

Hardness of linearly ordered 4-colouring of 3-colourable 3-uniform hypergraphs

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

DOI:

[10.48550/arXiv.2312.12981](https://doi.org/10.48550/arXiv.2312.12981)

Citation for published version (Harvard):

Filakovský, M, Nakajima, T-V, Opršal, J, Tasinato, G & Wagner, U 2023 'Hardness of linearly ordered 4-colouring of 3-colourable 3-uniform hypergraphs' arXiv, pp. 1-39. <https://doi.org/10.48550/arXiv.2312.12981>

[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 ✉ 

Charles University, Prague, Czech Republic

Tamio-Vesa Nakajima ✉ 

University of Oxford, Oxford, UK

Jakub Opršal ✉ 

University of Birmingham, 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 → Approximation algorithms; Mathematics of computing → Algebraic topology; Theory of computation → Problems, reductions and completeness

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

Funding For the purpose of Open Access, the authors have applied a CC-BY public copyright licence to any Author Accepted Manuscript version arising from this submission. All data is provided in full in the results section of this paper.

Marek Filakovský: This research was supported by Charles University (project PRIMUS/21/SCI/014) and the Austrian Science Fund (FWF project P31312-N35).

Tamio-Vesa Nakajima: This research was funded by UKRI EP/X024431/1 and by a Clarendon Fund Scholarship.

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).

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 [15]); 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 [33]). 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 [31], 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 $k \geq c \geq 3$ were obtained by Dinur, Mossel, and Regev [12], Guruswami and Sandeep [14], 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, 31, 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 [32], 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 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 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 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 PCSP($\mathbf{LO}_2, \mathbf{LO}_3$). They conjectured that this problem is NP-hard, and more generally that 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]): PCSP($\mathbf{LO}_2, \mathbf{B}$) is either solved by the *affine integer programming relaxation*, or NP-hard (see Ciardo, Kozik, Krokhn, 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 PCSP($\mathbf{LO}_k, \mathbf{LO}_c$) for $4 \leq k \leq c$ follows relatively easily from NP-hardness of the approximate graph colouring PCSP(K_{k-1}, K_{c-1}), as was observed by Nakajima and Živný and by Austrin (personal communications).²

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

► **Theorem 1.1.** 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

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.

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].

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 $\text{hom}(\mathbf{A}, \mathbf{B})$.

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.

Our focus is on structures with a single ternary relation R , i.e., pairs $(A; R^A)$ with $R^A \subseteq A^3$. Moreover, most structures in this paper have a *symmetric* relation, i.e., the relation R^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*

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

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

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*.

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.

► **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.

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 \mathbf{K} is finite system of sets that is downward closed under inclusion, i.e., $F \subseteq G \in \mathbf{K}$ implies $F \in \mathbf{K}$. The (finite) set $V = \bigcup \mathbf{K}$ is called the set of

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.

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 $f \in \text{hom}(\mathbf{A}, \mathbf{B})$ 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.*

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

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].

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

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 $\omega^3 = 1_A$ (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.

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 $gf \sim 1_X$.

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

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 $(\mathbb{Z}_3\text{-})$ 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, 31, 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 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 $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$.

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



■ **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.

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 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{X}_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 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

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.

required minion isomorphism, we provide a minion homomorphism

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

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 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 $\chi: \text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \mathcal{L}_3$ where \mathcal{L}_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 [30], 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 $\alpha_i = \alpha_k = 1$ 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. \blacktriangleleft

As mentioned before, the above corollary combined with Theorem 2.6 provides the main result of this paper, the NP-completeness of 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 $f(b) = g(b), f(c) = g(c)$.

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.

$c = 1$. 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 $((1, 2), (x, y), (x', y')) = ((1, 2), (1, 3), (x', 1)) \in R^{\mathbf{LO}_3}$, $f(x, y) = 1$, and $f(x', y') = 1$. $f(1, 2) = 2$ is not possible since $((1, 2), (1, 1), (x', y')) = ((1, 1), (1, 2), (x', 1)) \in R^{\mathbf{LO}_3}$, and $f(x', y') = 1, f(1, 1) = 2$. $f(1, 2) = 3$ is not possible since $((1, 2), (2, 2), (x, y)) = ((1, 2), (2, 2), (1, 3)) \in R^{\mathbf{LO}_3}$, 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 $((1, 3), (1, 1), (x', y')) = ((1, 3), (1, 1), (2, 1)) \in R^{\mathbf{LO}_3}$ 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 $((1, 3), (x, y), (2, 2)) = ((1, 3), (1, 2), (2, 2)) \in R^{\mathbf{LO}_3}$ 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 $((1, 2), (1, 3), (3, 1)) \in R^{\mathbf{LO}_3}$, which yields our contradiction.

$c = 2$. 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. \blacktriangleleft

► **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. \blacktriangleleft

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 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. \blacktriangleleft


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.*

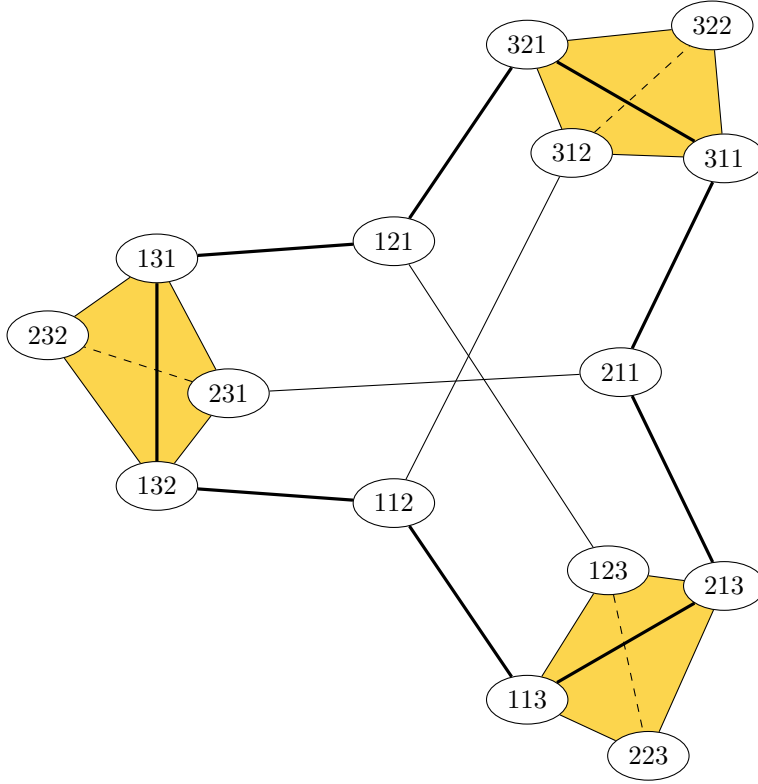
A Structure of $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$

In this section we will investigate the structure of $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$. Luckily the first of these is easy to just draw. We can see the structure of $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ in Figure 3. In this figure, the vertices corresponding to homomorphisms are highlighted, and labelled by the triple of the values of the homomorphism. Each of the tetrahedrons corresponds to a multihomomorphism consisting of all the edges where the unique maximum is 3 in a fixed position $i \in [3]$. The \mathbb{Z}_3 action on $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ is given by rotation.

► Lemma A.1. *There exists a \mathbb{Z}_3 -equivariant map $S^1 \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$.*



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



■ **Figure 3** $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ where only the homomorphisms are explicitly marked.

Proof. Consider the cycle in $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ given by the vertices¹²

$$\begin{aligned} (2, 1, 1) &\rightarrow (2, 1, 3) \rightarrow (1, 1, 3) \rightarrow (1, 1, 2) \rightarrow (1, 3, 2) \\ &\rightarrow (1, 3, 1) \rightarrow (1, 2, 1) \rightarrow (3, 2, 1) \rightarrow (3, 1, 1) \rightarrow (2, 1, 1). \end{aligned}$$

This cycle is highlighted in Figure 3. Observe that it is invariant under the \mathbb{Z}_3 action. Thus we see that there is a \mathbb{Z}_3 -equivariant map $S^1 \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$. ◀

Now, we turn to $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$. Unfortunately, this is significantly more difficult to “just see”, as it has many overlapping simplices of dimensions greater than 1. However, luckily, for our argument it turns out we only need to find a (\mathbb{Z}_3 -equivariant) map from $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ to some sufficiently homotopically simple \mathbb{Z}_3 space P^2 (we will expand on the meaning of “simple” later in Section A.1). This is because the template we are looking at is $(\mathbf{LO}_3, \mathbf{LO}_4)$; thus \mathbf{LO}_4 is “on the right”.

However, constructing a map directly from $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ into P^2 is not entirely straightforward; instead, it is much easier to define an intermediate \mathbb{Z}_3 -space L_4 and show that $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow L_4$ and $L_4 \rightarrow P^2$ equivariantly.

► **Lemma A.2.** *Let L_4 be the order complex of the poset depicted in Fig. 1c, i.e., the poset on $[3] \times \mathbb{Z}_3$ where $(i, \omega^\alpha) < (j, \omega^\beta)$ if $i < j$. Then there is a \mathbb{Z}_3 -map $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow L_4$.*

¹²To be precise, every edge $f_0 \rightarrow f_1$ of the cycle as it is written is not an edge of the complex, but it “goes through” the vertex corresponding to the multihomomorphism $m: i \mapsto f_0(i) \cup f_1(i)$. For ease of readability, these vertices have been suppressed from the notation.

Proof. In order to provide the required equivariant map, we will provide an equivariant simplicial map to the barycentric subdivision of L_4 . Moreover, since both $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ and the barycentric subdivision of L_4 can be expressed as order complexes (barycentric subdivision is the order complex of the face poset of L_4 ; see [25, Section 1.7]), it is enough to provide a monotone equivariant map $\phi: \text{mhom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow \mathcal{F}$, where $\text{mhom}(\mathbf{R}_3, \mathbf{LO}_4)$ is ordered by inclusion and \mathcal{F} is the face poset of L_4 .

The face poset \mathcal{F} of L_4 can be described explicitly as follows: Its elements are those subsets $F \subseteq [3] \times \mathbb{Z}_3$ such that there are no two elements with the same first coordinate, i.e.,

$$\mathcal{F} = \{F \subseteq [3] \times \mathbb{Z}_3 \mid \text{if } (i, \omega^\alpha) \text{ and } (j, \omega^\beta) \text{ are two distinct elements of } F, \text{ then } i \neq j\}$$

ordered by inclusion.

