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

MAREK FILAKOVSKÝ, Masaryk University, Brno, Czech Republic and Charles University, Prague, Czech Republic
TAMIO-VESA NAKAJIMA, Department of Computer Science; Trinity College, University of Oxford, Oxford, United Kingdom of Great Britain and Northern Ireland and Department of Mathematics and Computer Science, University of Marburg, Marburg, Germany
JAKUB OPRŠAL, School of Computer Science, University of Birmingham, Birmingham, United Kingdom of Great Britain and Northern Ireland
GIANLUCA TASINATO, Institute of Science and Technology Austria, Klosterneuburg, Austria
ULI WAGNER, Institute of Science and Technology Austria, Klosterneuburg, Austria

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).

CCS Concepts: • **Mathematics of computing** → *Approximation algorithms*; *Algebraic topology*; • **Theory of computation → Problems, reductions and completeness;**

Additional Key Words and Phrases: Constraint satisfaction problem, hypergraph colouring, promise problem, topological methods

This research was supported by the Charles University project PRIMUS/21/SCI/014, by the Ministry of Education, Youth and Sports of the Czech Republic under the project MSCAfellow5_MUNI (CZ.02.01.01/00/22_010/0003229), and by the Austrian Science Fund (FWF project P31312-N35). This research was funded by UKRI EP/X024431/1 and by a Clarendon Fund Scholarship. 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.
Authors' Contact Information: Marek Filakovský, Masaryk University, Brno, South Moravian Region, Czech Republic and Charles University, Prague, Prague, Czech Republic; e-mail: filakovsky@fi.muni.cz; Tamio-Vesa Nakajima, Department of Computer Science; Trinity College, University of Oxford, Oxford, England, United Kingdom of Great Britain and Northern Ireland and Department of Mathematics and Computer Science, University of Marburg, Marburg, HE, Germany; e-mail: nakajima@uni-marburg.de; Jakub Opršal (corresponding author), School of Computer Science, University of Birmingham, Birmingham, United Kingdom of Great Britain and Northern Ireland; e-mail: j.oprsal@bham.ac.uk; Gianluca Tasinato, Institute of Science and Technology Austria, Klosterneuburg, Lower Austria, Austria; e-mail: gianluca.tasinato@ist.ac.at; Uli Wagner, Institute of Science and Technology Austria, Klosterneuburg, Lower Austria, Austria; e-mail: uli@ist.ac.at.

ACM 1942-3454/2026/05-ART10
https://doi.org/10.1145/3779121

ACM Reference Format:
Marek Filakovský, Tamio-Vesa Nakajima, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. 2026. Hardness of linearly ordered 4-colouring of 3-colourable 3-uniform hypergraphs. *ACM Trans. Comput. Theory* 18, 2, Article 10 (May 2026), 37 pages. https://doi.org/10.1145/3779121

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 [19]. 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 $O(n(\log\log n)^2/(\log n)^3)$ (Halldórsson [16]); conversely, it is known that the chromatic number cannot be approximated in polynomial time within a factor of $n^{1-\varepsilon}$, for any fixed $\varepsilon > 0$, unless P = NP (Zuckerman [36]). 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 [22] 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 Barto, Bulín, Krokhin, and Opršal [3]. On the other hand, the best polynomial-time algorithm, due to Kawarabayashi, Thorup and Yoneda [21], 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 [3] generalise to give hardness for $k = 2c - 1$ and all $c \geq 3$. For $c \geq 6$, this was improved by Krokhin, Opršal, Wrochna, and Živný [26], who showed that it is hard to colour c-colourable graphs with $k = \binom{c}{\lfloor c/2 \rfloor} - 1$ colours. We remark that conditional hardness (assuming different variants of Khot's *Unique Games Conjecture*) for approximate graph colouring for all $k \geq c \geq 3$ was obtained by Dinur, Mossel, and Regev [13], Guruswami and Sandeep [15], and Braverman, Khot, Lifshitz, and Minzer [7].

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 article 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*** (***PCSPs***); in the process, we considerably extend a topological approach and toolkit for studying approximate colouring that was introduced by Krokhin, Opršal, Wrochna, and Živný [26].

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) $G \to K_k$, where K_k is the complete graph with k vertices. The general formulation of the CSP is then as follows (see Section 2.1 for more details): Fix a relational structure $\mathbf{A}$ (e.g., a graph, or a uniform hypergraph), which parametrises the problem. CSP($\mathbf{A}$) is then the problem of deciding whether a given structure $\mathbf{X}$ allows a homomorphism $\mathbf{X} \to \mathbf{A}$. One of the celebrated results in the complexity theory of CSPs is the Dichotomy Theorem of Bulatov [9] and Zhuk [35], which asserts that for every finite relational structure $\mathbf{A}$, $\mathrm{CSP}(\mathbf{A})$ is either NP-complete, or solvable in polynomial time.

The framework of CSPs can be extended to *PCSP*s, 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 [6], 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 $\mathbf{A} \to \mathbf{B}$. Given an input structure $\mathbf{X}$, the goal is then to distinguish between the case that there is a homomorphism $\mathbf{X} \to \mathbf{A}$, and the case that there does not even exist a homomorphism $\mathbf{X} \to \mathbf{B}$ (these cases are distinct but not necessarily complementary, and no output is required in case neither holds); we denote this decision problem by $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$. For example, $\mathrm{PCSP}(K_3, K_k)$ 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 even 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 [14] and Brakensiek and Guruswami [5]. A variant of hypergraph colouring that is closely connected to approximate graph colouring and generalises (monotone[1]) *1-in-3-SAT* is **linearly ordered hypergraph colouring (LO 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, LO colouring is the same as ordinary graph colouring. Moreover, LO 2-colouring of 3-uniform hypergraphs corresponds to (monotone) 1-in-3-SAT, by viewing the edges of the hypergraph as clauses. In the present article, we focus on 3-uniform hypergraphs: whether such a graph has an LO k-colouring can be expressed as $\mathrm{CSP}(\mathbf{LO}_k)$ for a specific relational structure $\mathbf{LO}_k$ with one ternary relation (the vertex set of $\mathbf{LO}_k$ is $\{1, \dots, k\}$ and the relations are triples with unique maximum, see Section 2.1 for the full definition). In particular, 1-in-3-SAT corresponds to $\mathrm{CSP}(\mathbf{LO}_2)$. Observe that the notion of LO colouring corresponds exactly to the notion of *Unique Maximum Colouring* [10] — however, in the context of promise CSPs these were first identified by [2], and thus we follow their terminology.[2]

