathbf{C} .

► **Definition 2.8.** *Let \mathbf{A} and \mathbf{B} be two structures of the same signature. The homomorphism complex $\text{Hom}(\mathbf{A}, \mathbf{B})$ is the order complex of the poset of multihomomorphisms from \mathbf{A} to \mathbf{B} , i.e., the vertices of this simplicial complex are multihomomorphisms from \mathbf{A} to \mathbf{B} , and faces correspond to chains $f_1 < f_2 < \dots < f_k$ of such multihomomorphisms.*

⁸ There are several alternative definitions of homomorphism complexes that lead to topologically equivalent spaces; e.g., the definition given here is the *barycentric subdivision* of the version of the homomorphism complex defined in [23, Definition 3.3].

By the discussion above, every homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ induces a simplicial map $f_*: \text{Hom}(\mathbf{C}, \mathbf{A}) \rightarrow \text{Hom}(\mathbf{C}, \mathbf{B})$ between homomorphism complexes, and hence a continuous map between the corresponding spaces (defined on vertices by mapping a multihomomorphism m to the composition $f \circ m$, and then extended linearly).

In the case of graphs, the homomorphism complex $\text{Hom}(K_2, G)$ is commonly used⁹ to study graph colourings, including in [23]. In the present paper, we work instead with the homomorphism complex $\text{Hom}(\mathbf{R}_3, \mathbf{A})$ where \mathbf{R}_3 is the structure with 3 elements and all *rainbow tuples*, i.e., tuples (a, b, c) such that a, b , and c are pairwise distinct; this structure is a hypergraph analogue of the graph K_2 .

Note that a homomorphism $h: \mathbf{R}_3 \rightarrow \mathbf{A}$ can be identified with a triple $(h(1), h(2), h(3)) \in R^{\mathbf{A}}$; conversely, every triple $(a, b, c) \in R^{\mathbf{A}}$ also corresponds to a homomorphism as long as $R^{\mathbf{A}}$ is symmetric. Similarly, a multihomomorphism m can be identified with a triple $(m(1), m(2), m(3))$ of subsets of A such that $m(1) \times m(2) \times m(3) \subseteq R^{\mathbf{A}}$.

2.4 Group actions

Throughout this paper, we will work with actions of the cyclic group \mathbb{Z}_3 on various objects (relational structures, simplicial complexes, topological spaces, groups, etc.) by structure-preserving maps (homomorphisms, simplicial maps, continuous maps, etc.). Thinking of \mathbb{Z}_3 as the multiplicative group with three elements $1, \omega, \omega^2$ (with the understanding that $\omega^i \cdot \omega^j = \omega^{i+j \pmod{3}}$ and $\omega^0 = 1$), such an action is described by describing the action of the generator ω . Thus, specifying the action of \mathbb{Z}_3 on a structure \mathbf{A} amounts to specifying a homomorphism $\omega: \mathbf{A} \rightarrow \mathbf{A}$ such that $\omega^3 = 1_{\mathbf{A}}$ (hence, ω is necessarily an isomorphism; note that we are abusing notation here, writing ω both for the generator of the group and the isomorphism by which it acts). Analogously, an action of \mathbb{Z}_3 on a simplicial complex (or a topological space) is described by specifying a simplicial isomorphism (respectively, a homeomorphism) ω of order 3 from the complex (or space) to itself. We will mostly work with actions that are *free*, which in our special case of \mathbb{Z}_3 -actions simply means that ω has no fixed points.

In particular, consider the action of \mathbb{Z}_3 that acts on \mathbf{R}_3 by cyclically permuting elements. This action induces an action on multihomomorphisms $h: \mathbf{R}_3 \rightarrow \mathbf{A}$ by pre-composition, and this action extends naturally to an action of \mathbb{Z}_3 on $\text{Hom}(\mathbf{R}_3, \mathbf{A})$.¹⁰

It is not hard to show that the action on $\text{Hom}(\mathbf{R}_3, \mathbf{A})$ is free as long as \mathbf{A} has no constant tuples: If a multihomomorphism m is a fixed point of a non-trivial element of \mathbb{Z}_3 , then $m(1) = m(2) = m(3)$, and since $m(1) \neq \emptyset$ and $m(1) \times m(2) \times m(3) \subseteq R^{\mathbf{A}}$ then $R^{\mathbf{A}}$ contains a constant tuple (a, a, a) for any $a \in m(1)$. Consequently, we may observe that the action does not fix any face of the complex.

For every homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$, the induced simplicial map $f_*: \text{Hom}(\mathbf{R}_3, \mathbf{A}) \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{B})$ (defined on vertices by mapping multihomomorphism m to the composition $f \circ m$) is *equivariant*; as remarked above, we will often identify f_* with the corresponding continuous map between the underlying spaces.

⁹ Some papers use a different complex, the so-called *box complex*, that leads to a homotopically equivalent (see below) space.

¹⁰ This is analogous to the action of \mathbb{Z}_2 on graph homomorphism complexes $\text{Hom}(K_2, G)$ used, for example, in [23].

2.5 Homotopy

Two continuous maps $f, g: X \rightarrow Y$ between topological spaces are called *homotopic*, denoted $f \sim g$, if there is a continuous map $h: X \times [0, 1] \rightarrow Y$ such that $h(x, 0) = f(x)$ and $h(x, 1) = g(x)$; the map h is called a *homotopy* from f to g . Note that a homotopy can also be thought of as a family of maps $h(\cdot, t): X \rightarrow Y$ that varies continuously with $t \in [0, 1]$. In what follows, X and Y will often be given as simplicial complexes, but we emphasize that we will generally not assume that the maps (or homotopies) between them are simplicial maps. Two spaces X and Y are said to be *homotopy equivalent* if there are continuous maps $f: X \rightarrow Y$ and $g: Y \rightarrow X$ such that $fg \sim 1_Y$ and $gf \sim 1_X$.