First, we define an auxiliary function $h: \text{hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow [3] \times \mathbb{Z}_3$ by

$$h(f) = (f(j) - 1, \omega^j)$$

where $j \in [3]$ is such that $f(j)$ is the unique maximum of f . Using this function, we define ϕ by

$$\phi(m) = \{h(f) \mid f \in \text{hom}(\mathbf{R}_3, \mathbf{LO}_4), f \leq m\}.$$

It is straightforward to check that ϕ is monotone and equivariant. We will prove that it is well-defined, i.e., that $\phi(m) \in \mathcal{F}$.

Observe that $F \in \mathcal{F}$ if and only if $X \in \mathcal{F}$ for any $X \in \binom{F}{2}$; hence, to show that $\phi(m) \in \mathcal{F}$, it is enough to check that $\{h(f_0), h(f_1)\} \in \mathcal{F}$ for any two distinct homomorphisms $f_0, f_1 \leq m$. Furthermore, $\phi(m) \notin \mathcal{F}$ if and only if f_0 and f_1 share a maximum, but the maximum is attained at a different element. Say $f_0(1) = f_1(2)$ is the unique maximum, then $g \leq m$ defined by $g(1) = f_0(1)$, $g(2) = f_1(2)$, $g(3) = f_0(3)$ is not a homomorphism, which yields the contradiction. \blacktriangleleft

A.1 Constructing Eilenberg-MacLane spaces

The goal of this section is to build explicitly our target space P^2 . However, the bulk of our work will be to show that the standard methods from obstruction theory can be adapted to work in our equivariant setting without fundamental alterations. This task has been already carried out in much bigger generality by Bredon in [7]; the explicit construction in section A.1.3 will be simply a (slightly modified) presentation of the general construction.

It is still useful to build the space P^2 and carry out all the computations explicitly for our context: in this way we can maintain a finer control over the topological structure of the space; control that we will use to construct explicitly the minion isomorphism in Appendix B.

Before exploring the equivariant version of obstruction theory, it might be useful to first invest a bit of time in outlining the basic results in the non-equivariant setting.

A.1.1 Beyond Simplicial Complexes: CW-complexes

While the idea of simplicial complex is a powerful tool that allow for a discretization of topological spaces, for the purpose of homotopy theory they are often too rigid and unwieldy. Instead, the concept of a CW-complex is much more natural and easy to work with. There are multiple equivalent definitions of what a CW-complex is, we will present a constructive definition that is well suited to our use case. For a more complete and thorough presentation, we refer to [16].

The basic step in constructing a CW-complex is the gluing of a disc: suppose we are given a space Y , then we can use a map $\phi: S^n \rightarrow Y$ (called attaching map) to define a new space

$$Z = Y \cup D^{n+1} / \sim$$

where $x \sim y$ if $x \in S^n$ and $y = \phi(x)$. The space Z is usually denoted as $Y \cup_\phi D^{n+1}$. Geometrically, we are identifying the boundary of the disc with the corresponding (via ϕ) point on Y .

► **Definition A.3.** A CW-complex (of finite type) is the union X of spaces $\text{sk}_d X$ defined inductively as follows: We start with $\text{sk}_0 X = \{p_0, \dots, p_k\}$, which is a disjoint union of points. Then, if we are given $\text{sk}_d X$, we can attach a finite number of $(d+1)$ -discs $D_1^{d+1}, \dots, D_\ell^{d+1}$ along some attaching maps $\phi_i: \partial D_i^{d+1} = S^d \rightarrow \text{sk}_d X$ to obtain the new space $\text{sk}_{d+1} X$ that extends $\text{sk}_d X$. By construction, $\text{sk}_0 X \subseteq \text{sk}_1 X \subseteq \dots \subseteq \text{sk}_d X \subseteq \dots$ hence it is possible to define the union

$$X = \bigcup_{d=0}^{\infty} \text{sk}_d X$$

For any $d \geq 0$, the space $\text{sk}_d X$ is d -dimensional and it is called d -skeleton of X . A CW-complex X is n -dimensional if there is $n \geq 0$ such that $\text{sk}_i X = \text{sk}_{i+1} X$ for any $i \geq n$ (i.e., no cells of dimension higher than n were used in “building” X).

- **Example A.4.** ■ S^n is a CW-complex by gluing a n -disc to a point, or by gluing, for all $k \leq n$ two k -discs to an “equatorial” S^{k-1} .
- Every simplicial complex is also a CW-complex: any k -simplex in a simplicial complex L can be seen as glued to its faces via linear maps.
- If X and Y are CW-complexes, then $X \times Y$ is again a CW-complex.

► **Definition A.5.** A map $f: X \rightarrow Y$ between CW-complexes is cellular if for all $d \geq 0$

$$f(\text{sk}_d X) \subseteq \text{sk}_d Y.$$

CW-complexes are the natural way to “discretize” homotopy theory: this is because they are much more flexible than simplicial complexes while still giving an manageable description of the available maps thanks to the Cellular Approximation Theorem [16, Theorem 4.8].

► **Theorem A.6 (Cellular Approximation Theorem).** Let X, Y be CW-complex and $f: X \rightarrow Y$ a continuous map. Then there is a cellular map $\tilde{f}: X \rightarrow Y$ that is homotopic to f .

Furthermore, if there is a subcomplex A in X where f is already cellular then the homotopy can be taken to be stationary on A , i.e., $h(t, x) = f(x)$ for any $x \in A$, $0 \leq t \leq 1$.

This means that, as far as homotopy theory is concerned, every map between CW-complex is a cellular map.

► **Remark A.7.** In a CW-complex, the d -skeleton gives generators and the $d+1$ skeleton gives relations for the d -homotopy group. Formally, let X a connected CW-complex with base point $x_0 \in X$, pick $d \geq 0$ and let $\iota_d: \text{sk}_d X \hookrightarrow X$. Then $(\iota_d)_* : \pi_d(\text{sk}_d X, x_0) \rightarrow \pi_d(X, x_0)$ is surjective while $(\iota_{d+1})_* : \pi_d(\text{sk}_{d+1} X, x_0) \rightarrow \pi_d(X, x_0)$ is an isomorphism.

A.1.2 Classical versus equivariant obstruction theory

One of the most important concept at the core of obstruction theory is the definition of Eilenberg-MacLane spaces, spaces that have exactly one non-zero homotopy group in a prescribed dimension and isomorphism class.

► **Definition A.8.** *Let G be a group and $n \geq 1$ (assume that G is Abelian if $n \geq 2$). Then a space is an Eilenberg-MacLane space $K(G, n)$ with base point $x_0 \in K(G, n)$ if it is connected and*

$$\pi_k(K(G, n), x_0) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

It is a standard homotopic argument to show that there is an Eilenberg-MacLane space for any possible choice of G and n (see, e.g., [16, Section 4.2]). The key property that we are interested in is the following well-known representability result (see, e.g., [16, Theorem 4.57]).

► **Theorem A.9.** *Let X be a CW complex with base point $x_0 \in X$ and G an Abelian group. There is a natural isomorphism¹³ between base point preserving maps up to base point preserving homotopies and cohomology of the space*

$$[X, K(G, n)]_* \cong H^n(X; G).$$

Therefore, it is natural to try and use an equivariant version of Eilenberg-MacLane spaces to obtain a similar isomorphism for equivariant homotopy classes and a suitable version of equivariant cohomology classes.

The main difference in the equivariant setting is that there is no map that naturally plays the role of “0” as the constant map does for the classical setting; as a result, to obtain the desired isomorphism it is necessary to first choose a “reference” map (see [29, Theorem II.3.17]):¹⁴

► **Theorem A.10.** *Let X be a CW-complex with a free \mathbb{Z}_3 action, let G be an Abelian group with an action of \mathbb{Z}_3 and let $K(G, n)$ be an Eilenberg-MacLane space with a free \mathbb{Z}_3 action that induces on $\pi_n(K(G, n)) \cong G$ the same \mathbb{Z}_3 -structure.*

Then, given an equivariant map $f : X \rightarrow K(G, n)$, there is an isomorphism

$$\phi_f : [X, K(G, n)]_{\mathbb{Z}_3} \rightarrow H_{\mathbb{Z}_3}^n(X; G)$$

where $H_{\mathbb{Z}_3}^n(X; G)$ is the Bredon cohomology, defined below, in dimension n (ϕ_f in general depends on the equivariant homotopy class of f).

It is now clear what we want to achieve: we want to construct an equivariant Eilenberg-MacLane space $P^2 = K(G, n)$, for some $n \geq 1$ and Abelian group G , with a free \mathbb{Z}_3 -action such that $L_4 \rightarrow K(G; n)$ equivariantly, and compute the Bredon cohomology module $H_{\mathbb{Z}_3}^n(T^r; G)$ for all $r \geq 1$.

¹³To be precise, $K(-, n)$ is a functor from (Abelian) groups into pointed topological spaces up to base point preserving homotopy equivalences and the stated isomorphism is natural both in X and G .

¹⁴This theorem is shown also in [7] under slightly different topological hypothesis that makes it awkward to use in our setting.

A.1.3 Constructing P^2

We will build the space P^2 by inductively gluing higher dimensional cells on the lower dimensional skeleta.

We start with a 2-dimensional CW-complex Y_2 with three vertices $\{v_0, v_1, v_2\}$, three edges $\{e_0, e_1, e_2\}$ and three 2-discs $\{d_0, d_1, d_2\}$ glued according to the following maps:

$$\phi_i: \partial e_i \rightarrow \text{sk}_0(Y_2)$$

$$\phi_i(\partial e_i) = v_{i+1} - v_i$$

and

$$\psi_i: \partial d_i \rightarrow \text{sk}_1(Y_2)$$

$$\psi_i(\partial d_i) = e_i + e_{i+1} + e_{i+2}.$$

Geometrically, Y_2 is a sphere S^2 with a disc glued along an equator. The \mathbb{Z}_3 -action is also geometrically pretty simple: on the equator ($S^1 = \text{sk}_1 Y_2$) the generator ω acts as a rotation by $2\pi/3$, while it cyclically exchanges the three discs around.