The promise version of LO hypergraph colouring was introduced by Barto, Battistelli, and Berg [2], who studied the *promise 1-in-3-SAT* problem. More precisely, let $\mathbf{B}$ be a fixed ternary structure such that there is a homomorphism $\mathbf{LO}_2 \to \mathbf{B}$. Then $\mathrm{PCSP}(\mathbf{LO}_2, \mathbf{B})$ is the following decision problem: Given an instance $\mathbf{X}$ of 1-in-3-SAT, distinguish between the case that $\mathbf{X}$ is satisfiable, and the case that there is not even a homomorphism $\mathbf{X} \to \mathbf{B}$. 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 $\mathrm{PCSP}(\mathbf{LO}_2, \mathbf{LO}_3)$. They conjectured that this problem is NP-hard, and more generally that $\mathrm{PCSP}(\mathbf{LO}_c, \mathbf{LO}_k)$ 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 [30]): $\mathrm{PCSP}(\mathbf{LO}_2, \mathbf{B})$ is either solved by the *affine integer programming relaxation*, or NP-hard. (See Ciardo, Kozik, Krokhin, Nakajima, and Živný [11] for recent progress in this direction.)

Promise LO hypergraph colouring was further studied by Nakajima and Živný [29], 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

[1]In the present article, we will only consider the monotone version of 1-in-3-SAT, i.e., the case where clauses contain no negated variable, and we will often omit the adjective "monotone" in what follows.
[2]We thank Dömötör Pálvölgyi for informing us that LO colourings were also studied under this name.

hypergraphs with a superconstant number of colours, by adapting methods used for similar algorithms for approximate graph colouring, e.g., Blum [4], Kawarabayashi and Thorup [20]. The number of colours used was then reduced by Håstad, Martinsson, Nakajima, and Živný [17] to $\log_2(n)$. In the other direction, it was observed by Nakajima and Živný and by Austrin (personal communications) that $\mathrm{PCSP}(K_{k-1}, K_{c-1})$ reduces to $\mathrm{PCSP}(\mathbf{LO}_k, \mathbf{LO}_c)$ for all $c \geq k \geq 4$ via a simple gadget reduction. Consequently, the hardness of these problems can be derived from known results about hardness of approximate graph colouring (which we briefly outlined above).

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

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

Apart from the intrinsic interest of LO hypergraph colouring, we believe that the main contribution of this article is on a technical level, by extending the topological approach of [26] and bringing to bear more advanced methods from algebraic topology, in particular *equivariant obstruction theory*. To our knowledge, this article is the first that uses these methods in the PCSP context; we view this as a "proof of concept" and believe these tools will be useful to make further progress on approximate graph colouring and related problems.

We remark that, with some additional work, our method could be extended to prove NP-hardness of $\mathrm{PCSP}(\mathbf{LO}_k, \mathbf{LO}_{2k-2})$, $k \geq 3$. As remarked above, this already follows from known hardness results for approximate graph colouring.

After the announcement of our results, Krokhin and Vagnozzi [25] proved that also $\mathrm{PCSP}(\mathbf{LO}_2, \mathbf{LO}_3)$ is NP-hard from which our main theorem follows by a simple reduction.

Plan of the article. In the next section, we present basic notions of the algebraic theory of PCSPs and topology — the main objects and concepts we will be working with. We continue with an overview of the proof of Theorem 1.1, which is presented in Section 3. Sections 4–7 then include the details of the proof; we present a more detailed overview of their content at the end of Section 3.

2 Preliminaries

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

2.1 Promise CSPs

We start by recalling some fundamental notions from the theory of PCSPs, following the presentation of Barto et al. [3] and Krokhin and Opršal [24].

A *relational structure* is a tuple $\mathbf{A} = (A; R_1^{\mathbf{A}}, \ldots, R_k^{\mathbf{A}})$, where A is a set, and $R_i^{\mathbf{A}} \subseteq A^{\mathrm{ar}(R_i)}$ is a relation of arity $\mathrm{ar}(R_i)$. The *signature* of $\mathbf{A}$ is the tuple $(\mathrm{ar}(R_1), \ldots, \mathrm{ar}(R_k))$. For two relational structures $\mathbf{A} = (A; R_1^{\mathbf{A}}, \ldots, R_k^{\mathbf{A}})$ and $\mathbf{B} = (B; R_1^{\mathbf{B}}, \ldots, R_k^{\mathbf{B}})$ with the same signature, a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$, denoted $h\colon \mathbf{A} \to \mathbf{B}$, is a function $h\colon A \to B$ that preserves all relations: for each $i \in [k]$ and $a \in R_i^{\mathbf{A}}$, if we let $h(a)$ denote the componentwise application of h on the elements of a, then $h(a) \in R_i^{\mathbf{B}}$. To express the mere existence of such a homomorphism, we will use the notation $\mathbf{A} \to \mathbf{B}$. We denote the set of all homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ by $\hom(\mathbf{A}, \mathbf{B})$.

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

Definition 2.1 (Promise CSP). Fix two relational structures such that $\mathbf{A} \to \mathbf{B}$. The *promise CSP* with template $(\mathbf{A}, \mathbf{B})$, denoted by $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$, 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 $\mathbf{X} \to \mathbf{A}$, and we are tasked with finding a homomorphism $h\colon \mathbf{X} \to \mathbf{B}$.
— In the *decision* version of the problem, we are given a relational structure $\mathbf{X}$, and we must answer *Yes* if $\mathbf{X} \to \mathbf{A}$, and *No* if $\mathbf{X} \not\to \mathbf{B}$. (These cases are mutually exclusive since $\mathbf{A} \to \mathbf{B}$ 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 $\mathbf{LO}_k$ that appears in our main result. Let $k \in \mathbb{N}$, $k \geq 2$. Then the domain of $\mathbf{LO}_k$ is $[k]$, 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 $a = b < c$, or $a = c < b$, or $b = c < a$, 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 by Barto et al. [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

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

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 $\mathrm{pol}(\mathbf{A}, \mathbf{B})$, and the set of all n-ary polymorphisms by $\mathrm{pol}^{(n)}(\mathbf{A}, \mathbf{B})$.[3]

Concretely, in the special case of structures with a ternary relation, a polymorphism is a mapping $f\colon A^n \to B$ such that, for all sequences of 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 determine the complexity of a PCSP up to log-space reductions. Loosely speaking, the richer the algebraic structure of polymorphisms is, the easier the problem is. We will use a hardness criterion that essentially states that the problem is hard if the polymorphisms have no interesting structure. To define what do we mean by interesting structure, we have to define the notions of *minor*, *minion* and *minion homomorphism.*

Definition 2.3. Fix two sets A and B, and let $f\colon A^n \to B$, $\pi\colon [n] \to [m]$ be functions. The *π-minor* of f is the function $g\colon A^m \to B$ defined by $g(x) = f(x \circ \pi)$, or equivalently

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

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

[3]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.