These notions naturally generalize to the setting of spaces with group actions. If \mathbb{Z}_3 acts on two spaces X and Y then a continuous map $f: X \rightarrow Y$ is (\mathbb{Z}_3) -*equivariant* if f preserves the action, i.e., $f \circ \omega = \omega \circ f$. Two equivariant maps $f, g: X \rightarrow Y$ are said to be *equivariantly homotopic*, denoted by $f \sim_{\mathbb{Z}_3} g$, if there exists an *equivariant homotopy* between them, i.e., a homotopy $h: X \times [0, 1] \rightarrow Y$ from f to g such that all maps $h(\cdot, t): X \rightarrow Y$ are equivariant. We denote by $[X, Y]_{\mathbb{Z}_3}$ the set of all equivariant homotopy classes of (equivariant) maps from X to Y , i.e.,

$$[X, Y]_{\mathbb{Z}_3} = \{[f] \mid f: X \rightarrow Y \text{ is equivariant}\}$$

where $[f]$ denotes the set of all equivariant maps g such that $f \sim_{\mathbb{Z}_3} g$.

3 Overview of the proof

We give a brief overview of the proof and the core techniques used. The result is proved by a combination of topological, combinatorial, and algebraic methods. In particular, the hardness is provided by analysing the polymorphisms of the template together with a hardness criterion from [3], see Theorem 2.6. In short, our goal is to provide a minion homomorphism from the polymorphism minion $\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4)$ to the minion of projections \mathcal{P} (which is incidentally isomorphic to the polymorphism minion of 3SAT). The core of our contribution is in providing deep-enough understanding of polymorphisms of our template so that the minion homomorphism follows. The proof has two parts: topological and combinatorial.

3.1 Topology

The first part builds on the topological method introduced by Krokhin, Opršal, Wrochna, and Živný [21, 29, 23]. The core idea is that, similarly to approximate graph colouring, there are unreasonably many polymorphisms between the linear-ordering hypergraphs, but most of them are very similar. This means that each polymorphism contains a lot of noise but relatively little information. We use topology to remove this noise, and uncover a signal. This is done by considering the polymorphisms “up to homotopy” – essentially claiming that the homotopy class of a polymorphism carries the information and everything else is noise.

In order to formalise this idea, we consider for each hypergraph \mathbf{A} the topological space $\text{Hom}(\mathbf{R}_3, \mathbf{A})$. Consequently, we get that each the homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ induces a continuous map $f_*: \text{Hom}(\mathbf{R}_3, \mathbf{A}) \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{B})$. Consequently, we identify two homomorphisms f, g if f_* and g_* are homotopic to each other. The same is then extended to polymorphisms, although this requires to overcome a few subtle technical issues. The key observation for this extension is that the power of a homomorphism complex is homotopically equivalent to a homomorphism complex of the corresponding power (see, e.g., [20]).

This general idea requires a refinement to avoid trivial collapses, i.e., we have to avoid the case when f_* is homotopic to a constant map which is always a continuous map between topological spaces. In our case, this is avoided by keeping track of the action of \mathbb{Z}_3 on the spaces $\text{Hom}(\mathbf{R}_3, \mathbf{A})$ and $\text{Hom}(\mathbf{R}_3, \mathbf{B})$ described in the preliminaries. Consequently, we consider maps only up to \mathbb{Z}_3 -equivariant homotopy (note that the map f_* induced by a homomorphism is always equivariant). Further in this exposition, we will silently assume that the action is always present, and all notions are equivariant – the formal proof below is presented with the action in mind.

At this point we sketched how to construct a map that assigns to a polymorphism $f: \mathbf{A}^n \rightarrow \mathbf{B}$, an equivariant continuous map $f_*: \text{Hom}(\mathbf{R}_3, \mathbf{A}^n) \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{B})$. This map does not necessarily preserve minors, nevertheless, it preserves minors *up to homotopy*, i.e., for each $\pi: [n] \rightarrow [m]$, we have that $(f^\pi)_*$ and $(f_*)^\pi$ are equivariantly homotopic (this is since $\text{Hom}(\mathbf{R}_3, \mathbf{A})^n$ and $\text{Hom}(\mathbf{R}_3, \mathbf{A}^n)$ are only homotopically equivalent and not homeomorphic). This allows us to define a minion homomorphism between the polymorphism minion $\text{pol}(\mathbf{A}, \mathbf{B})$ and the minion of “homotopy classes of continuous maps” from powers of $\text{Hom}(\mathbf{R}_3, \mathbf{A})$ to $\text{Hom}(\mathbf{R}_3, \mathbf{B})$.

► **Definition 3.1.** *Let X and Y be two topological spaces with an action of G . The minion of homotopy classes of equivariant polymorphisms from X to Y is the minion $\text{hpol}(X, Y)$ defined by*

$$\text{hpol}^{(n)}(X, Y) = [X^n, Y]_G$$

and $[f]^\pi = [f^\pi]$.

Note that minors are well-defined in this minion since if f and g are equivariantly homotopic, then so are f^π and g^π for all maps π . Hence, we have a minion homomorphism


$$\zeta: \text{pol}(\mathbf{A}, \mathbf{B}) \rightarrow \text{hpol}(\text{Hom}(\mathbf{R}_3, \mathbf{A}), \text{Hom}(\mathbf{R}_3, \mathbf{B})).$$

This part of the proof follows [23], namely this minion homomorphism can be constructed by following the proof of [23, Lemma 3.22] while substituting \mathbf{R}_3 for \mathbf{K}_2 , and \mathbb{Z}_3 for \mathbb{Z}_2 . We give a general categorical proof in the full version of this paper [14, Appendix C].