Formally, if $R: D^2 \rightarrow D^2$ is the rotation by $2\pi/3$ around the origin, then the generator ω acts cellularly as (the addition in the indexes is considered modulo 3):

$$\omega \cdot v_i = v_{i+1}$$

$$\omega \cdot e_i = e_{i+1}$$

$$\omega \cdot d_i = R(d_{i+1})$$

► **Lemma A.11.** *The space Y_2 is simply connected and the action on it is free.*

Proof. The action is cellular and no cell is mapped to itself, therefore it is free. Up to homotopy, Y_2 is equivalent to $S^2 \vee_{v_0} S^2$, the space obtained by glueing two spheres at one point (the equivalence contracts the “middle” disc of Y_2 to the vertex v_0).

By the Cellular Approximation Theorem, every map $\gamma: (S^1, p) \rightarrow (Y_2, v_0) \sim (S^2 \vee_{v_0} S^2, v_0)$ can be assumed to map in $\text{sk}_1(S^2 \vee_{v_0} S^2) = \{v_0\}$ and thus nullhomotopic. ◀

► **Lemma A.12.** *There is an equivariant map $L_4 \rightarrow Y_2$.*

Proof. We want to build the map by induction on the skeleton of L_4 , if there were no actions involved, this would be a straightforward obstruction theoretic argument. However, since the action on L_4 is free, it is enough to take care of one representative per orbit and “translate” along the orbit. We will carry out the argument explicitly since it exemplifies well the standard obstruction theoretic arguments.

We start by picking a representative V_i in every orbit of vertices of L_4 and define $f_0(V_i) = v_0$. Then, by equivariance there is a unique way to extend f_0 to the whole vertex set: given $V \in \text{sk}_0 L_4$, there is a unique ω^α and unique V_i such that $\omega^\alpha \cdot V_i = V$, hence define $f_0(V) = v_\alpha$. The map is well defined on $\text{sk}_0 L_4$ and it is also equivariant.

Now, pick a representative E_j for every orbit of edges in L_4 and let V_j, W_j be its endpoints. The vertices $f_0(V_j)$ and $f_0(W_j)$ are already defined but, since Y_2 is connected, there is a path γ_j connecting $f_0(V_j)$ and $f_0(W_j)$; hence we can map the edge E_j along the chosen curve γ_j . Again, by equivariance we have a unique way to extend f_1 to $\text{sk}_1 L_4$: if $\omega^\alpha E_j$ is an edge, its endpoints are $\omega^\alpha \cdot V_j, \omega^\alpha \cdot W_j$ and $\omega^\alpha \cdot \gamma_j$ is a path connecting their images; thus we can map the edge $\omega^\alpha \cdot E_j$ to the path $\omega^\alpha \cdot \gamma_j$.

Finally, pick a representative D_k for every orbit of triangles in L_4 . The map on the boundary $\partial D_k \cong S^1$ is already defined so $f_1|_{\partial D_k}: S^1 \rightarrow Y_2$ is a loop in Y_2 . Since Y_2 is simply

connected, $f_1|_{\partial D_k}$ extends to a map defined on the whole disc $f_2: D_k \cong D^2 \rightarrow Y_2$. Once again, completing the definition of f_2 along the orbits in the same way as before we obtain an equivariant map $f = f_2: L_4 \rightarrow Y_2$, as desired. ◀

We can now build the rest of the space P^2 by inductively glueing $(i + 1)$ -discs to the i -skeleton Y_i obtained at the previous step. First, we fix $Y_3 = Y_2$, i.e., we don't glue any 3-discs. We work out the details in the first step of the construction to exemplify the procedure.

The goal is to glue cells to representatives of generators of $\pi_3(Y_3)$ so that they become nullhomotopic in the new space. By a result of Serre [28], all the homotopy groups of a CW-complex (of finite type) are finitely generated so we can fix representatives $\gamma_0, \dots, \gamma_{d_4}$ for a set of generators $[\gamma_0], \dots, [\gamma_{d_4}]$ of $\pi_3(Y_3)$.

We can glue three 4-dimensional discs $\{D_{0,i}^4, D_{1,i}^4, D_{2,i}^4\}$ per generator γ_i using as attaching maps $\gamma_i, \omega \cdot \gamma_i$ and $\omega^2 \cdot \gamma_i$; the action extends naturally and so we obtain a new \mathbb{Z}_3 -CW-complex

$$Y_4 = \bigcup_{0 \leq i \leq d_4} (Y_3 \cup_{\gamma_i} D_{0,i}^4 \cup_{\omega \cdot \gamma_i} D_{1,i}^4 \cup_{\omega^2 \cdot \gamma_i} D_{2,i}^4)$$

Since we glued disc of dimension 4, $\text{sk}_3 Y_4 = Y_3$; by cellular approximation, $\pi_i(Y_4) \cong \pi_i(Y_3)$ for $i \leq 2$, while $\iota_*: \pi_3(Y_3) \rightarrow \pi_3(Y_4)$ is surjective. For any $[\gamma] \in \pi_3(Y_3)$, $[\gamma] = \sum n_i [\gamma_i]$ hence $\iota_*([\gamma]) = \sum n_i \iota_*([\gamma_i]) = 0$. Therefore, $\pi_3(Y_4) = 0$. What is more, Y_4 is still a finite CW-complex hence Serre theorem still applies and $\pi_\ell(Y_4)$ remains finitely generated for any ℓ . Finally, the inclusion $\iota: Y_3 \hookrightarrow Y_4$ is \mathbb{Z}_3 -equivariant by construction and the action on Y_4 remains consistent with the action on the 3-skeleton.

Analogously, given Y_k we can build the space Y_{k+1} by fixing representatives $\gamma_0, \dots, \gamma_{d_k}$ for a set of generators of $\pi_k(Y_k)$ and gluing three $(k + 1)$ -discs $\{D_{0,i}^{k+1}, D_{1,i}^{k+1}, D_{2,i}^{k+1}\}$ per generator in a similar fashion:

$$Y_{k+1} = \bigcup_{0 \leq i \leq d_k} (Y_k \cup_{\gamma} D^{k+1} \cup_{\omega \cdot \gamma} D^{k+1} \cup_{\omega^2 \cdot \gamma} D^{k+1})$$

By cellular approximation, we have not changed anything in the homotopy groups in dimension less than k and $\iota_*: \pi_k(Y_k) \rightarrow \pi_k(Y_{k+1})$ is surjective. This again implies that $\pi_k(Y_{k+1}) = 0$ and by Serre's Theorem every homotopy group remains finitely generated.

Finally, we define P^2 as the union of Y_k 's.

► **Definition A.13.** *Let Y_k be as above. Using $Y_k \subseteq Y_{k+1}$ for all k , we define P^2 by*

$$P^2 = \bigcup_{k \geq 2} Y_k.$$

By construction, P^2 is a CW-complex with $\text{sk}_{k+1} P^2 = Y_{k+1}$ and consequently, for all $k \geq 1$,

$$\pi_k(P^2) \cong \begin{cases} \pi_2(Y_3) & \text{if } k = 2, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

We have constructed an Eilenberg-MacLane space $K(\pi_2(Y_2), 2)$ but P^2 has some additional properties that we can use (note that $Y_2 = Y_3$, and hence $\pi_2(Y_2) = \pi_2(Y_3)$). In particular, P^2 has a free cellular \mathbb{Z}_3 -action that acts transitively on vertices, edges and 2-cells. Furthermore, P^2 does not have any 3-cell in its CW-decomposition, which we will use later. We also have the following.

► **Lemma A.14.** *There is an equivariant map $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow P^2$.*

Proof. We have $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow L_4 \rightarrow Y_2 \rightarrow P^2$ where the first two maps follow from Lemmata A.2 and A.12, and the last one from the definition of P^2 (Definition A.13). ◀

B

 Computation of $[T^n, P^2]_{\mathbb{Z}_3}$

The goal of this appendix is to describe the minion $\text{hpol}(S^1, P^2)$. In particular, we classify classes of equivariant maps $[T^n, P^2]_{\mathbb{Z}_3}$ by providing the following lemma which is a key step in proving Lemma 3.2.

► **Lemma B.1.** *There is a minion isomorphism*

$$\mu: \mathcal{Z}_3 \rightarrow \text{hpol}(S^1, P^2).$$

Recall that \mathcal{Z}_3 is the minion of affine maps $f: \mathbb{Z}_3^n \rightarrow \mathbb{Z}_3$ such that $f(1, \dots, 1) = 1$. Since each such map is uniquely described by its coefficients we get that there are in bijections with the n -tuples $(\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_3^n$ such that $\sum_i \alpha_i = 1$. In this appendix, we view \mathcal{Z}_3 as an abstract minion whose elements are tuples of the form above, in particular, we have $\mathcal{Z}_3^{(n)} \subseteq \mathbb{Z}_3^n$.

The proof of Lemma B.1, while straightforward, it is rather technical so it might be useful to outline the structure of the argument. We start with describing a family of equivariant maps that are expressed as multivariate monomials over \mathbb{C} . In particular observe that each monomial restricts to a map $S^1 \rightarrow S^1$ (where $S^1 \subset \mathbb{C}$ is the unit circle), and moreover the monomials whose total degree is congruent to 1 modulo 3 induce \mathbb{Z}_3 -equivariant map. These maps are then composed with a fixed embedding $S^1 \rightarrow P^2$ to obtain an equivariant map $T^n = (S^1)^n \rightarrow P^2$ (the choice of a concrete embedding is irrelevant since P^2 is simply connected).

► **Definition B.2.** *Fix an embedding $\iota: S^1 \hookrightarrow P^2$ (e.g., the inclusion of the 1-skeleton), and let $n \geq 1$. We assigns to each tuple $\alpha \in \mathbb{Z}^n$ with $\sum_i \alpha_i = 1 \pmod{3}$, a monomial map $m_\alpha: T^n \rightarrow P^2$, defined by*

$$m_\alpha(z_1, \dots, z_n) = z_1^{\alpha_1} \cdots z_n^{\alpha_n}$$

It is straightforward to check that each m_α is equivariant, indeed

$$\begin{aligned} m_\alpha(\omega \cdot (z_1, \dots, z_n)) &= m_\alpha(\omega z_1, \dots, \omega z_n) = \iota(\prod_j \omega^{\alpha_j} z_j^{\alpha_j}) = \iota(\omega \prod_j z_j^{\alpha_j}) \\ &= \omega \iota(\prod_j z_j^{\alpha_j}) = \omega \cdot m_\alpha(z_1, \dots, z_n) \end{aligned}$$

where the third equality used that $\omega^{3k+1} = \omega$.