Abstracting from the fact that the polymorphism of any template are closed under taking minors leads to the following notion of *(abstract) minions*:[4]

Definition 2.4. An *(abstract) minion* $\mathscr{M}$ is a collection of non-empty sets $\mathscr{M}^{(n)}$, where $n > 0$ is an integer, and mappings $\pi^{\mathscr{M}}: \mathscr{M}^{(n)} \to \mathscr{M}^{(m)}$ for $\pi\colon [n] \to [m]$. These satisfy $\pi^{\mathscr{M}} \circ \sigma^{\mathscr{M}} = (\pi \circ \sigma)^{\mathscr{M}}$ for each π and σ for which $\pi \circ \sigma$ is defined, and $(1_{[n]})^{\mathscr{M}} = 1_{\mathscr{M}^{(n)}}$.[5]

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

An important example is the minion of projections denoted by $\mathscr{P}$. Abstractly, it can be defined by $\mathscr{P}^{(n)} = [n]$ and $\pi^{\mathscr{P}} = \pi$. Equivalently, and perhaps more concretely, $\mathscr{P}$ can also be described as follows: Given a finite set A with at least two elements and integers $i \le n$, the ith n-ary *projection* on A is the function $p_i\colon A^n \to A$ 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 $\mathscr{P}$. In particular, $\mathscr{P}$ is also isomorphic to the polymorphism minion $\mathrm{pol}(\mathbf{LO}_2, \mathbf{LO}_2)$.

Definition 2.5. A *minion homomorphism* from a minion $\mathscr{M}$ to a minion $\mathscr{N}$ is a collection of mappings $\xi_n\colon \mathscr{M}^{(n)} \to \mathscr{N}^{(n)}$ that preserve taking minors, i.e., such that for each $\pi\colon [n] \to [m]$, $\xi_m \circ \pi^{\mathscr{M}} = \pi^{\mathscr{N}} \circ \xi_n$. We denote such a homomorphism simply by $\xi\colon \mathscr{M} \to \mathscr{N}$, and write $\xi(f)$ instead of $\xi_n(f)$ when the index is clear from the context.[6]

Using the minion $\mathscr{P}$, 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* $\xi\colon \mathrm{pol}(\mathbf{A}, \mathbf{B}) \to \mathscr{P}$, $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ *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 Matoušek [27] for a detailed and accessible introduction (see also Krokhin et al. [26], in particular for further background on the use of homomorphism complexes in PCSPs).

A *(finite, abstract) simplicial complex* K is finite system of sets that is downward closed under inclusion: if $F \subseteq G \in \mathsf{K}$ then $F \in \mathsf{K}$. The (finite) set $V = \bigcup \mathsf{K}$ is called the set of *vertices* of K, and the sets in K are called *simplices* or *faces* of the simplicial complex. A *simplicial map* $f\colon \mathsf{K} \to \mathsf{L}$ between simplicial complexes is a map between the vertex sets that preserves simplices: if $F \in \mathsf{K}$ then $f(F) \in \mathsf{L}$.

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 $p_0 < p_1 < \cdots < p_k$. 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.

[4]Abstract minions as defined here are a generalisation of so-called *function minions* [3, Definition 2.20]; the relation between a function minion and an abstract minion is analogous to the distinction between a permutation group and a group.
[5]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: $\mathscr{M}(X) = \emptyset$ if and only $X = \emptyset$. The definition given abuses the fact that the sets $[n]$ form a skeleton of the category of finite sets.
[6]A minion homomorphism is a natural transformation between the two functors.

With every simplicial complex K, one can associate a topological space $|\mathsf{K}|$, called the *underlying space* or *geometric realisation* 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 $F, G \in \mathsf{K}$.) Then, in particular, the convex hull $\mathrm{conv}(F)$ is a geometric simplex for every $F \in \mathsf{K}$, and the geometric realisation can be defined as the union $|\mathsf{K}| = \bigcup_{F \in \mathsf{K}} \mathrm{conv}(F)$ of these geometric simplices (see, e.g., [27, Lemma 1.6.2]). We also say that the simplicial complex K is a triangulation of the space $|\mathsf{K}|$. Every simplicial map $f \colon \mathsf{K} \to \mathsf{L}$ between abstract simplicial complexes induces a continuous map $|f| \colon |\mathsf{K}| \to |\mathsf{L}|$ between their geometric realisations. In what follows, we will often blur the distinction between a simplicial complex and its geometric realisation (especially when considering properties that do not depend on a particular triangulation).

Following Matoušek [27, Section 5.9], we define homomorphism complexes as order complexes of the poset of *multihomomorphisms* from one structure to another.[7]

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

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

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

Multihomomorphisms are partially ordered by component-wise comparison: we say that $f \leq g$ if $f(u) \subseteq g(u)$ for all $u \in A$. They can also be composed in a natural way: if $f \in \hom(\mathbf{A}, \mathbf{B})$ and $g \in \mathrm{mhom}(\mathbf{B}, \mathbf{C})$, 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* $\mathrm{Hom}(\mathbf{A}, \mathbf{B})$ is the order complex of the poset of multihomomorphisms from $\mathbf{A}$ to $\mathbf{B}$: the vertices of this simplicial complex are multihomomorphisms from $\mathbf{A}$ to $\mathbf{B}$, and the faces correspond to chains $f_0 < f_1 < \cdots < f_k$ of such multihomomorphisms.

Every homomorphism $f \colon \mathbf{A} \to \mathbf{B}$ induces a simplicial map $f_* \colon \mathrm{Hom}(\mathbf{C}, \mathbf{A}) \to \mathrm{Hom}(\mathbf{C}, \mathbf{B})$ 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 $\mathrm{Hom}(K_2, G)$ is commonly used to study graph colourings.[8] In the present article, we work instead with the homomorphism complex $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$ where $\mathbf{R}_3$ is the structure with three elements and all *rainbow tuples*, i.e., tuples (a, b, c) such that a, b, and c are pairwise distinct. This structure can be seen as a hypergraph analogue of the graph K_2.