In order to describe the minion $\text{hpol}(\text{Hom}(\mathbf{R}_3, \mathbf{A}), \text{Hom}(\mathbf{R}_3, \mathbf{B}))$, we need to classify all homotopy classes of maps between the corresponding topological spaces. The problem of classifying maps between two spaces up to homotopy is well-studied in algebraic topology, although it can be immensely difficult, e.g., maps between spheres of dimensions m and n (i.e., $[S^n, S^m]$) has been classified for many pairs m, n , but the classification for infinitely many remaining cases is still open, and it is considered to be a central open problem in algebraic topology. We take advantage of the topological methods developed to solve these problems. Moreover, we may simplify the matters considerably by replacing the spaces $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ with spaces that allow equivariant maps to and from, resp., these spaces, and are better behaved from the topological perspective. This is due to the fact that if there are equivariant maps $X' \rightarrow X$ and $Y \rightarrow Y'$, then there is a minion homomorphism

$$\eta: \text{hpol}(X, Y) \rightarrow \text{hpol}(X', Y').$$

This minion homomorphism is defined in the same way as a minion homomorphism from $\text{pol}(\mathbf{A}, \mathbf{B})$ to $\text{pol}(\mathbf{A}', \mathbf{B}')$ if \mathbf{A}', \mathbf{B}' is a *homomorphic relaxation* of \mathbf{A}, \mathbf{B} , i.e., if $\mathbf{A}' \rightarrow \mathbf{A}$ and $\mathbf{B} \rightarrow \mathbf{B}'$ [3, Lemma 4.8(1)]. To substantiate our choice of X' and Y' , let us start with describing some topological properties of the spaces $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$.



■ **Figure 1** Some representations of spaces $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ up to \mathbb{Z}_3 homotopy equivalence.

It may be observed that the space $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ is homotopically equivalent to the simplicial complex depicted in Fig. 1a where the \mathbb{Z}_3 acts on each row cyclically. We choose X' so that $X' \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$, and its powers are topologically simple but non-trivial. A natural choice is S^1 which can be obtained from the simplicial complex by removing all horizontal edges. The action of \mathbb{Z}_3 on the circle can be then equivalently described as a rotation by $2\pi/3$. Consequently, the powers of this space are n -dimensional tori T^n with component-wise (diagonal) action of \mathbb{Z}_3 ; by definition T^n is the n -th power of S^1 .

The space $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ is a bit more complicated, in particular it is not homotopically equivalent to a 1-dimensional space. Up to equivalence, it is the order complex of the partial order depicted in Fig. 1c where \mathbb{Z}_3 acts on rows cyclically.¹¹ It may be observed that $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ is simply connected, i.e., that $\pi_1(\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)) = 0$, and that $\pi_2(\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4))$ is a non-trivial group. Moreover, the action of \mathbb{Z}_3 on $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ induces a non-trivial action of \mathbb{Z}_3 on $\pi_2(\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4))$. The precise group and action is described in the full version of this article [14, Appendix A], nevertheless it is irrelevant for us at this point. The space that we use to replace $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$, and denote by P^2 , shares these two properties with $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$, and moreover $\pi_n(P^2) = 0$ for all $n > 2$. Spaces which have only one non-trivial homotopy group (and are sufficiently “nice”) are called *Eilenberg-MacLane spaces*, and denoted by $K(G, n)$ where $\pi_n(K(G, n)) = G$ is the only non-trivial homotopy group. These spaces are well-defined up to homotopy equivalence. They are also closely connected with cohomology: One of the core statements of obstruction theory provides a bijection $[X, K(G, n)] \simeq H^n(X; G)$ for each Abelian group G and $n \geq 1$. Consequently, it is much easier to classify maps into an Eilenberg-MacLane space up to homotopy. The space P^2 is in fact an Eilenberg-MacLane space $K(G, 2)$ where G is a suitable group with a free action of \mathbb{Z}_3 , it is chosen in such a way that it allows an equivariant homomorphism $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow P^2$ while allowing for much easier classification of maps into it.

Next, we prove that the minion $\text{hpol}(S^1, P^2)$ is isomorphic to the minion \mathcal{Z}_3 of affine maps modulo 3, i.e., maps $\mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$ of the form $(x_1, \dots, x_n) \mapsto \sum_{i=1}^n \alpha_i x_i$ where $\alpha_1, \dots, \alpha_n \in \mathbb{Z}_3$ are fixed constants such that $\sum_{i=1}^n \alpha_i = 1 \pmod{3}$. We construct this minion homomorphism by classifying equivariant continuous maps from T^n with the diagonal action of \mathbb{Z}_3 to P^2 . Since we are interested in equivariant maps and equivariant homotopy, we use a version of equivariant cohomology, called *Bredon cohomology*, introduced in [7]. For our purpose, this equivariant cohomology is defined analogously to regular cohomology except the coefficients

¹¹In the proof we will not need such a precise description of the space, and we will only provide an equivariant map $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow L_4$ where L_4 is the space represented by the poset in Fig. 1c.