► **Remark B.3.** It is not hard to observe that the assignment $\alpha \mapsto m_\alpha$ preserves minors when α is interpreted as a function $f: \mathbb{Z}^n \rightarrow \mathbb{Z}$. While we implicitly use this minion homomorphism, this is not the minion homomorphism we are looking for — importantly, \mathcal{Z}_3 is *not* a subminion of the minion of tuples $\alpha \in \mathbb{Z}^n$ with $\sum \alpha_i = 1 \pmod{3}$ since, e.g., the unary minor of $(2, 2)$ disagrees in the two minions.

We then prove that (in Section B.1), for $\alpha, \beta \in \mathbb{Z}^n$, as long as $\alpha_i \neq \beta_i \pmod{3}$, then m_α and m_β are not equivariantly homotopic. This then allow us to define an injective mapping $\mu: \mathcal{Z}_3 \rightarrow \text{hpol}(S^1, P^2)$ — we view each element $\alpha \in \mathcal{Z}_3^{(n)}$ as a tuple in \mathbb{Z} with $\alpha_i \in \{0, 1, 2\}$.

In Section B.2, we show that any equivariant map $T^n \rightarrow P^2$ is homotopic to some monomial map m_α for $\alpha \in \mathcal{Z}_3^{(n)}$. This is done by showing that $\#[T^n, P^2]_{\mathbb{Z}_3} = 3^{n-1}$ which

we show via obstruction theory — the bijection is given by a bijection between $[T^n, P^2]_{\mathbb{Z}_3}$ and the Bredon cohomology $H_{\mathbb{Z}_3}^2(T^n, G)$ (here use that P^2 is an Eilenberg-MacLane space), and a computation that gives that $H_{\mathbb{Z}_3}^2(T^n, G) \simeq \mathbb{Z}_3^{n-1}$. Consequently, we get that μ is a bijection (or more precisely, that $\mu_n: \mathcal{L}^{(n)} \rightarrow [T^n, P^2]_{\mathbb{Z}_3}$ is a bijection for all n).

Finally, we show that μ is a minion homomorphism, i.e., that it preserves minors. This is done by observing that the *inverse* of μ (which is in fact described in Section B.1) preserves minors. Similarly to homomorphisms of algebraic structures (e.g., groups), an inverse of a bijective minion homomorphism is a homomorphism as well.

B.1 Monomial maps are all homotopically distinct

Here, we prove the following:

► **Lemma B.4.** *Let $\alpha, \beta \in \mathcal{L}_3^{(n)}$. If m_α is equivariantly homotopic to m_β , then $\alpha = \beta$.*

The proof of the lemma is build on the following informal geometric intuition. By cellular approximation, a homotopy between two maps $f_0, f_1: T^n \rightarrow P^2$ “drags” around the image of any of the coordinate circles (i.e., the image of $c_i: S^1 \hookrightarrow T^n$ defined by $c_i(x) = b$ where $b_i = x$ and b_j is constant for each $j \neq i$) of the torus along some 2-disc in P^2 . However since it has to be an equivariant homotopy, it has to drag in the same way the rest of the orbit of the coordinate cycle equivariantly along the other discs in the image. Therefore the number of times a coordinate cycle wraps around can only change by a multiple of 3.

The formalization of this idea, however, requires some work. In particular, we need a way to make the idea of a “degree” of a coordinate cycle precise:

► **Definition B.5.** *Let $e^j \in C^1(P^2)$ and $d^j \in C^2(P^2)$ be the dual of e_j and d_j respectively (i.e., $e^j(e_i) = 0$ if $i \neq j$, and $e^j(e_i) = 1$ if $i = j$). Denote by $x_i \in C_1(T^n)$ the i -th coordinate cycle in T^n and $b_i, B_i \in C_2(T^n)$ fillings for $x_i - \omega \cdot x_i$ and $x_i - \omega^2 \cdot x_i$ respectively.*

Then, the i -degree of an equivariant map $f: T^n \rightarrow P^2$ is

$$\deg_i(f) = (f^*(e^0)(x_i) + f^*(d^0)(b_i) + f^*(d^0)(B_i)) \bmod 3$$

A key observation is that this quantity does not change along equivariant homotopies:

► **Lemma B.6.** *Let $f_0, f_1: T^n \rightarrow P^2$ equivariant maps homotopic via an equivariant homotopy h . Then for all $i \in [n]$, $\deg_i(f_0) = \deg_i(f_1)$.*

Proof. On the cochains of dimension 1 we have:

$$f_0^*(e_0)(x_i) - f_1^*(e_0)(x_i) = (\delta h^*(e^0))(x_i) + (h^*\delta(e^0))(x_i) = 0 + h(d^0 + d^1 + d^2)(x_i)$$

where the second equality is obtained by using the fact that $\partial x_i = 0$ and $\delta e^0 = d^0 + d^1 + d^2$.

On the 2-cochains we have:

$$\begin{aligned} f_0^*(d^0)(b_i) - f_1^*(d^0)(b_i) &= (\delta h^*(d^0))(b_i) + (h^*\delta(d^0))(b_i) = h^*(d^0)(x_i - \omega \cdot x_i) + 0 \\ f_0^*(d^0)(B_i) - f_1^*(d^0)(B_i) &= (\delta h^*(d^0))(B_i) + (h^*\delta(d^0))(B_i) = h^*(d^0)(x_i - \omega^2 \cdot x_i) + 0 \end{aligned}$$

where we use that $\partial b_i = x_i - \omega \cdot x_i$, $\partial B_i = x_i - \omega^2 \cdot x_i$ and $\delta d^0 = 0$.

However, since the homotopy is equivariant, it has to commute with the action and thus $h^*(d^0)(\omega \cdot x_i) = h^*(\omega \cdot d^0)(x_i)$. Summing everything together, using this fact and that $\omega^i \cdot d^0 = d^i$, we obtain that

$$\begin{aligned} \deg_i(f_0) - \deg_i(f_1) &= h(d^0 + d^1 + d^2)(x_i) + h^*(d^0)(x_i - \omega \cdot x_i) + h^*(d^0)(x_i - \omega^2 \cdot x_i) \\ &= h(d^0) \left((x_i + \omega \cdot x_i + \omega^2 \cdot x_i) + (x_i - \omega \cdot x_i) + (x_i - \omega^2 \cdot x_i) \right) \\ &= h^*(d^0)(3x_i) = 3h^*(d^0)(x_i) = 0 \pmod{3}. \end{aligned} \quad \blacktriangleleft$$

Furthermore, we can show the following which will be used to prove that the mapping μ is injective, and also (later) to show that its “inverse” is a minion homomorphism.

► **Lemma B.7.** *Let $\alpha \in \mathbb{Z}^n$ such that $\sum_i \alpha_i = 1 \pmod{3}$. Then*

$$\deg_i(m_\alpha) = \alpha_i \pmod{3}$$

Proof. Since the image of m_α is contained in the 1-skeleton, $m_\alpha^*(d^0) \equiv 0$. Moreover, $(m_\alpha)_*(x_i) = \alpha_i e_0 + \alpha_i e_1 + \alpha_i e_2$, hence $e^0((m_\alpha)_*(x_i)) = \alpha_i$ and $\deg_i(m_\alpha) = \alpha_i \pmod{3}$. ◀

We can now conclude the proof of Lemma B.4 which follows immediately from the above.

Proof of Lemma B.4. If $\alpha \neq \beta$, then they differ in at least one coordinate, i.e., $\alpha_i \neq \beta_i$ for some i . Then $\deg_i(m_\alpha) \neq \deg_i(m_\beta)$ by Lemma B.7, and therefore m_α and m_β are not equivariantly homotopic by Lemma B.6. ◀

Before, we progress further, let us discuss one more consequence of Lemma B.4, namely, the following.

► **Lemma B.8.** *The mapping $\gamma: \text{hpol}(S^1, P^2) \rightarrow \mathcal{Z}_3$ defined by*

$$[f] = (\deg_1(f), \dots, \deg_n(f))$$

satisfies, for each $\alpha \in \mathbb{Z}^n$ with $\sum_i \alpha_i = 1$, and $\pi: [n] \rightarrow [m]$,

$$\gamma([m_\alpha]^\pi) = \gamma([m_\alpha]^\pi),$$

i.e., it preserves minors when restricted to classes of monomial maps.

Proof. Note that γ is well-defined since the degrees do not depend on the choice of representative (Lemma B.6). Further, using Lemma B.7, we have that

$$\gamma([m_\alpha]) = (\alpha_1 \pmod{3}, \dots, \alpha_n \pmod{3})$$

and hence $\gamma([m_\alpha]^\pi) = \beta$ where $\beta_j = (\sum_{i \in \pi^{-1}(j)} \alpha_i) \pmod{3}$. Furthermore, $m_\alpha^\pi = m_{\beta'}$ where $\beta'_j = \sum_{i \in \pi^{-1}(j)} \alpha_i$, and consequently

$$\gamma([m_\alpha]^\pi) = \gamma([m_\alpha^\pi]) = \beta = \gamma(m_\alpha). \quad \blacktriangleleft$$

In the next subsection, we show that μ and γ are in fact inverses to each other, and hence the previous lemma applies for all elements of $\text{hpol}(S^1, P^2)$, i.e., γ (and consequently μ) preserves all minors.

B.2 Monomial maps cover all possible maps up to homotopy

To show that there are no other maps besides the monomial maps m_α , $\alpha \in \mathcal{Z}^3$ up to equivariant homotopy, we will rely on Theorem A.10 in the particular case of $X = T^n$ and P^2 playing the role of $K(\pi_2(P^2), 2)$, which asserts that there is a bijection $[T^n, P^2]_{\mathbb{Z}_3} \cong H_{\mathbb{Z}_3}^2(T^n; \pi_2(P^2))$. The rest of this section will be dedicated to precisely define the objects involved in Theorem A.10 and compute the Bredon cohomology of the torus, i.e., to prove the following proposition.