Note that a homomorphism $h \colon \mathbf{R}_3 \to \mathbf{A}$ 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 article, 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

[7] 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 complex defined in [26, Definition 3.3].

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

maps (homomorphisms, simplicial maps, continuous maps, etc.). Writing $\mathbb{Z}_3$ multiplicatively as $\{1, \omega, \omega^2\}$, where $\omega^i \cdot \omega^j = \omega^{(i+j) \bmod 3}$, such an action is then 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 $\omega\colon \mathbf{A} \to \mathbf{A}$ 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. We will often write $\omega \cdot x$ instead of $\omega(x)$ for such actions. 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 $h\colon \mathbf{R}_3 \to \mathbf{A}$ by pre-composition, and this action extends naturally to an action of $\mathbb{Z}_3$ on $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$.[9]

It is not hard to show that the action on $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$ 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 $f\colon \mathbf{A} \to \mathbf{B}$, the induced simplicial map $f_*\colon \mathrm{Hom}(\mathbf{R}_3, \mathbf{A}) \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{B})$ (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

Throughout Sections 5.1 and 6, we assume some familiarity with fundamental notions of algebraic topology such as homotopy, homology, and cohomology. We refer to Hatcher [18] for background on the more basic non-equivariant setting, and to May et al. [28, Chapters I and II] for the equivariant setting. Here, we briefly outline a few necessary notions.

Two continuous maps $f, g\colon X \to Y$ between topological spaces are called *homotopic*, denoted $f \sim g$, if there is a continuous map $h\colon X \times [0,1] \to Y$ 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 $h(\cdot, t)\colon X \to Y$ 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 $f\colon X \to Y$ and $g\colon Y \to X$ such that $f \circ g \sim 1_Y$ and $g \circ f \sim 1_X$.

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

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

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

[9] This is analogous to the action of $\mathbb{Z}_2$ on graph homomorphism complexes $\mathrm{Hom}(K_2, G)$ used, for example, by Krokhin et al. [26].

2.6 Equivariant Cohomology

Our proof uses (equivariant) cohomology and obstruction theory. For the definition of regular cohomology recall [18, Chapter 3]. Here we only explain the difference between the equivariant and non-equivariant cases. We introduce equivariant cohomology in a specific case that is sufficient for our purposes, and refer to Bredon [8, Section 1], or tom Dieck [32, Section 2.3] for a detailed treatment.

Consider the ring $\mathbb{Z}[\mathbb{Z}_3]$ defined as the formal integer combinations of elements of $\mathbb{Z}_3 = \{1, \omega, \omega^2\}$ with the multiplication defined by setting $\omega^i \omega^j = \omega^{(i+j) \bmod 3}$ and extending bilinearly. Observe that choosing an action of $\mathbb{Z}_3$ on an abelian group G is the same as choosing a structure of $\mathbb{Z}[\mathbb{Z}_3]$-module M on G. We will henceforth treat abelian groups with such an action as $\mathbb{Z}[\mathbb{Z}_3]$-modules.

Definition 2.9 (Equivariant cohomology). Let X be a topological space with a free $\mathbb{Z}_3$-action. The $\mathbb{Z}_3$-action on X induces a similar action on the chain groups $C_d(X)$, which commutes with the boundary maps. Hence, $C_d(X)$ is a chain of $\mathbb{Z}[\mathbb{Z}_3]$-modules. The *equivariant (or Bredon) cohomology* with coefficients in a $\mathbb{Z}[\mathbb{Z}_3]$-module M is the homology of the cochain complex of $\mathbb{Z}[\mathbb{Z}_3]$-modules defined by

$$C^i_{\mathbb{Z}_3}(X; M) = \hom_{\mathbb{Z}[\mathbb{Z}_3]}\left(C_i^{\mathbb{Z}_3}(X), M\right)$$

with the standard coboundary maps. The cohomology groups are denoted by $H^n_{\mathbb{Z}_3}(X; M)$.[10]

An important concept at the core of obstruction theory is *Eilenberg–MacLane spaces*. An Eilenberg–MacLane space has exactly one non-zero homotopy group.

Definition 2.10. Let G be a non-trivial group and $n \geq 1$ (assume that G is Abelian if $n \geq 2$). Then a space is an *Eilenberg–MacLane space* $K(G, n)$ if it is connected and

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

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., Hatcher [18, Section 4.2]). The key property that we are interested in is that homotopy classes of maps from a topological space X to Eilenberg–MacLane spaces are easily classified using the cohomology groups of X, see [18, Theorem 4.57]. We will make use of an equivariant version of this theorem. A key difference of the equivariant setting is that there is no map that plays the role of the constant map, which would correspond to the 0 of the cohomology group. To obtain the desired isomorphism it is necessary to first choose a reference map (see tom Dieck [32, Theorem II.3.17]).[11] Finally, the statement is true for nice-enough topological spaces X; here we phrase it for *CW-complexes* (see Definition 5.4).

Theorem 2.11. *Let X be a CW-complex with a free $\mathbb{Z}_3$ action, $K(G, n)$ be an Eilenberg–MacLane space with a free $\mathbb{Z}_3$ action, and consider the structure of $\mathbb{Z}[\mathbb{Z}_3]$-module on $G = \pi_n(K(G, n))$ induced by the action on $K(G, n)$. Then, given an equivariant map $f : X \to K(G, n)$, there is an isomorphism*

$$\phi_f\colon\ [X, K(G, n)]_{\mathbb{Z}_3} \to H^n_{\mathbb{Z}_3}(X; G)$$

which in general depends on the equivariant homotopy class of f.

[10]We use a simplified definition of equivariant cohomology which aligns with Bredon cohomology in our specific context. In the general case, the Bredon cohomology has a richer structure than simply that of $\mathbb{Z}[\mathbb{Z}_3]$-module.
[11]This theorem is due to Bredon [8], who uses slightly different topological hypothesis that makes it awkward to use in our setting.

3 Overview of the Proof of Theorem 1.1

We give a detailed overview of the proof and the core techniques used. The result is proved by a combination of topological, combinatorial, and algebraic methods.

In short, our overall goal is to provide a minion homomorphism from the polymorphism minion $\mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4)$ to the minion of projections $\mathscr{P}$ (which is incidentally isomorphic to the polymorphism minion of 3-SAT). The result then follows from the hardness criterion given in Theorem 2.6 [3, Theorem 3.1].