have a \mathbb{Z}_3 -action, and this action together with the action of \mathbb{Z}_3 on the space is taken into account in all computations. The space P^2 has the property that for every \mathbb{Z}_3 -space X such that there is an equivariant map $X \rightarrow P^2$, there is a bijection $[X, P^2]_{\mathbb{Z}_3} \simeq H_{\mathbb{Z}_3}^2(X; G)$ where G is the group with a \mathbb{Z}_3 action described above. Again, this is a consequence of the equivariant obstruction theory. We then compute that

$$H_{\mathbb{Z}_3}^2(T^n; G) \simeq \mathbb{Z}_3^{n-1},$$

and hence observe that there are 3^{n-1} elements in $\text{hpol}^{(n)}(S^1, P^2)$. This means that $\text{hpol}(S^1, P^2)$ and \mathcal{X}_3 have the same number of elements of each arity. To obtain the required minion isomorphism, we provide a minion homomorphism

$$\mathcal{X}_3 \rightarrow [S^1, P^2]_{\mathbb{Z}_3},$$

and show that it is injective. More precisely, this homomorphism is given by assigning to each affine map $f: \mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$ (or a tuple of its coefficients), a continuous map $\mu(f): T^n \rightarrow P^2$, and showing that if $f \neq g$ then $\mu(f)$ and $\mu(g)$ are not equivariantly homotopic, and that $\mu(f^\pi)$ and $\mu(f)^\pi$ are equivariantly homotopic for all π . Since both minions have the same number of elements of each arity (and this number is finite), μ is bijective, and hence a minion isomorphism. All these computations are presented in detail in the full version of this paper [14, Appendix B].

The above isomorphism together with the composition of ζ , η , and ξ provides the following lemma.

► **Lemma 3.2.** *There is a minion homomorphism $\chi: \text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \mathcal{X}_3$ where \mathcal{X}_3 denotes the minion of affine maps over \mathbb{Z}_3 .*

This minion homomorphism is not enough to prove NP-hardness. Although we could conclude from it, for example, that $\text{PCSP}(\mathbf{LO}_3, \mathbf{LO}_4)$ is not solved by any level of Sherali-Adams hierarchy (this is a direct consequence of [11, Theorems 3.3 and 5.2]). To provide hardness, we need to further analyse the image of χ which is done using combinatorial arguments.

3.2 Combinatorics

In the second part, which is a combinatorial argument presented in Section 4, we show that the image of χ avoids all the affine maps except of projections. This is done by analysing binary polymorphisms from \mathbf{LO}_3 to \mathbf{LO}_4 .

We use the notion of *reconfiguration* of homomorphisms to achieve this. Loosely speaking, a homomorphism f is reconfigurable to a homomorphism g if there is a path of homomorphism starting with f and ending with g such that neighbouring homomorphisms differ in at most one value. (For graphs and hypergraphs without tuples with repeated entries this can be taken as a definition, but with repeated entries there are two sensible notions of reconfigurations that do not necessarily align.) The connection between reconfigurability and topology was described by Wrochna [28], and we use these ideas to connect reconfigurability with our minion homomorphism ξ .

We show that any binary polymorphism $f: \mathbf{LO}_3^2 \rightarrow \mathbf{LO}_4$ is reconfigurable to an essentially unary polymorphism, i.e., that there is an increasing function $h: \mathbf{LO}_3 \rightarrow \mathbf{LO}_4$ such that f is reconfigurable to the map $(x, y) \mapsto h(x)$ or to the map $(x, y) \mapsto h(y)$. Further, we show that if f and g are reconfigurable to each other, then $\chi(f) = \chi(g)$. Together with the above,

this means the image of $\chi_2: \text{hpol}^{(2)}(S^1, P^2) \rightarrow \mathcal{L}_3^{(2)}$ omits an element. More precisely, we have the following lemma where \mathcal{P}_3 denotes the minion of projections on a three element set (which is a subminion of \mathcal{L}_3).

► **Lemma 3.3.** *For each binary polymorphism $f \in \text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$, $\chi(f) \in \mathcal{P}_3^{(2)}$.*

This lemma is then enough to show that the image of χ omits all affine maps except projections.

► **Corollary 3.4.** *χ is a minion homomorphism $\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \mathcal{P}_3$.*

Proof. We show that if a subminion $\mathcal{M} \subseteq \mathcal{L}_3$ contains any non-projection then it contains the map $g: (x, y) \mapsto 2x + 2y$. Let $f \in \mathcal{M}^{(n)}$ depends on at least 2 coordinates, and let $f(x_1, \dots, x_n) = \alpha_1 x_1 + \dots + \alpha_n x_n$. First assume that $\alpha_i = 2$ for some i . Then the binary minor given by $\pi: [n] \rightarrow [2]$ defined by $\pi(i) = 1$ and $\pi(j) = 2$ if $j \neq i$ is g since its first coordinate is 2 and the second is $1 - 2 = 2 \pmod{3}$. Otherwise, we have that $\alpha_i \in \{0, 1\}$ for all i . In particular, there are $i \neq j$ such that $\alpha_i = \alpha_j = 1$ since f depends on at least 2 coordinates. Consequently, the minor defined by $\pi': [n] \rightarrow [2]$ where $\pi'(i) = \pi'(j) = 1$ and $\pi'(k) = 2$ for $k \notin \{i, j\}$ is again g by a similar argument.

Finally, the image of ξ is a subminion of \mathcal{L}_3 , and since it omits g and every subminion of \mathcal{L}_3 contains \mathcal{P}_3 , it is equal to \mathcal{P}_3 which yields the desired. ◀

As mentioned before, the above corollary combined with Theorem 2.6 provides the main result of this paper, the NP-completeness of $\text{PCSP}(\mathbf{LO}_3, \mathbf{LO}_4)$ (Theorem 1.1).