► **Proposition B.9.** *For all $n, d \geq 1$,*

$$H_{\mathbb{Z}_3}^d(T^n; \pi_2(P^2)) \cong \mathbb{Z}_3^{\binom{n-1}{d-1}}.$$

B.2.1 Equivariant cohomology: a primer

We introduce equivariant cohomology in a specific case that is sufficient for our purposes, and refer to Bredon [7, Section 1], or tom Dieck [29, Section II.3] for a detailed treatment.

As we have previously observed, if a CW-complex X has a free cellular \mathbb{Z}_3 -action, then there is an induced action on the group of d -dimensional integer chains $C_d(X)$ that commutes with the boundary maps. The main insight of Bredon equivariant (co)homology is a shift in perspective: specifying an action on the chains is exactly the same as giving the chain group a structure of Λ -module for a suitable constructed ring (the integer group ring $\mathbb{Z}[\mathbb{Z}_3]$).¹⁵

More explicitly, the ring Λ is defined as the formal integer combinations of elements of $\mathbb{Z}_3 = \{1, \omega, \omega^2\}$ with the multiplication defined by setting $\omega^i \omega^j = \omega^{i+j}$ and extending bilinearly. The key observation is that prescribing a \mathbb{Z}_3 -action on an Abelian group M is exactly the same as choosing a structure of Λ -module on M . In particular, if X is a CW-complex with a cellular \mathbb{Z}_3 -action, for any $d \geq 0$, the corresponding chain group $C_d(X)$ can be seen as a Λ -module with multiplication given on a cell σ by

$$(n_0 + n_1\omega + n_2\omega^2)\sigma = n_0\sigma + n_1(\omega \cdot \sigma) + n_2(\omega^2 \cdot \sigma)$$

and extended linearly.

Since the all the boundary maps commute with the action, these are also maps of Λ -modules hence $C_\bullet(X)$ is a chain complex of Λ -modules. We will denote this chain complex by $C_\bullet^{\mathbb{Z}_3}(X)$. The homology associated to this chain complex is the *equivariant homology* (or Bredon homology) of X , denoted by $H_\bullet^{\mathbb{Z}_3}(X)$.

At the same time, if we fix a Λ -module N , we can consider the equivariant cochain complex:

$$C_{\mathbb{Z}_3}^i(X; N) = \text{Hom}_\Lambda(C_i^{\mathbb{Z}_3}(X), N)$$

with the standard coboundary maps. The cohomology of this cochain complex is the Bredon cohomology, denoted by $H_{\mathbb{Z}_3}^\bullet(X; N)$.

► **Lemma B.10.** *If the action on X is free and cellular, then $C_\bullet^{\mathbb{Z}_3}(X)$ is a chain complex of free Λ -module.*

Proof. For every orbit of d -cells in X choose a representative σ . Then the map

$$\begin{aligned} C_d\left(\frac{X}{\mathbb{Z}_3}; \Lambda\right) &\longrightarrow C_d^{\mathbb{Z}_3}(X) \\ [\sigma] &\longmapsto \sigma \end{aligned}$$

is an isomorphism of Λ -modules; moreover, since the action is cellular, these isomorphisms commute with the boundary maps hence the two chain complexes are isomorphic. ◀

This implies that the functor $\text{Hom}_\Lambda(C_d^{\mathbb{Z}_3}(X), -)$ is exact for all $d \geq 0$. Therefore, if we have a short exact sequence of Λ -modules

$$0 \rightarrow N \rightarrow M \rightarrow Q \rightarrow 0,$$

there is a corresponding short exact sequence of cochain complexes

$$0 \rightarrow C_{\mathbb{Z}_3}^\bullet(X; N) \rightarrow C_{\mathbb{Z}_3}^\bullet(X; M) \rightarrow C_{\mathbb{Z}_3}^\bullet(X; Q) \rightarrow 0,$$

¹⁵This holds in our specific context. In the general case, the Bredon cohomology has a richer structure than simply that of Λ -module.

and thus a long exact sequence in cohomology

$$\dots \rightarrow H_{\mathbb{Z}_3}^{i-1}(X; Q) \rightarrow H_{\mathbb{Z}_3}^i(X; N) \rightarrow H_{\mathbb{Z}_3}^i(X; M) \rightarrow H_{\mathbb{Z}_3}^i(X; Q) \rightarrow \dots$$

The key technical lemma we are going to use to compute the equivariant cohomology of the torus T^n follows after we fix some notation.

Assume that X is a CW complex with a free \mathbb{Z}_3 -action. Let $p: X \rightarrow X/\mathbb{Z}_3$ be the projection map. Let $I = \Lambda(1 + \omega + \omega^2)$ be the ideal in Λ generated by the element $1 + \omega + \omega^2$, and let $\iota: I \hookrightarrow \Lambda$ denote the inclusion. Let Z be the Λ -module defined as $Z = \mathbb{Z}$ with the trivial multiplication by ω and let Z_{cyc} the Λ -module $Z_{\text{cyc}} = \mathbb{Z} \oplus \mathbb{Z} \oplus \mathbb{Z}$ where the multiplication by ω is a cyclic shift, i.e., $\omega \cdot (n_0, n_1, n_2) = (n_2, n_0, n_1)$. Furthermore, we have the following homomorphisms:

1. an inclusion of Λ -modules $d: Z \hookrightarrow Z_{\text{cyc}}$ defined as $d(n) = (n, n, n)$;
2. an isomorphism of Λ -modules $\phi_1: Z \rightarrow I$ defined as $\phi_1(n) = n + n\omega + n\omega^2$;
3. an isomorphism of Λ -modules $\phi_2: Z_{\text{cyc}} \rightarrow \Lambda$ defined by $\phi_2(n_0, n_1, n_2) = n_0 + n_1\omega + n_2\omega^2$;
4. an isomorphism of Abelian groups $h_1: C^\bullet(X/\mathbb{Z}_3) \rightarrow C_{\mathbb{Z}_3}^\bullet(X; Z)$ defined by $h_1(\alpha) = (\sigma \mapsto \alpha(p(\sigma)))$; and finally
5. an isomorphism of Abelian groups $h_2: C^\bullet(X) \rightarrow C_{\mathbb{Z}_3}^\bullet(X; Z_{\text{cyc}})$ defined by

$$h_2(\alpha) = (\sigma \mapsto (\alpha(\sigma), \alpha(\omega \cdot \sigma), \alpha(\omega^2 \cdot \sigma))).$$

All of these claims are straightforward to check.

► **Lemma B.11.** *Let X be a CW complex with a free \mathbb{Z}_3 -action, then the following diagram commutes*

$$\begin{array}{ccccc} C^\bullet(X/\mathbb{Z}_3) & \xrightarrow{h_1} & C_{\mathbb{Z}_3}^\bullet(X; Z) & \xrightarrow{-\circ\phi_1} & C_{\mathbb{Z}_3}^\bullet(X; I) \\ \downarrow p^* & & \downarrow -\circ d & & \downarrow -\circ \iota \\ C^\bullet(X) & \xrightarrow{h_2} & C_{\mathbb{Z}_3}^\bullet(X; Z_{\text{cyc}}) & \xrightarrow{-\circ\phi_2} & C_{\mathbb{Z}_3}^\bullet(X; \Lambda) \end{array}$$

Proof. The proof is a straightforward unravelling of definitions. ◀

B.2.2 Computing the cohomology of the torus

We are now going to specify the previous construction to our case, that is $X = T^n$ with the diagonal action (that is cellular for the product CW-structure on the torus) and coefficients $\pi_2(P^2)$ (that is a Λ -module with the induced action).

The first ingredient is to explicitly study the structure of Λ -module on $\pi_2(P^2)$.

► **Lemma B.12.** *The second homotopy group $\pi_2(P^2)$ is isomorphic as a Λ -module to $M = \mathbb{Z} \oplus \mathbb{Z}$ where the multiplication by ω is given by the matrix*

$$A = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$$

Proof. Since P^2 is simply connected, by Hurewicz theorem [16, see Section 4.2], there is an isomorphism between $\pi_2(P^2)$ and $H_2(P^2)$ that is an isomorphism of Λ -modules.

By direct computation, it is easy to show that the chains $\beta_i = d_i - d_{i+1} \in C_2(P^2)$ are cycles and any two out of $\{[\beta_i]\}_{0 \leq i \leq 2}$ form a basis for $H_2(P^2) \cong \mathbb{Z} \oplus \mathbb{Z}$. Fix $[\beta_0]$ and $[\beta_1]$, then the multiplication by ω is

$$\begin{aligned} \omega \cdot [\beta_0] &= [\omega \cdot \beta_0] = [\beta_1] \\ \omega \cdot [\beta_1] &= [\omega \cdot \beta_1] = [\beta_2] = -[\beta_0] - [\beta_1]. \end{aligned}$$

as claimed. ◀

The module M has some very useful properties that will allow us to use Lemma B.11. In particular, we have the following lemma.