The core of our contribution lies providing deep-enough understanding of polymorphisms of our template so that the minion homomorphism follows. More concretely, in the second and third steps of the proof we present techniques of equivariant obstruction theory in conjunction with a combinatorial argument to obtain the overall result.

We use the topological method, introduced by Krokhin, Opršal, Wrochna, and Živný [26]. The crucial observation of this method is: 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 thus "de-noise" the signal and consider the polymorphisms "up to homotopy" — essentially claiming that the (equivariant) homotopy class of a polymorphism carries the information and everything else is noise. This idea is then formalised using homomorphism complexes and their geometric representations. We present the proof in three steps.

The first step replaces hypergraphs with topological spaces using homomorphism complexes and translates the problem of understanding their polymorphisms into understanding (equivariant) homotopy classes of continuous maps $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)^n \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$. We further simplify the situation by replacing $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ by a circle S^1 and $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ with a (in some sense, topologically simpler) space P. This replacement is necessary since characterising homotopy classes of continuous maps between the homomorphism complexes is still a highly non-trivial problem.

The second step is a classification of equivariant continuous maps $(S^1)^n \to P$ up to homotopy. Using the equivariant obstruction theory, these classes are characterised via a *Bredon cohomology group* of the torus $(S^1)^n$. We are ultimately left with a minion homomorphism χ from $\mathrm{pol}(\mathbf{LO}3, \mathbf{LO}4)$ to a minion $\mathscr{Z}_3$ (defined precisely below).

The third and the final step is a combinatorial argument that studies in detail the minion homomorphism χ. As a conclusion, we are able to show that the image of χ misses everything except the minion of projections $\mathscr{P}$ (which is a subminion of $\mathscr{Z}_3$).

We now explore these steps in more detail:

3.1 Step 1: From Hypergraphs to Topological Spaces

To translate hypergraphs into topological spaces, we consider for each hypergraph $\mathbf{A}$ the topological space $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$. Consequently, we get that each homomorphism $f\colon \mathbf{A} \to \mathbf{B}$ induces a continuous map $f_*\colon \mathrm{Hom}(\mathbf{R}_3, \mathbf{A}) \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{B})$. 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., Kozlov [23]).

This general idea requires a refinement to avoid trivial collapses: 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 $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$ and $\mathrm{Hom}(\mathbf{R}_3, \mathbf{B})$ 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 $f\colon \mathbf{A}^n \to \mathbf{B}$, an equivariant continuous map $f_*\colon \operatorname{Hom}(\mathbf{R}_3, \mathbf{A}^n) \to \operatorname{Hom}(\mathbf{R}_3, \mathbf{B})$. This map does not necessarily preserve minors; nevertheless, it preserves minors *up to homotopy*: for each $\pi\colon [n] \to [m]$, we have that $(f^\pi)_*$ and $(f_*)^\pi$ are equivariantly homotopic. (This is since $\operatorname{Hom}(\mathbf{R}_3, \mathbf{A})^n$ and $\operatorname{Hom}(\mathbf{R}_3, \mathbf{A}^n)$ are only homotopically equivalent and not homeomorphic). Instead, we define a minion homomorphism between the polymorphism minion $\operatorname{pol}(\mathbf{A}, \mathbf{B})$ and the minion of homotopy classes of continuous maps from powers of $\operatorname{Hom}(\mathbf{R}_3, \mathbf{A})$ to $\operatorname{Hom}(\mathbf{R}_3, \mathbf{B})$ as defined in the following definition.

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 $\operatorname{hpol}(X, Y)$ defined by

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

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

$$\zeta\colon\ \operatorname{pol}(\mathbf{A}, \mathbf{B}) \to \operatorname{hpol}(\operatorname{Hom}(\mathbf{R}_3, \mathbf{A}), \operatorname{Hom}(\mathbf{R}_3, \mathbf{B})).$$

This part of the proof follows Krokhin et al. [26], namely this minion homomorphism can be constructed by following the proof of Krokhin et al. [26, Lemma 3.22] while substituting $\mathbf{R}_3$ for $\mathbf{K}_2$, and $\mathbb{Z}_3$ for $\mathbb{Z}_2$.

In order to describe the minion $\operatorname{hpol}(\operatorname{Hom}(\mathbf{R}_3, \mathbf{A}), \operatorname{Hom}(\mathbf{R}_3, \mathbf{B}))$, 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: even maps between spheres of dimensions m and n (i.e., $[S^n, S^m]$), while classified for many pairs (m, n), remain unclassified for infinitely many open cases. Finding such a classification is considered to be a central open problem in algebraic topology. We take advantage of the topological methods developed to solve such problems. Moreover, we may simplify matters considerably by replacing the spaces $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ 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 $X' \to X$ and $Y \to Y'$, then there is a minion homomorphism

$$\eta\colon\ \operatorname{hpol}(X, Y) \to \operatorname{hpol}(X', Y').$$

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

It may be observed that the space $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ is homotopically equivalent to the simplicial complex depicted in Figure 1(a) where the $\mathbb{Z}_3$ acts on each row cyclically. We choose X' so that $X' \to \operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$, 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 vertical 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 nth power of S^1.

The space $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ is a bit more complicated, in particular it is not homotopically equivalent to a 1-dimensional space. Up to equivalence, it is the order complex of the partial order

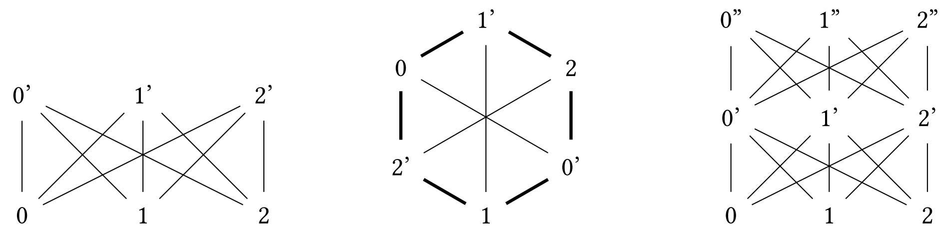

(a) $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ (b) S^1 in $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ (c) $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$

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

depicted in Figure 1(c) where $\mathbb{Z}_3$ acts on rows cyclically.[12] It may be observed that $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ is simply connected, i.e., that $\pi_1(\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)) = 0$, and that $\pi_2(\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4))$ is a non-trivial group. Moreover, the action of $\mathbb{Z}_3$ on $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ induces a non-trivial action of $\mathbb{Z}_3$ on $\pi_2(\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4))$. The precise group and action is irrelevant for us at this point. The space that we use to replace $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$, and denote by P, shares these two properties with $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$; moreover $\pi_n(P) = 0$ for all $n > 2$. Spaces which have only one non-trivial homotopy group (and are sufficiently "nice") are called *Eilenberg–MacLane spaces*, and are 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 is in fact an Eilenberg–MacLane space $K(G, 2)$ where G is a suitable group with a free action of $\mathbb{Z}_3$, chosen in such a way that it allows an equivariant homomorphism $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \to P$ while allowing for much easier classification of maps into it.