4 Combinatorics of reconfigurations

The goal of this section is a careful combinatorial analysis of the binary polymorphisms. In particular, we will describe how the minion homomorphism $\xi: \text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \mathcal{P}$ acts on binary polymorphisms. This is the key to the argument that the image of ξ is the projection minion and the whole of $\text{pol}(\mathbb{Z}_3)$.

We say that two polymorphisms $f, g \in \text{pol}^{(n)}(\mathbf{LO}_3, \mathbf{LO}_4)$ are *reconfigurable* one to the other if a path between f and g exists within the homomorphism complex $\text{Hom}(\mathbf{LO}_3^n, \mathbf{LO}_4)$. (Note that every polymorphism is a homomorphism $\mathbf{LO}_3^n \rightarrow \mathbf{LO}_4$, and hence a vertex of the homomorphism complex.)

We will use the following combinatorial criterion that ensures that two polymorphisms are reconfigurable to each other. The proof is subtly dependent on some properties of the structure \mathbf{LO}_4 .

► **Lemma 4.1.** *Let \mathbf{A} be a symmetric relational structure. If $f, g: \mathbf{A} \rightarrow \mathbf{LO}_4$ are two homomorphisms such that f and g differ in exactly one value, i.e., there is $d \in A$ such that for all $a \in A \setminus \{d\}$ we have $f(a) = g(a)$, then f and g are reconfigurable.*

Proof. We first claim that under the above assumption, the multifunction $m: A \rightarrow 2^{[4]}$ given by $m(a) = \{f(a), g(a)\}$ is a multihomomorphism. Assume that $(a, b, c) \in R^{\mathbf{A}}$. Observe that for any $x \in A \setminus \{d\}$ we have $f(x) = g(x)$ and hence $m(x) = \{f(x)\} = \{g(x)\}$. We now have cases depending on how many times d appears in $\{a, b, c\}$.

d does not appear. In this case $m(a) \times m(b) \times m(c) = \{(f(a), f(b), f(c))\} \subseteq R^{\mathbf{LO}_4}$.

d appears once. Suppose $d = a, d \neq b, d \neq c$; then $m(a) \times m(b) \times m(c) = \{f(a), g(a)\} \times \{f(b)\} \times \{f(c)\} = \{(f(a), f(b), f(c)), (g(a), g(b), g(c))\} \subseteq R^{\mathbf{LO}_4}$, as $f(b) = g(b), f(c) = g(c)$.

34:14 Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs

d appears twice. Suppose $d = a = b, d \neq c$; then as $(f(a), f(b), f(c)) = (f(d), f(d), f(c)) \in R^{\mathbf{LO}_4}$ and likewise $(g(d), g(d), g(c)) \in R^{\mathbf{LO}_4}$, we have $f(d) < f(c)$ and $g(d) < g(c) = f(c)$. Consequently, $m(a) \times m(b) \times m(c) = \{f(d), g(d)\}^2 \times \{f(c)\} \subseteq R^{\mathbf{LO}_4}$, since every tuple has a unique maximum, namely $f(c)$.

d appears thrice. This case (i.e., $d = a = b = c$) is impossible, as $\mathbf{A} \rightarrow \mathbf{LO}_4$, and thus \mathbf{A} has no constant tuples.

Thus m is a multihomomorphism in all cases.

We can now define a path $p: [0, 1] \rightarrow \text{Hom}(\mathbf{LO}_3, \mathbf{LO}_4)$ by $p(0) = f, p(1/2) = m, p(1) = g$, and extending linearly. \blacktriangleleft

We note, without a proof, if f and g are reconfigurable, then there is a sequence $f = f_0, \dots, f_k = g$ such that f_i and f_{i+1} differ in exactly one point. A polymorphism $f \in \text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$ has, as its domain, the set $[3]^2$, and thus it can naturally be represented as a matrix:

$$\begin{array}{ccc} f(1, 1) & f(1, 2) & f(1, 3) \\ f(2, 1) & f(2, 2) & f(2, 3) \\ f(3, 1) & f(3, 2) & f(3, 3) \end{array}.$$

When we speak of “rows” or “columns” of f this is what is meant.

We show the following lemma from which we will be able to derive that each binary polymorphism is reconfigurable to an essentially unary one. (Recall that a function $f: A^n \rightarrow B$ is essentially unary if it depends on at most one input coordinate.) The lemma is an analogue of the *Trash Colour Lemma* for polymorphisms from K_d to K_{2d-2} .

► **Lemma 4.2.** *For each $f \in \text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$ there exists an increasing function $h \in \text{pol}^{(1)}(\mathbf{LO}_3, \mathbf{LO}_4)$, a coordinate $i \in \{1, 2\}$, and a colour $t \in [4]$ (called trash colour) such that*

$$f(x_1, x_2) \in \{h(x_i), t\}$$

for all $x_1, x_2 \in [3]$.

Proof. Throughout we will implicitly use the fact that if $a < b$ and $c < d$ then $f(a, c) < f(b, d)$, as $((a, c), (a, c), (b, d)) \in R^{\mathbf{LO}_3}$.

First, we claim that every colour $c \in [4]$ appears inside only one row or only one column of f , i.e., that either there is $a \in [3]$ such that $f(x, y) = c$ implies $x = a$, or there is $b \in [3]$ such that $f(x, y) = c$ implies $y = b$. For contradiction, assume that this is not the case, i.e., there are x, y and $x', y' \in [3]$ such that $f(x, y) = f(x', y') = c$, $x \neq x'$, and $y \neq y'$. The claim is proved by case analysis as follows. First, observe that either $x < x'$ and $y > y'$, or $x > x'$ and $y < y'$, since otherwise (x, y) and (x', y') are comparable, and hence $f(x, y) \neq f(x', y')$. Since the two cases are symmetric, we may assume without loss of generality that $x < x'$ and $y > y'$. Furthermore, since $((x, y), (x', y'), (x, y')) \in R^{\mathbf{LO}_3}$, and $f(x, y) = f(x', y') = c$, we have $f(x, y') > c$. Similarly, as $x' > x, y > y'$ we have that $f(x', y) > f(x, y') > c$. This means that $c \in \{1, 2\}$. We consider each case separately.