► **Lemma B.13.** *The module M is generated by a single element and $\text{Ann}(M) = I$ the ideal generated by $1 + \omega + \omega^2$.*

Proof. M is generated over Λ by $m = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. As a consequence, $\lambda \in \text{Ann}(M)$ if and only if $\lambda m = 0$; hence, if $\lambda = n_0 + n_1\omega + n_2\omega^2$, then

$$\lambda m = n_0 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + n_1 \begin{pmatrix} 0 \\ 1 \end{pmatrix} + n_2 \begin{pmatrix} -1 \\ -1 \end{pmatrix} = \begin{pmatrix} n_0 - n_2 \\ n_1 - n_2 \end{pmatrix}.$$

Therefore, $\lambda m = 0$ if and only if $n_0 = n_1 = n_2$ if and only if $\lambda \in I$. ◀

The last ingredient needed is to compute what the projection map p^* does on the level of cohomology. While it is possible to compute p^* directly, it is easier to change a little the action on the torus to simplify the calculations:

► **Lemma B.14.** *Let X be the torus $T^n \subseteq \mathbb{C}^n$ with the diagonal \mathbb{Z}_3 -action given by the multiplication with $\omega = e^{2\pi i/3}$ (i.e., $\omega \cdot (z_1, \dots, z_n) = (\omega z_1, \dots, \omega z_n)$). Let Y be the same torus but with \mathbb{Z}_3 acting only on the first coordinate (i.e., $\omega \cdot (z_1, \dots, z_n) = (\omega z_1, z_2, \dots, z_n)$). Then there is a \mathbb{Z}_3 -equivariant homeomorphism $X \rightarrow Y$.*

Proof. The maps $h: X \rightarrow Y$ and $h': Y \rightarrow X$ defined by

$$\begin{aligned} h: (z_1, \dots, z_n) &\mapsto (z_1, z_1^{-1}z_2, \dots, z_1^{-1}z_n) \\ h': (z_1, \dots, z_n) &\mapsto (z_1, z_1z_2, \dots, z_1z_n) \end{aligned}$$

are clearly continuous and mutually inverse. We will show that h preserve the actions involved, and hence that h is an equivariant homeomorphism:

$$h(\omega \cdot_X (z_1, \dots, z_n)) = h(\omega z_1, \dots, \omega z_n) = (\omega z_1, z_1^{-1}\omega z_2, \dots, z_1^{-1}\omega z_n) = \omega \cdot_Y h(z_1, \dots, z_n) \quad \blacktriangleleft$$

► **Remark B.15.** If we see a torus T^n as a quotient of \mathbb{R}^n over the standard lattice, Lemma B.14 shows that factoring out the action is the same as factoring out the lattice generated by $\{\frac{1}{3}e_1, e_2, \dots, e_n\}$. This means that, topologically, the quotient is still a torus.

For all future homological calculations, we can assume that \mathbb{Z}_3 acts by multiplying by ω on the first coordinate. Using this simplified action on the torus it is much easier to compute the quotient map $p^*: H^\bullet(T^n/\mathbb{Z}_3) \rightarrow H^\bullet(T^n)$. To achieve this objective, it is necessary to fix a basis for the cohomology of the torus. The ideal choice would be a basis that is “easy” to evaluate on homology classes in order to compute easily the image of p^* . In the case of the torus, a direct application of the universal coefficient theorem ([16, see Section 3.1]) show that homology and cohomology in dimension 1 are dual to each other; hence we can choose as basis for the first cohomology group the dual of a clever basis for the first homology group. In particular, let $\{x_i\}$ be the basis for $H_1(T^n)$ corresponding to the coordinate loops in T^n and denote by $\{x^i\}$ the dual basis in $H^1(T^n) \cong \text{Hom}(H_1(T^n), \mathbb{Z})$; define analogously $\{q_i\}$ basis for $H_1(T^n/\mathbb{Z}_3)$ and $\{q^i\}$ for $H^1(T^n/\mathbb{Z}_3)$. Then

$$p_1^*(q^i) = \begin{cases} 3x^1 & \text{if } i = 1 \\ x^i & \text{otherwise.} \end{cases}$$

The ring structure on cohomology (see [16, Section 3.2]) of the torus allows us to build a convenient basis for all the other cohomology groups out of $\{x^i\}$. In fact, elements of the form $x^I = x^{i_1} \smile \dots \smile x^{i_d}$, where $I = (i_1, \dots, i_d)$ and $i_1 < \dots < i_d$, form a basis for $H^d(T^n)$. Let q^I denote the analogous basis for $H^d(T^n/\mathbb{Z}_3)$. Since p^* is a ring map, it commutes with the cup product, hence it can be explicitly computed on such a basis. We have that, for all d ,

$$p_d^*(q^I) = p_1^*(q^{i_1}) \smile \dots \smile p_1^*(q^{i_d}) = \begin{cases} 3x^I & \text{if } i_1 = 1 \\ x^I & \text{else.} \end{cases}$$

In particular, p_d^* is injective for all $d \geq 0$ and, in this choice of basis, p^* is the diagonal matrix with $\binom{n-1}{d-1}$ 3's and $\binom{n-1}{d}$ 1's on the diagonal.

We are finally ready to compute the equivariant cohomology group of the torus T^n and prove Proposition B.9 which claims that $H_{\mathbb{Z}_3}^d(T^n; \pi_2(P^2)) \cong \mathbb{Z}_3^{\binom{n-1}{d-1}}$.

Proof of Proposition B.9. Fix $n \geq 2$. By Lemma B.13, we have a short exact sequence

$$0 \rightarrow I \rightarrow \Lambda \rightarrow M \rightarrow 0$$

which induces short exact sequence of cochain complexes

$$0 \rightarrow C_{\mathbb{Z}_3}^\bullet(T^n; I) \rightarrow C_{\mathbb{Z}_3}^\bullet(T^n; \Lambda) \rightarrow C_{\mathbb{Z}_3}^\bullet(T^n; M) \rightarrow 0$$

Using Lemma B.11, we get that the following short sequence is also exact

$$0 \rightarrow C^\bullet(T^n/\mathbb{Z}_3) \xrightarrow{p^*} C^\bullet(T^n) \rightarrow C_{\mathbb{Z}_3}^\bullet(T^n; M) \rightarrow 0$$

This short exact sequence induces the following long exact sequence in cohomology

$$\dots \rightarrow H^d(T^n/\mathbb{Z}_3) \xrightarrow{p_d^*} H^d(T^n) \rightarrow H_{\mathbb{Z}_3}^d(T^n; M) \rightarrow H^{d+1}(T^n/\mathbb{Z}_3) \xrightarrow{p_{d+1}^*} H^{d+1}(T^n) \rightarrow \dots$$

Since p_d^* is injective for any $d \geq 1$, by exactness we have that $H_{\mathbb{Z}_3}^d(T^n; M) \cong \text{coker } p_d^*$. Furthermore, $\text{coker } p_d^* = \mathbb{Z}^{\binom{n}{d}} / \text{im } p_d^* \simeq \mathbb{Z}_3^{\binom{n-1}{d-1}}$ which yields the desired. \blacktriangleleft

B.3 The minion isomorphism

Putting everything together, we can now provide the required isomorphism of minions \mathcal{Z}_3 and $\text{hpol}(S^1, P^2)$. Recall the definition of maps $\mu: \mathcal{Z}_3 \rightarrow \text{hpol}(S^1, P^2)$ which maps α to the homotopy class of the monomial map m_α , and $\gamma: \text{hpol}(S^1, P^2) \rightarrow \mathcal{Z}_3$ which maps a homotopy class of a map f to the sequence of its degrees.

Proof of Lemma B.1. We show that μ and γ are mutually inverse minion homomorphisms. Recall that μ is injective by Lemma B.4. We first show that it is surjective. Indeed, we have that, for each $n \geq 1$,

$$\#[T^n, P^2]_{\mathbb{Z}_3} = 3^{n-1}$$

by a combination of Proposition B.9 and Theorem A.10 since $H^2(T^n, \pi_2(P^2)) \simeq \mathbb{Z}_3^{n-1}$, and the latter has 3^{n-1} elements.

The above, in particular, means that $\mu_n: \mathcal{Z}_3^{(n)} \rightarrow \text{hpol}^{(n)}(S^1, P^2)$ is onto, and hence a bijection. Consequently, γ is a minion homomorphism by Lemma B.8 since every class in

$[T^n, P^2]_{\mathbb{Z}_3}$ contains a monomial map. Observe that $\mu \circ \gamma$ is the identity map by Lemma B.7, which implies that γ is the inverse of the bijection μ .

Finally, an inverse of a bijective minion homomorphism γ is a minion homomorphism since, for all $\alpha \in \mathcal{L}_3$,

$$\mu(\alpha^\pi) = \mu(\gamma(\mu(\alpha)^\pi)) = \mu(\gamma(\mu(\alpha)^\pi)) = \mu(\alpha)^\pi.$$

This concludes that μ and γ are the required minion isomorphisms. \blacktriangleleft

C A minion homomorphism

The goal of this appendix is to provide minion homomorphisms

$$\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \text{hpol}(\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3), \text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)) \rightarrow \text{hpol}(S^1, P^2)$$

using the results of previous appendices. In particular, the second of the above homomorphism uses the equivariant maps $S^1 \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\text{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \rightarrow P^2$. Let us note that similar minion homomorphisms have been constructed in [23, Lemma 3.22], and the proof can be altered to construct the homomorphisms, we require here. Rather than doing that, we provide two general lemmata that cover both of these uses as well as any possible future uses. We formulate and prove these two lemmata in the language of category theory which is well suited. We do not assume deep knowledge of category theory and recall the necessary core notions.

Categories, functors, and natural transformations

Instead of giving precise definitions, which can be found in any category theory textbook, we will give examples of the basic notions of category theory.

Let us recall that a *category* \mathbf{cat} is a class of *objects* (e.g., relational structures of a given signature, topological spaces, groups, sets, etc.) together with a set $\text{hom}(A, B)$ of *morphisms* for each pair of objects $A, B \in \mathbf{cat}$ (e.g., homomorphisms from A to B , etc.). Morphisms can be composed in a natural way, and for each object A there is a (two-sided) identity morphism $1_A \in \text{hom}(A, A)$. Also recall that an *isomorphism* is an invertible morphism, i.e., $f \in \text{hom}(A, B)$ is an isomorphism if there is $g \in \text{hom}(B, A)$ such that $f \circ g = 1_A$ and $g \circ f = 1_B$.

► **Example C.1.** The most basic category is the category of sets: objects are sets, and morphisms are maps between the two given sets. We denote this category by \mathbf{set} . Relational structures of a fixed signature Σ together with homomorphisms form a category which we denote by \mathbf{rel}_Σ or usually just \mathbf{rel} when the signature is clear from the context. The categorical notion of isomorphism coincides with structural isomorphism.

► **Example C.2.** The usual structure of a category on topological spaces has continuous maps as morphisms, and homeomorphisms as isomorphisms. Nevertheless, we will work with a different structure of a category on topological spaces. Namely, the category \mathbf{htop}_G whose objects are topological spaces with an action of a fixed group G , and morphisms are *equivariant homotopy classes of maps*, i.e., $\text{hom}(X, Y) = [X, Y]_G$. The composition is defined by $[f] \circ [g] = [f \circ g]$.¹⁶ We will write just \mathbf{htop} instead of $\mathbf{htop}_{\mathbb{Z}_3}$.

¹⁶ Do not confuse this category with the *homotopy category* which identifies *weakly* homotopy equivalent spaces.