In sum after this step, we are given minion homomorphisms

$$\mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathrm{hpol}(\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3), \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)) \to \mathrm{hpol}(S^1, P)$$

Full details on the spaces $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ are to be found in Section 5 including an explicit construction of the replacement space P in Section 5.1. The final minion homomorphism due to a categorical argument given in Section 4.

3.2 Step 2: Classification of Continuous Maps up to Equivariant Homotopy

Next, we prove that the minion $\mathrm{hpol}(S^1, P)$ is isomorphic to the minion $\mathscr{Z}_3$ of affine maps modulo 3, i.e., maps $\mathbb{Z}_3^n \to \mathbb{Z}_3$ of the form $(x_1, \ldots, x_n) \mapsto \sum_{i=1}^n \alpha_i x_i$ where $\alpha_1, \ldots, \alpha_n \in \mathbb{Z}_3$ 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. Since we are interested in equivariant maps and equivariant homotopy, we use a version of equivariant cohomology, called *Bredon cohomology*, introduced by Bredon [8]. 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 has the property that for every $\mathbb{Z}_3$-space X such that there is an equivariant map $X \to P$, there is a bijection $[X, P]_{\mathbb{Z}_3} \simeq H^2_{\mathbb{Z}_3}(X; G)$ where G is the group with a $\mathbb{Z}_3$-action described above. Again, this

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

is a consequence of the equivariant obstruction theory. We then compute that

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

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

$$\mathscr{Z}_3 \to \mathrm{hpol}(S^1, P),$$

and show that it is injective. More precisely, this homomorphism is given by assigning to each affine map $f\colon \mathbb{Z}_3^n \to \mathbb{Z}_3$ (or a tuple of its coefficients) a continuous map $\mu(f)\colon T^n \to P$. We then show 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 Section 6.

The above isomorphism together with homomorphisms ζ and η provides the following lemma (explicitly proven in Section 4).

Lemma 3.2. *There is a minion homomorphism $\chi\colon \mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathscr{Z}_3$ where $\mathscr{Z}_3$ denotes the minion of affine maps over $\mathbb{Z}_3$.*

3.3 Step 3: Combinatorics of Reconfigurations

The minion homomorphism χ is not enough to prove NP-hardness. Although we could conclude from it, for example, that $\mathrm{PCSP}(\mathbf{LO}_3, \mathbf{LO}_4)$ is not solved by any level of Sherali–Adams hierarchy (this is a direct consequence of Reference [12, Theorems 3.3 and 5.2]). To provide hardness, we use combinatorial arguments presented in full detail in Section 7. The main point is to 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 necessary align.) The connection between reconfigurability and topology was described by Wrochna [33], and we use these ideas to connect reconfigurability with our minion homomorphism χ.

We show that any binary polymorphism $f\colon \mathbf{LO}_3^2 \to \mathbf{LO}_4$ is reconfigurable to an essentially unary polymorphism. In particular, we show that there is an increasing function $h\colon \mathbf{LO}_3 \to \mathbf{LO}_4$ 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 $\chi(f) = \chi(g)$ and demonstrate that the image of $\chi_2\colon \mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathscr{Z}_3^{(2)}$ omits an element. More precisely, we have the following lemma, where $\mathscr{P}_3$ denotes the minion of projections on a three element set (which is a subminion of $\mathscr{Z}_3$).

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

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 $\mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathscr{P}_3$.*

Proof. We show that if a subminion $\mathscr{M} \subseteq \mathscr{Z}_3$ contains any non-projection then it contains the map $g\colon (x, y) \mapsto 2x + 2y$. Let $f \in \mathscr{M}^{(n)}$ depend on at least 2 coordinates, and let $f(x_1, \dots, x_n) = \alpha_1 x_1 + \cdots + \alpha_n x_n$. Observe that $\alpha_1, \dots, \alpha_n \in \{0, 1, 2\}$, and also $\sum_i \alpha_i \geq 2$ (else f does not depend on at least 2 coordinates). Hence there exists $S \subseteq [n]$ such that $\sum_{i \in S} \alpha_i = 2$. Since

$\sum_{i=1}^{n} \alpha_i \equiv 1 \pmod 3$, it follows that $\sum_{i \notin S} \alpha_i \equiv 2 \pmod 3$. Hence, letting $\pi \colon [n] \to [2]$ be defined by $\pi(S) = 2$ and $\pi([n] \setminus S) = 1$, we have $f^{\pi}(x, y) = 2x + 2y$, i.e., $f = g$. So $\mathscr{M}$ cannot contain any non-projections.

Finally, the image of ξ is a subminion of $\mathscr{Z}_3$, and since it omits g and every subminion of $\mathscr{Z}_3$ contains $\mathscr{P}_3$, it is equal to $\mathscr{P}_3$ which yields our result. □

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

Full details on proof are presented in the remaining sections below. Step 1 is covered by Sections 4 and 5; the former section provides abstract arguments for existence of the required minion homomorphism, and the latter section provides the description of all the topological spaces involved. The classification of continuous maps up to homotopy required for Step 2 is provided in Section 6. The final combinatorial argument of Step 3 is provided in Section 7.

4 Minion Homomorphisms

The goal of this section is to provide necessary tools and techniques that will allow us to provide minion homomorphisms

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

We remark that the second of the above homomorphisms uses the existence of equivariant maps $S^1 \to \operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \to P$, which are proved in the following section (see Lemma 5.1 and 5.11, respectively). Let us note that similar minion homomorphisms have been constructed in [26, 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. We give slightly vague overview of the basic notions of category theory (precise definitions can be found in any category theory textbook) supported by examples relevant for our proofs.

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

Example 4.1. 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 Set. Relational structures of a fixed signature Σ together with homomorphisms form a category which we denote by Str_Σ or usually just Str when the signature is clear from the context. The categorical notion of isomorphism coincides with structural isomorphism.