Case $c = 1$. We claim that $x = y' = 1$ since if $x > 1$, then $f(1, y') < f(x, y) = 1$, and similarly if $y' > 1$. This implies that $f(1, 1) > 1$ since $((1, 1), (x, x'), (y, y')) = ((1, 1), (1, y), (x', 1)) \in R^{\mathbf{LO}_3}$ and $f(x, y) = f(x', y') = 1$. As $1 < f(1, 1) < f(2, 2) < f(3, 3) \leq 4$, we have that $f(1, 1) = 2, f(2, 2) = 3$, and $f(3, 3) = 4$. We now have three cases.

- $y = 3$. We argue that $f(1, 2)$ has no possible value. First, the value 1 is not possible since $((1, 2), (x, y), (x', y')) = ((1, 2), (1, 3), (x', 1)) \in R^{\mathbf{LO}_3^2}$, $f(x, y) = 1$, and $f(x', y') = 1$. $f(1, 2) = 2$ is not possible since $((1, 2), (1, 1), (x', y')) = ((1, 1), (1, 2), (x', 1)) \in R^{\mathbf{LO}_3^2}$, and $f(x', y') = 1, f(1, 1) = 2$. $f(1, 2) = 3$ is not possible since $((1, 2), (2, 2), (x, y)) = ((1, 2), (2, 2), (1, 3)) \in R^{\mathbf{LO}_3^2}$, and $f(x, y) = 1, f(2, 2) = 3$. Finally, $f(1, 2) < f(3, 3) = 4$, so $f(1, 2) \neq 4$.
- $x' = 3$. Here the contradiction follows analogously to the previous case.
- $x' = y = 2$. We consider the pair of values $f(1, 3)$ and $f(3, 1)$. First, we have $f(1, 3) > f(1, 2) = f(x, y) = 1$ and $f(3, 1) > f(2, 1) = f(x', y') = 1$. As $((1, 3), (1, 1), (x', y')) = ((1, 3), (1, 1), (2, 1)) \in R^{\mathbf{LO}_3^2}$ and $f(1, 1) = 2, f(x', y') = 2$ we have that $f(1, 3) \neq 2$; symmetrically $f(3, 1) \neq 2$. We also have $f(1, 3) \neq 3$ since $((1, 3), (x, y), (2, 2)) = ((1, 3), (1, 2), (2, 2)) \in R^{\mathbf{LO}_3^2}$ and $f(1, 2) = 1, f(2, 2) = 3$; symmetrically $f(3, 1) \neq 2$. Thus $f(1, 3) = f(3, 1) = 4$. However, then $(f(1, 2), f(1, 3), f(3, 1)) = (1, 4, 4) \notin R^{\mathbf{LO}_4}$, which is not possible, as $((1, 2), (1, 3), (3, 1)) \in R^{\mathbf{LO}_3^2}$, which yields our contradiction.

Case $c = 2$. As $f(x', y) > f(x, y') > c = 2$, we have that $f(x, y') = 3$ and $f(x', y) = 4$. Since $f(x, y') = 3$ then either $x > 1$ or $y' > 1$, otherwise $f(3, 3) > f(2, 2) > f(1, 1) = 3$ yields a contradiction. By symmetry it is enough to discuss the case $y' = 2$ and $y = 3$. Finally, we have $f(x, 1) < f(x', 2) = 2$, hence $f(x, 1) = 1$ which is in contradiction with

$$(1, 2, 2) = (f(x, 1), f(x', 2), f(x, 3)) \in R^{\mathbf{LO}_4}.$$

Thus we get a contradiction in all cases, and hence each colour appears in only one row or only one column.

We say that a colour $c \in [4]$ is of *column* type if $f(x, y) = c$ implies $x = a_c$ for some fixed $a_c \in [3]$, and is of *row* type if $f(x, y) = c$ implies $y = b_c$ for some $b_c \in [3]$. Note that a colour can be both row and column type, in which case we may choose either. We claim that there are at least three colours that share a type – otherwise there are two colours of row type and two colours of column type which would leave an element of \mathbf{LO}_3^2 uncoloured. A similar observation also yields that there has to be three colours of the same type that cover all rows or all columns, i.e., such that the constants a_c or b_c (depending on the type) are pairwise distinct. Let us assume they are of the column type; the other case is symmetric. Further, we may assume that the fourth colour is of the row type, since if two colours share a column, then one of the colours appears only once, and can be therefore considered to be of row type.