Isomorphisms in this category coincide with equivariant homotopy equivalences: If $f: X \rightarrow Y$ and $g: Y \rightarrow X$ witness a homotopy equivalence, then both $[f]$ and $[g]$ are isomorphisms in \mathbf{htop}_G , since we have $f \circ g \sim 1_X$ and $g \circ f \sim 1_Y$, and consequently $[f] \circ [g] = 1_X$ and $[g] \circ [f] = 1_Y$. Loosely speaking, this category is obtained from the category of topological spaces with a G -action by identifying equivariantly homotopic spaces.

► **Example C.3.** Another useful category is the category whose objects are non-negative integers, and morphisms from n to m are maps $[n] \rightarrow [m]$, i.e., $\text{hom}(n, m) = \{f: [n] \rightarrow [m]\}$. We denote this subcategory of the category of sets by \mathbf{nat} . Note that \mathbf{nat} contains all important information about finite sets since it has an object for each isomorphism class (cardinality) of finite sets.

A *functor* is a natural notion of a morphism between categories. A functor $\mathcal{F}: \mathbf{cat}_1 \rightarrow \mathbf{cat}_2$ is a mapping between objects and morphisms that preserves composition and identity morphisms. More specifically, we require that $\mathcal{F}(A) \in \mathbf{cat}_2$ for each $A \in \mathbf{cat}_1$, and $\mathcal{F}(f) \in \text{hom}(\mathcal{F}(A), \mathcal{F}(B))$ for each $f \in \text{hom}(A, B)$ where $A, B \in \mathbf{cat}_1$.

► **Example C.4.** We will view the assignment $\mathcal{B}: \mathbf{A} \mapsto \text{Hom}(\mathbf{R}_3, \mathbf{A})$ as a functor $\mathbf{rel} \rightarrow \mathbf{htop}$. This functor is defined on morphism by mapping a homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ to the class $[f_*]$ where $f_*: \text{Hom}(\mathbf{R}_3, \mathbf{A}) \rightarrow \text{Hom}(\mathbf{R}_3, \mathbf{B})$ is the induced continuous map.

► **Example C.5.** An abstract minion is a functor $\mathcal{M}: \mathbf{nat} \rightarrow \mathbf{set}$ such that $\mathcal{M}(n) \neq \emptyset$ for all $n \neq 0$.

Finally, *natural transformations* are morphisms between functors.

► **Example C.6.** A natural transformation between two minions coincides with the notion of minion homomorphism: a natural transformation $\eta: \mathcal{M} \rightarrow \mathcal{N}$ is a collection of maps η_n where n ranges through objects of \mathbf{nat} such that, for each $n, m \in \mathbf{nat}$ and $\pi \in \text{hom}(n, m)$, the following square commutes.

$$\begin{array}{ccc} \mathcal{M}(n) & \xrightarrow{\pi^{\mathcal{M}}} & \mathcal{M}(m) \\ \eta_n \downarrow & & \downarrow \eta_m \\ \mathcal{N}(n) & \xrightarrow{\pi^{\mathcal{N}}} & \mathcal{N}(m) \end{array}$$

It is easy to observe that this square commutes if and only if η preserves minors.

In general, we will say that a map η_A that depends on an object A of some category is *natural in A* if some square, obtained by varying A by a morphism $a: A \rightarrow A'$, commutes. Which square commutes is usually clear from the context.

Products and polymorphisms

A polymorphism is a homomorphism from a power, and hence if we want to define polymorphisms in a category, we need to have a notion of a power, or more generally *product*. Products in an arbitrary category are defined by how they interact with other objects. More precisely, we can observe that products of sets, structures, or topological spaces there is a bijection

$$\text{hom}(A, B_1) \times \cdots \times \text{hom}(A, B_n) \simeq \text{hom}(A, B_1 \times \cdots \times B_n)$$

for each object A , i.e., in order to define a morphism from an object A to a product of objects B_1, \dots, B_n , it is enough to specify an n -tuple of morphisms from A to each B_i . (In order to

define products precisely, we would also require that the above bijection is natural in A .) Formally, the products are defined as follows.

► **Definition C.7.** Let $B_1, \dots, B_n \in \mathbf{cat}$. We say that an object P together with morphisms $p_i \in \text{hom}(P, B_i)$ for $i \in [n]$ is the product of B_1, \dots, B_n if, for each $A \in \mathbf{cat}$ and morphisms $a_i \in \text{hom}(A, B_i)$, there is a unique morphism $a \in \text{hom}(A, P)$ such that $a_i = p_i \circ a$. The maps p_i are called projections.

Let us note that the projections can be obtained from the above bijection by substituting P for A and considering the identity map on P which is assigned a tuple of morphisms $(p_1, \dots, p_n) \in \text{hom}(P, B_1) \times \dots \times \text{hom}(P, B_n)$. Observe that products are well-defined only up to isomorphisms. Indeed, if (P, p_i) and (P', p'_i) are both products of B_1, \dots, B_n , then there are morphisms $f \in \text{hom}(P, P')$ and $g \in \text{hom}(P', P)$ such that $p'_i = p_i \circ f$ and $p_i = p'_i \circ g$. Since 1_P is the unique morphisms such that $p_i = p_i \circ 1_P$, we get that $g \circ f = 1_P$. Analogously, $f \circ g = 1_{P'}$.

► **Example C.8.** Products do not always exist. Nevertheless, all the categories we work with in this paper (**nat**, **set**, **rel**, and **htop_G**) have all (finite) products which coincide with the usual definitions.

Let us comment on products in **htop_G** which are a bit subtle. The product of G -spaces X_1, \dots, X_n is the space $X_1 \times \dots \times X_n$ with the usual product topology and the coordinatewise (diagonal) action of G . The projections are then the homotopy classes of usual projections. Let us show the universal property of the product. Let Y be a G -space, and let $f_i: Y \rightarrow X_i$ be continuous maps. We claim that there is a map $f: Y \rightarrow X_1 \times \dots \times X_n$ such that $p_i \circ f$ and f_i are equivariantly homotopic, and this map is unique up to equivariant homotopy. The existence is straightforward since we can define f in the usual way:

$$f(y_1, \dots, y_n) = (f_1(y_1), \dots, f_n(y_n))$$

It is easy to check that this map is indeed G -equivariant since the f_i 's are, moreover, we have $p_i \circ f = f_i$. To show that f is unique up to equivariant homotopy, let g be such that $p_i \circ g$ and f_i are homotopic for all i . This means we have homotopies $h_i: [0, 1] \times Y \rightarrow X_i$ between f_i and $p_i \circ g$ for each i . By the universal property of product in topological spaces, these define a homotopy $h: [0, 1] \times Y \rightarrow X_1 \times \dots \times X_n$, i.e., we define h by the same prescription as f using h_i instead of f_i . This h is the required homotopy between f and g .

We can now define the notion of a polymorphism in any category with products.

► **Definition C.9.** Let \mathbf{cat} be a category with finite products, $A, B \in \mathbf{cat}$ be two objects, and $n \geq 0$. We define a polymorphism from A to B of arity n to be any element $f \in \text{hom}(A^n, B)$ where A^n is the n -fold power of A .

In order to define the *polymorphism minion* $\text{pol}(A, B)$ in such a category \mathbf{cat} , we need a functor from **nat** to **set** that assigns to each n , the set $\text{hom}(A^n, B)$. An easy way to observe that such a functor can be always defined is to decompose it as two contravariant (i.e., arrow-reversing) functors **nat** \rightarrow **cat** and **cat** \rightarrow **set**: the first of which assigns to n the n -fold power A^n of A , and the second of which assigns to this A^n the set $\text{hom}(A^n, B)$.

Let us first observe that the assignment $n \mapsto A^n$ can be extended to a functor $A^-: \mathbf{nat} \rightarrow \mathbf{cat}$ for each $A \in \mathbf{cat}$. This can be relatively easily observed in all concrete cases, e.g., if A is a structure \mathbf{A} , then, for each $\pi: [n] \rightarrow [m]$, the mapping $\mathbf{A}^\pi: A^m \rightarrow A^n$ is defined by $a \mapsto a \circ \pi$ which is clearly a homomorphism. To give a general definition, we use the universal property of products. Let $\pi: [n] \rightarrow [m]$ be a mapping. We want to define a homomorphism

$p_\pi^A: A^m \rightarrow A^n$. Using the definition of the n th power, it is enough to give an n -tuple of homomorphisms $a_i: A^m \rightarrow A$, $i \in [n]$. We let $a_i = p_{\pi(i)}$ where p_j denote projections of the m th power of A . Finally, it is straightforward to check that the assignment $\mathcal{A}(\pi) = p_\pi^A$ preserves compositions and identities.

The second of these assignments is known as *contravariant hom-functor* $\text{hom}(-, B)$. It maps an object $A \in \mathbf{cat}$ to the set $\text{hom}(A, B)$, and a morphism $f \in \text{hom}(A, A')$ to the mapping $- \circ f: \text{hom}(A', B) \rightarrow \text{hom}(A, B)$ defined by $g \mapsto g \circ f$.

► **Remark C.10.** If $\mathbf{cat} = \mathbf{set}$, then the functor A^- can be also described as a restriction of the contravariant hom-functor $\text{hom}(-, A)$ to \mathbf{nat} viewed as a subcategory of \mathbf{set} .

► **Definition C.11.** Let \mathbf{cat} be a category with finite products, and $A, B \in \mathbf{cat}$ be such that $\text{hom}(A, B)$ is non-empty. We define the polymorphism minion $\text{pol}_{\mathbf{cat}}(A, B)$ as the composition of functors A^- and $\text{hom}(-, B)$. We will omit the index \mathbf{cat} whenever the category is clear from the context.

► **Example C.12.** The polymorphism minion in the category \mathbf{htop} coincides with the minion of homotopy classes of polymorphisms, i.e., for all spaces $X, Y \in \mathbf{htop}$, we have $\text{pol}_{\mathbf{htop}}(X, Y) = \text{hpol}(X, Y)$.

Note that the definition of this polymorphism minion does not depend (up to natural equivalence) on which of the realisation of the power functor A^- we take. Finally, we recall the categorical definition of preserving products, which ensures in particular that $\mathcal{F}(A)^-$ is naturally equivalent to the composition of A^- and \mathcal{F} .