Example 4.2. 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 naïve homotopy category 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., $\hom(X, Y) = [X, Y]_G$. The composition is defined by $[f] \circ [g] = [f \circ g]$. We will write just hTop instead of $\mathsf{hTop}_{\mathbb{Z}_3}$.[13]

Isomorphisms in this category coincide with equivariant homotopy equivalences: If $f\colon X \to Y$ and $g\colon Y \to X$ witness a homotopy equivalence, then both $[f]$ and $[g]$ are isomorphisms in 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 4.3. Another useful category is the category whose objects are non-negative integers, and morphisms from n to m are maps $[n] \to [m]$, i.e., $\hom(n, m) = \{f\colon [n] \to [m]\}$. We denote this subcategory of the category of sets by Fin. Note that Fin 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 $\mathscr{F}\colon C_1 \to C_2$ is a mapping between objects and morphisms that preserves composition and identity morphisms. More specifically, we require that $\mathscr{F}(A) \in C_2$ for each $A \in C_1$, and $\mathscr{F}(f) \in \hom(\mathscr{F}(A), \mathscr{F}(B))$ for each $f \in \hom(A, B)$ where $A, B \in C_1$.

Example 4.4. We will view the assignment $\mathscr{B}\colon \mathbf{A} \mapsto \mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$ as a functor $\mathsf{Str} \to \mathsf{hTop}$. This functor is defined on morphism by mapping a homomorphism $f\colon \mathbf{A} \to \mathbf{B}$ to the class $[f_*]$ where $f_*\colon \mathrm{Hom}(\mathbf{R}_3, \mathbf{A}) \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{B})$ is the induced continuous map.

Example 4.5. An abstract minion is a functor $\mathscr{M}\colon \mathsf{Fin} \to \mathsf{Set}$ such that $\mathscr{M}(n) \neq \emptyset$ for all $n > 0$.

Finally, *natural transformations* are morphisms between functors.

Example 4.6. A natural transformation between two minions coincides with the notion of minion homomorphism: a natural transformation $\eta\colon \mathscr{M} \to \mathscr{N}$ is a collection of maps η_n where n ranges through objects of Fin such that, for each $n, m \in \mathsf{Fin}$ and $\pi \in \hom(n, m)$, the following square commutes.

$$\begin{tikzcd} \mathscr{M}(n) \arrow[r, "\pi^{\mathscr{M}}"] \arrow[d, "\eta_n"'] & \mathscr{M}(m) \arrow[d, "\eta_m"] \\ \mathscr{N}(n) \arrow[r, "\pi^{\mathscr{N}}"'] & \mathscr{N}(m) \end{tikzcd}$$

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 $a\colon A \to A'$, 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, in the case of products of sets, structures, or topological spaces, there is a bijection

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

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

[13]We note for the knowledgeable reader that the proper *homotopy category*, which identifies *weakly* homotopy equivalent spaces, would be suitable here as well.

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

Definition 4.7. Let $B_1, \dots, B_n \in \mathcal{C}$. We say that an object P together with morphisms $p_i \in \hom(P, B_i)$ for $i \in [n]$ is the *product* of $B_1, \dots, B_n$ if, for each $A \in \mathcal{C}$ and morphisms $a_i \in \hom(A, B_i)$, there is a unique morphism $a \in \hom(A, P)$ 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 $(p_1, \dots, p_n) \in \hom(P, B_1) \times \dots \times \hom(P, B_n)$. 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 $f \in \hom(P, P')$ and $g \in \hom(P', P)$ such that $p'_i = p_i \circ f$ and $p_i = p'_i \circ g$. Since 1_P is the unique morphism such that $p_i = p_i \circ 1_P$, we get that $g \circ f = 1_P$. Analogously, $f \circ g = 1_{P'}$.

Example 4.8. Products do not always exist. Nevertheless, all the categories we work with in this article (Fin, Set, Str, 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 $f_i \colon Y \to X_i$ be continuous maps. We claim that there is a map $f \colon Y \to X_1 \times \dots \times X_n$ 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 $h_i \colon [0,1] \times Y \to X_i$ between f_i and $p_i \circ g$ for each i. By the universal property of product in topological spaces, these define a homotopy $h \colon [0,1] \times Y \to X_1 \times \dots \times X_n$, 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 4.9. Let $\mathcal{C}$ be a category with finite products, $A, B \in \mathcal{C}$ be two objects, and $n \geq 0$. We define a *polymorphism* from A to B of arity n to be any element $f \in \hom(A^n, B)$ where A^n is the n-fold power of A.

In order to define the *polymorphism minion* $\operatorname{pol}(A, B)$ in such a category $\mathcal{C}$, we need a functor from Fin to Set that assigns to each n, the set $\hom(A^n, B)$. 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 $\mathsf{Fin} \to \mathcal{C}$ and $\mathcal{C} \to \mathsf{Set}$: 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 $\hom(A^n, B)$.

Let us first observe that the assignment $n \mapsto A^n$ can be extended to a functor $A^- \colon \mathsf{Fin} \to \mathcal{C}$ for each $A \in \mathcal{C}$. This can be relatively easily observed in all concrete cases, e.g., if A is a structure $\mathbf{A}$, then, for each $\pi \colon [n] \to [m]$, the mapping $\mathbf{A}^\pi \colon A^m \to A^n$ 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 $\pi \colon [n] \to [m]$ be a mapping. We want to define a homomorphism $p^A_\pi \colon A^m \to A^n$. Using the definition of the nth power, it is enough to give an n-tuple of homomorphisms $a_i \colon A^m \to A$, $i \in [n]$. We let $a_i = p_{\pi(i)}$ where p_j denote projections of the mth power of A. Finally, it is straightforward to check that the assignment $\mathscr{A}(\pi) = p^A_\pi$ preserves compositions and identities.

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

Remark 4.10. If $C = \mathsf{Set}$, then the functor A^- can be also described as a restriction of the contravariant hom-functor $\hom(-, A)$ to Fin viewed as a subcategory of Set.

Definition 4.11. Let C be a category with finite products, and $A, B \in C$ be such that $\hom(A, B)$ is non-empty. We define the *polymorphism minion* $\mathrm{pol}_C(A, B)$ as the composition of functors A^- and $\hom(-, B)$. We will omit the index C whenever the category is clear from the context.

Example 4.12. The polymorphism minion in the category hTop coincides with the minion of homotopy classes of polymorphisms, i.e., for all spaces $X, Y \in \mathsf{hTop}$, we have $\mathrm{pol}_{\mathsf{hTop}}(X, Y) = \mathrm{hpol}(X, Y)$.