We define $h(a)$ to be the colour c of column type with $a_c = a$, then we have $f(x, y) \in \{h(x), t\}$ where t is the colour of the row type. Finally, we argue that h is increasing. This is since there are $y < y'$ with $y \neq b_t$ and $y' \neq b_t$, and consequently

$$h(1) = f(1, y) < f(2, y') = h(2) = f(2, y) < f(3, y') = h(3).$$

This concludes the proof of the lemma. ◀

► **Lemma 4.3.** *Every binary polymorphism $f \in \text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$ is reconfigurable to an essentially unary polymorphism.*

Proof. The proof relies on Lemma 4.2. We prove our result by induction on the number of appearances of the trash colour. The result is clear if the trash colour never appears; so assume it appears at least once. Thus suppose without loss of generality that $f(x, y) \in \{h(x), t\}$ for some increasing $h \in \text{pol}^{(1)}(\mathbf{LO}_3, \mathbf{LO}_4)$, and that in particular $f(x_0, y_0) = t$. Furthermore,

suppose that among all such pairs, (x_0, y_0) is the one that maximises x_0 . We claim that $f'(x, y)$, which is equal to $f(x, y)$ everywhere except that $f'(x_0, y_0) = h(x_0)$ is also a polymorphism, which gives us our inductive step.

Consider any $((x, y), (x', y'), (x'', y'')) \in R^{\mathbf{LO}_3^2}$; if $(x_0, y_0) \notin \{(x, y), (x', y'), (x'', y'')\}$, then $(f'(x, y), f'(x', y'), f'(x'', y'')) = (f(x, y), f(x', y'), f(x'', y'')) \in R^{\mathbf{LO}_4}$, so assume without loss of generality that $(x'', y'') = (x_0, y_0)$. We now have two cases, depending on where the unique maximum of $(f(x, y), f(x', y'), f(x_0, y_0)) \in R^{\mathbf{LO}_4}$ falls.

- **$f(x, y)$ is the unique maximum.** In this case, $f(x, y) > f(x_0, y_0) = t$ and $f(x, y) > f(x', y')$. We must show that $f'(x_0, y_0) = h(x_0) \neq f(x, y)$. Since we know that $f(x, y) \neq t$ and thus $f(x, y) = h(x)$, and furthermore that h is increasing, this is the same as showing that $x \neq x_0$. Suppose for contradiction that $x = x_0$; thus $x' > x$. If $f(x', y) = h(x') > h(x)$, then $f(x, y)$ would not be the unique maximum, so $f(x', y) = t$. This contradicts the choice of (x_0, y_0) , as $x' > x_0$.
- **$f(x', y')$ is the unique maximum.** This case is identical to the previous case.
- **$f(x_0, y_0)$ is the unique maximum.** It follows that $f(x, y) < t$ and $f(x', y') < t$, hence $f(x, y) = h(x)$ and $f(x', y') = h(x')$. Thus since $(x, x', x_0) \in R^{\mathbf{LO}_3}$ and h is increasing, it follows that $(f'(x, y), f'(x', y'), f'(x_0, y_0)) = (h(x), h(x'), h(x_0)) \in R^{\mathbf{LO}_4}$.

Thus we see that this f' is indeed a polymorphism, and contains one fewer trash colour. Thus our conclusion follows. ◀

In Figure 2, we can see the reconfiguration graph of $\text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$. This shows how one can reconfigure all polymorphisms to essentially unary ones. In the diagram, we show a polymorphism in its matrix representation.

It can be also observed that unary polymorphisms that depend on the same coordinate are reconfigurable to each other. Moreover, since every connected component of $\text{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$ contains a homomorphism, and hence a unary one, we can derive from these observation that $\text{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$ has at most two connected components. In the full version [14, Appendix B], we also prove that it has at least two components using topological methods.

Finally, we conclude with the statement that we actually use in the proof, which follows from well-known properties of homomorphism complexes.

► **Lemma 4.4.** *Let \mathbf{A} , \mathbf{B} , and \mathbf{C} be three structures, G a group acting on \mathbf{A} , and assume that $f, g \in \text{hom}(\mathbf{B}, \mathbf{C})$ are reconfigurable. Then the induced maps $f_*, g_* : \text{Hom}(\mathbf{A}, \mathbf{B}) \rightarrow \text{Hom}(\mathbf{A}, \mathbf{C})$ are G -homotopic.*

Proof. First, observe that the composition of multihomomorphisms as a map $\text{mhom}(\mathbf{A}, \mathbf{B}) \rightarrow \text{mhom}(\mathbf{B}, \mathbf{C}) \rightarrow \text{mhom}(\mathbf{A}, \mathbf{C})$ is monotone. This means that the composition extends linearly to a continuous map

$$c : \text{Hom}(\mathbf{B}, \mathbf{C}) \times \text{Hom}(\mathbf{A}, \mathbf{B}) \rightarrow \text{Hom}(\mathbf{A}, \mathbf{C})$$

(see also [20, Section 18.4.3]). Since the composition is associative, we obtain that the map c is equivariant (under an action of any automorphism of \mathbf{A} on the second coordinate).

Finally, we have that $f_*(x) = c(f, x)$ by the definition of f_* , and analogously, $g_*(x) = c(g, x)$. Consequently, if $h : [0, 1] \rightarrow \text{Hom}(\mathbf{B}, \mathbf{C})$ is an arc connecting f and g , i.e., such that $h(0) = f$ and $h(1) = g$, then the map $H : [0, 1] \times \text{Hom}(\mathbf{A}, \mathbf{B}) \rightarrow \text{Hom}(\mathbf{A}, \mathbf{C})$ defined by

$$H(t, x) = c(h(t), x)$$

is a homotopy between f_* and g_* . This H is also equivariant since c is equivariant. ◀

The following corollary then follows directly from the above and Lemma 4.3.