► **Definition C.13.** We say that a functor $\mathcal{F}: \mathbf{cat}_1 \rightarrow \mathbf{cat}_2$ preserves finite products if, for each $A_1, \dots, A_n \in \mathbf{cat}_1$ and their product (P, p_i) , $(\mathcal{F}(P), \mathcal{F}(p_i))$ is a product of $\mathcal{F}(A_1), \dots, \mathcal{F}(A_n)$.

► **Lemma C.14.** If a functor \mathcal{F} preserves products, then $\mathcal{F} \circ A^-$ and $\mathcal{F}(A)^-$ are naturally equivalent.

Proof. We define a natural equivalence $\eta: \mathcal{F} \circ A^- \rightarrow \mathcal{F}(A)^-$ by components as $\eta_n: \mathcal{F}(A^n) \rightarrow \mathcal{F}(A)^n$ is defined as the map given by the n -tuple $\mathcal{F}(p_1), \dots, \mathcal{F}(p_n): \mathcal{F}(A^n) \rightarrow \mathcal{F}(A)$. Since η_n commutes with the projections of the two products, it is an isomorphism by the same argument as we used in showing that product is unique up to isomorphisms. It is also straightforward to check that η is natural using the universal property of the product. ◀

C.1 Two general lemmata

Given the definitions above, the first lemma is rather trivial. It claims that we can transfer polymorphisms through any functor that preserves products. It was mentioned in the context of promise CSPs in [31], although in different flavours it can be tracked down to Lawvere theories.

► **Lemma C.15.** If a functor $\mathcal{F}: \mathbf{cat}_1 \rightarrow \mathbf{cat}_2$ preserves products, then there is a minion homomorphism

$$\xi: \text{pol}(A, B) \rightarrow \text{pol}(\mathcal{F}(A), \mathcal{F}(B))$$

for all $A, B \in \mathbf{cat}_1$ such that $\text{hom}(A, B) \neq \emptyset$.

Proof. By definition $\text{pol}(\mathcal{F}(A), \mathcal{F}(B))$ is the composition of $\mathcal{F}(A)^-$ and $\text{hom}(-, \mathcal{F}(B))$. Since the first functor is naturally equivalent to $\mathcal{F} \circ A^-$, we may assume they are equal. We can then define a natural transformation $\xi: \text{hom}(A^-, B) \rightarrow \text{hom}(\mathcal{F}(A^-), \mathcal{F}(B))$ by

$$\xi_n(f) = \mathcal{F}(f).$$

To show ξ preserves minors, observe that f^π is defined as $f \circ p_\pi^A$, and consequently $\mathcal{F}(f^\pi) = \mathcal{F}(f \circ p_\pi^A) = \mathcal{F}(f) \circ \mathcal{F}(p_\pi^A) = \mathcal{F}(f) \circ p_\pi^{\mathcal{F}(A)} = \mathcal{F}(f)^\pi$. \blacktriangleleft

The second lemma is a direct generalisation of [3, Lemma 4.8(1)], and the proof is analogous.

► **Lemma C.16.** *Let $A, A', B, B' \in \mathbf{cat}$ be such that $\text{hom}(A, B)$ is non-empty. Then every pair $a \in \text{hom}(A', A)$ and $b \in \text{hom}(B, B')$ induces a minion homomorphism*

$$\xi: \text{pol}(A, B) \rightarrow \text{pol}(A', B').$$

Proof. First, observe that a induces a natural transformation $A^- \rightarrow (A')^-$, which can be shown by using the definition of products. We will denote its components by $a^n \in \text{hom}(A^n, (A')^n)$. The minion homomorphism is defined by $\xi_n(f) = b \circ f \circ a^n$. To show that ξ preserves minors, observe that the following diagram commutes.

$$\begin{array}{ccccc} (A')^n & \xrightarrow{a^n} & A^n & & \\ p_\pi^{A'} \downarrow & & p_\pi^A \downarrow & \searrow^{f^\pi} & \\ (A')^m & \xrightarrow{a^m} & A^m & \xrightarrow{f} & B \xrightarrow{b} B' \end{array}$$

Consequently, we have that $\xi(f^\pi) = b \circ f^\pi \circ a^n = (b \circ f \circ a^m) \circ p_\pi^{A'} = \xi_m(f)^\pi$. \blacktriangleleft

C.2 Homomorphism complexes preserve products

We prove that the functor $\text{Hom}(\mathbf{R}_3, -)$ preserves products in the categorical sense. Essentially, it follows by the same argument as the well-known fact that the product of homomorphism complexes is homotopically equivalent to the homomorphism complex of the product (see, e.g., [20, Section 18.4.2]). We provide a bit more detailed proof to show that the homotopy equivalence can be taken equivariant, and that the products are preserved in the categorical sense which is a slightly stronger statement. Let us first prove that there is an equivariant homotopy equivalence. To provide this equivalence, we use the following technical lemma.

► **Lemma C.17.** *Let \mathbf{K} be an order complex of a poset. If $\alpha, \beta: V(\mathbf{K}) \rightarrow V(\mathbf{K})$ are monotone maps such that $\alpha(x) \leq \beta(x)$ for each x , then their linear extensions are homotopic.*

Proof. Observe that the map $h: \{0, 1\} \times V(\mathbf{K}) \rightarrow V(\mathbf{K})$ defined by $h(0, x) = \alpha(x)$ and $h(1, x) = \beta(x)$ is monotone, and that the complex of the product poset is homeomorphic to $[0, 1] \times \mathbf{K}$. Consequently, h induces a homotopy between α and β by linear extension. \blacktriangleleft

► **Lemma C.18.** *For each G -structure \mathbf{A} , and structures $\mathbf{B}_1, \dots, \mathbf{B}_n$, there is a G -equivariant homotopy equivalence*

$$\text{Hom}(\mathbf{A}, \mathbf{B}_1) \times \dots \times \text{Hom}(\mathbf{A}, \mathbf{B}_n) \sim_G \text{Hom}(\mathbf{A}, \mathbf{B}_1 \times \dots \times \mathbf{B}_n).$$

Proof. For each functor \mathcal{F} , we always have a morphism

$$f \in \text{hom}(\mathcal{F}(\mathbf{B}_1 \times \cdots \times \mathbf{B}_n), \mathcal{F}(\mathbf{B}_1) \times \cdots \times \mathcal{F}(\mathbf{B}_n))$$

defined by the universal property of the product applied to the maps $\mathcal{F}(p_i)$. In the case $\mathcal{F} = \text{Hom}(\mathbf{R}_3, -)$, this morphism is represented by the continuous map

$$\alpha : \text{Hom}(\mathbf{A}, \mathbf{B}_1 \times \cdots \times \mathbf{B}_n) \rightarrow \text{Hom}(\mathbf{A}, \mathbf{B}_1) \times \cdots \times \text{Hom}(\mathbf{A}, \mathbf{B}_n)$$

defined by $\alpha(m) = (p_1 \circ m, \dots, p_n \circ m)$ for each $m \in \text{mhom}(\mathbf{A}, \mathbf{B}_1 \times \cdots \times \mathbf{B}_n)$ and extending linearly. It is straightforward to check that α is an equivariant continuous map. We prove that it has a homotopy inverse β defined on vertices by

$$\beta(m_1, \dots, m_n) : x \mapsto m_1(x) \times \cdots \times m_n(x),$$

and extending linearly. Clearly, β is G -equivariant and $\alpha \circ \beta = 1$.

We show that $\beta \circ \alpha$ is equivariantly homotopic to the identity. It is easy to observe that $\beta(\alpha(m)) \geq m$ for each $m \in \text{mhom}(\mathbf{A}, \mathbf{B}_1 \times \cdots \times \mathbf{B}_n)$. Since both α and β are monotone and equivariant, so is the composition, and the inequality gives an equivariant homotopy between $\beta \circ \alpha$ and the identity by the above lemma. \blacktriangleleft

The proof above immediately gives that $\text{Hom}(\mathbf{R}_3, -)$ preserves products. This is since we showed that $[\alpha]$, which commutes with projections, is an isomorphism.

► Corollary C.19. *For each G -structure \mathbf{A} , the functor $\text{Hom}(\mathbf{A}, -) : \mathbf{rel} \rightarrow \mathbf{htop}_G$ preserves products.*

C.3 The minion homomorphism $\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \mathcal{L}_3$

We can now conclude the proof of the lemma claiming that there is a minion homomorphism $\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \mathcal{L}_3$.

Proof of Lemma 3.2. We have the following minion homomorphisms

$$\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \rightarrow \text{hpol}(\text{Hom}(\mathbf{R}_3, \mathbf{LO}_3), \text{Hom}(\mathbf{R}_3, \mathbf{LO}_4)) \rightarrow \text{hpol}(S^1, P^2).$$

Where the first arrow follows from a combination of Corollary C.19, and Lemma C.15, and the second arrow follows from Lemmata C.16, A.1, and A.14. The last minion is isomorphic to \mathcal{L}_3 by Lemma B.1, which yields the required. \blacktriangleleft

References

- 1 Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$ -Sat is NP-hard. *SIAM Journal on Computing*, 46(5):1554–1573, 2017. URL: <https://eccc.weizmann.ac.il/report/2013/159/>, 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. URL: <http://link.springer.com/10.1007/BFb0082690>, 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. arXiv:2302.03456, doi:10.48550/arXiv.2302.03456.
- 11 Victor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems, 2023. arXiv:2301.05084, 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 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.
- 15 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.
- 16 Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- 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. arXiv:2208.13538, 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. arXiv:2003.11351, 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 Jean-Pierre Serre. Groupes d’homotopie et classes de groupes abeliens. *Annals of Mathematics*, 58(2):258–294, 1953. URL: <http://www.jstor.org/stable/1969789>.
- 29 Tammo tom Dieck. *Transformation Groups*. De Gruyter, Berlin, New York, 1987. URL: <https://doi.org/10.1515/9783110858372>, doi:doi:10.1515/9783110858372.
- 30 Marcin Wrochna. Homomorphism reconfiguration via homotopy. *SIAM J. Discret. Math.*, 34(1):328–350, 2020. doi:10.1137/17M1122578.
- 31 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.
- 32 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.
- 33 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.