► **Corollary 4.5.** *For every binary polymorphism $f \in \text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$, the induced map $f_*: \text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)^2 \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ is equivariantly homotopic either to the map $(x, y) \mapsto i_*(x)$, or to the map $(x, y) \mapsto i_*(y)$ where $i: \mathbf{LO}_3 \rightarrow \mathbf{LO}_4$ is the inclusion.*

References

- 1 Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$ -Sat is NP-hard. *SIAM Journal on Computing*, 46(5):1554–1573, 2017. doi:10.1137/15M1006507.
- 2 Libor Barto, Diego Battistelli, and Kevin M. Berg. Symmetric promise constraint satisfaction problems: Beyond the Boolean case. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16–19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPICs*, pages 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.STACS.2021.10.
- 3 Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–66, 2021. doi:10.1145/3457606.
- 4 Joshua Brakensiek and Venkatesan Guruswami. New hardness results for graph and hypergraph colorings. In Ran Raz, editor, *Conference on Computational Complexity (CCC 2016)*, volume 50 of *LIPICs*, pages 14:1–14:27, Dagstuhl, Germany, 2016. Schloss Dagstuhl–LZI. doi:10.4230/LIPICs.CCC.2016.14.
- 5 Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric Boolean dichotomy. *SIAM Journal on Computing*, 50(6):1663–1700, 2021. doi:10.1137/19M128212X.
- 6 Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. An invariance principle for the multi-slice, with applications. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021*, pages 228–236. IEEE Computer Soc., Los Alamitos, CA, 2022. doi:10.1109/FOCS52979.2021.00030.
- 7 Glen E. Bredon. *Equivariant Cohomology Theories*, volume 34 of *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, 1967. doi:10.1007/BFb0082690.
- 8 Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, Berkeley, CA, USA, 2017. IEEE. doi:10.1109/FOCS.2017.37.
- 9 Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. In *Proc. of the 51st Annual ACM-SIGACT Symposium on the Theory of Computing, STOC 2019*, pages 602–613, New York, USA, 2019. ACM. doi:10.1145/3313276.3316300.
- 10 Lorenzo Ciardo, Marcin Kozik, Andrei A. Krokhin, Tamio-Vesa Nakajima, and Stanislav Živný. On the complexity of the approximate hypergraph homomorphism problem, 2023. doi:10.48550/arXiv.2302.03456.
- 11 Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems, 2023. doi:10.48550/arXiv.2301.05084.
- 12 Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009. doi:10.1137/07068062X.
- 13 Irit Dinur, Oded Regev, and Clifford Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica*, 25(5):519–535, 2005. doi:10.1007/s00493-005-0032-4.
- 14 Marek Filakovský, Tamio-Vesa Nakajima, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. Hardness of linearly ordered 4-colouring of 3-colourable 3-uniform hypergraphs, 2023. doi:10.48550/arXiv.2312.12981.
- 15 Venkatesan Guruswami and Sai Sandeep. d -to-1 hardness of coloring 3-colorable graphs with $O(1)$ colors. In *47th International Colloquium on Automata, Languages, and Programming*, volume 168 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 62, 12. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020. doi:10.4230/LIPICs.ICALP.2020.62.

- 16 Magnús M. Halldórsson. A still better performance guarantee for approximate graph coloring. *Inform. Process. Lett.*, 45(1):19–23, 1993. doi:10.1016/0020-0190(93)90246-6.
- 17 Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger, editors, *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, Mar 20–22, 1972, Yorktown Heights, New York*, pages 85–103, Boston, MA, 1972. Springer US. doi:10.1007/978-1-4684-2001-2_9.
- 18 Ken-ichi Kawarabayashi and Mikkel Thorup. Coloring 3-colorable graphs with less than $n^{1/5}$ colors. *J. ACM*, 64(1):4:1–4:23, 2017. doi:10.1145/3001582.
- 19 Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the chromatic number. *Combinatorica*, 20(3):393–415, 2000. doi:10.1007/s004930070013.
- 20 Dmitry Kozlov. *Combinatorial algebraic topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, 2008. doi:10.1007/978-3-540-71962-5.
- 21 Andrei Krokhin and Jakub Opršal. The complexity of 3-colouring H -colourable graphs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS'19)*, pages 1227–1239, Baltimore, MD, USA, 2019. IEEE. doi:10.1109/FOCS.2019.00076.
- 22 Andrei Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, 2022. doi:10.1145/3559736.3559740.
- 23 Andrei A. Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM Journal on Computing*, 52(1):38–79, 2023. doi:10.1137/20M1378223.
- 24 László Lovász. Kneser’s conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, 25(3):319–324, 1978. doi:10.1016/0097-3165(78)90022-5.
- 25 Jiří Matoušek. *Using the Borsuk-Ulam Theorem. Lectures on Topological Methods in Combinatorics and Geometry*. Springer-Verlag Berlin Heidelberg, first edition, 2003. doi:10.1007/978-3-540-76649-0.
- 26 Tamio-Vesa Nakajima and Stanislav Živný. Linearly ordered colourings of hypergraphs. *ACM Trans. Comput. Theory*, 14(3–4), February 2023. doi:10.1145/3570909.
- 27 Tamio-Vesa Nakajima and Stanislav Živný. Boolean symmetric vs. functional PCSP dichotomy. In *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12, 2023. doi:10.1109/LICS56636.2023.10175746.
- 28 Marcin Wrochna. Homomorphism reconfiguration via homotopy. *SIAM J. Discret. Math.*, 34(1):328–350, 2020. doi:10.1137/17M1122578.
- 29 Marcin Wrochna and Stanislav Živný. Improved hardness for H -colourings of G -colourable graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’20)*, pages 1426–1435, Salt Lake City, UT, USA, 2020. SIAM. doi:10.1137/1.9781611975994.86.
- 30 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.
- 31 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.*, 3:103–128, 2007. doi:10.4086/toc.2007.v003a006.

A Graph of reconfigurations of binary polymorphisms




Figure 2 Graph of reconfigurations of $\text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$.