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 $\mathscr{F}(A)^-$ is naturally equivalent to the composition of A^- and $\mathscr{F}$.

Definition 4.13. We say that a functor $\mathscr{F}\colon C_1 \to C_2$ *preserves finite products* if, for each $A_1, \dots, A_n \in C_1$ and their product (P, p_i), $(\mathscr{F}(P), \mathscr{F}(p_i))$ is a product of $\mathscr{F}(A_1), \dots, \mathscr{F}(A_n)$.

Lemma 4.14. *If a functor $\mathscr{F}$ preserves products, then $\mathscr{F} \circ A^-$ and $\mathscr{F}(A)^-$ are naturally equivalent.*

Proof. We define a natural equivalence $\eta\colon \mathscr{F} \circ A^- \to \mathscr{F}(A)^-$ by components as $\eta_n\colon \mathscr{F}(A^n) \to \mathscr{F}(A)^n$ is defined as the map given by the n-tuple $\mathscr{F}(p_1), \dots, \mathscr{F}(p_n)\colon \mathscr{F}(A^n) \to \mathscr{F}(A)$. 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. □

4.1 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 [34, Lemma 4.8], although in different flavours it can be tracked down to Lawvere theories.

Lemma 4.15. *If a functor $\mathscr{F}\colon C_1 \to C_2$ preserves products, then there is a minion homomorphism*

$$\xi\colon \mathrm{pol}(A, B) \to \mathrm{pol}(\mathscr{F}(A), \mathscr{F}(B))$$

for all $A, B \in C_1$ such that $\hom(A, B) \neq \emptyset$.

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

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

To show ξ preserves minors, observe that f^π is defined as $f \circ p^A_\pi$, and consequently $\mathscr{F}(f^\pi) = \mathscr{F}(f \circ p^A_\pi) = \mathscr{F}(f) \circ \mathscr{F}(p^A_\pi) = \mathscr{F}(f) \circ p^{\mathscr{F}(A)}_\pi = \mathscr{F}(f)^\pi$. □

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

Lemma 4.16. *Let $A, A', B, B' \in C$ be such that $\hom(A, B)$ is non-empty. Then every pair $a \in \hom(A', A)$ and $b \in \hom(B, B')$ induces a minion homomorphism*

$$\xi\colon \mathrm{pol}(A, B) \to \mathrm{pol}(A', B').$$

Proof. First, observe that a induces a natural transformation $A^- \to (A')^-$, which can be shown by using the definition of products. We will denote its components by $a^n \in \hom(A^n, (A')^n)$. 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.

$$\begin{tikzcd}
(A')^n \arrow[r, "a^n"] \arrow[d, "p_\pi^{A'}"'] & A^n \arrow[d, "p_\pi^A"'] \arrow[dr, "f^\pi"] & & \\
(A')^m \arrow[r, "a^m"] & A^m \arrow[r, "f"'] & B \arrow[r, "b"] & B'
\end{tikzcd}$$

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

4.2 Homomorphism Complexes Preserve Products

We prove that the functor $\mathrm{Hom}(\mathbf{R}_3, -)$ 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., [23, 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 4.17. *Let* K *be an order complex of a poset. If $\alpha, \beta \colon V(\mathsf{K}) \to V(\mathsf{K})$ are monotone maps such that $\alpha(x) \le \beta(x)$ for each x, then their linear extensions are homotopic.*

Proof. Observe that the map $h \colon \{0,1\} \times V(\mathsf{K}) \to V(\mathsf{K})$ 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 $[0,1] \times \mathsf{K}$. Consequently, h induces a homotopy between α and β by linear extension. □

Lemma 4.18. *For each G-structure $\mathbf{A}$, and structures $\mathbf{B}_1, \dots, \mathbf{B}_n$, there is a G-equivariant homotopy equivalence*

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

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

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

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

$$\alpha \colon \mathrm{Hom}(\mathbf{A}, \mathbf{B}_1 \times \cdots \times \mathbf{B}_n) \to \mathrm{Hom}(\mathbf{A}, \mathbf{B}_1) \times \cdots \times \mathrm{Hom}(\mathbf{A}, \mathbf{B}_n)$$

defined by $\alpha(m) = (p_1 \circ m, \dots, p_n \circ m)$ for each $m \in \mathrm{mhom}(\mathbf{A}, \mathbf{B}_1 \times \cdots \times \mathbf{B}_n)$ 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

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

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 $\beta(\alpha(m)) \ge m$ for each $m \in \mathrm{mhom}(\mathbf{A}, \mathbf{B}_1 \times \cdots \times \mathbf{B}_n)$. 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. □

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

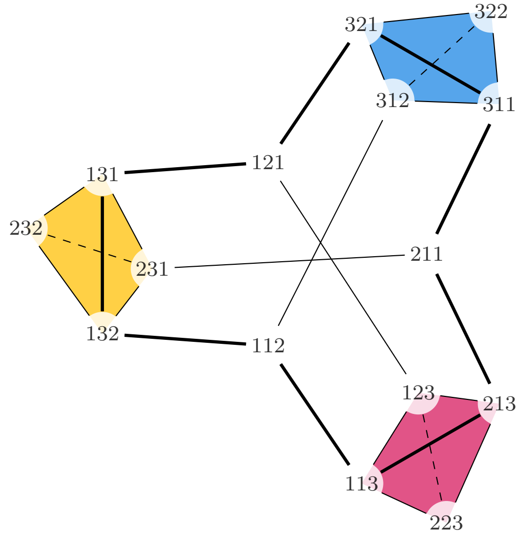

Fig. 2. $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ where only the homomorphisms are explicitly marked.

Corollary 4.19. *For each G-structure $\mathbf{A}$, the functor $\mathrm{Hom}(\mathbf{A}, -)\colon \mathsf{Str} \to \mathsf{hTop}_G$ preserves products.*

Consequently, we get a minion homomorphism between the polymorphism minion of a pair of structures and a polymorphism minion of their homomorphism complexes by a combination of the above and Lemma 4.15.

Corollary 4.20. *For each G-structure $\mathbf{C}$, and a pair of structures $\mathbf{A}$ and $\mathbf{B}$ of the same signature, there is a minion homomorphism*

$$\mathrm{pol}(\mathbf{A}, \mathbf{B}) \to \mathrm{hpol}(\mathrm{Hom}(\mathbf{C}, \mathbf{A}), \mathrm{Hom}(\mathbf{C}, \mathbf{B})).$$

5 Homomorphism Complexes of Linear Orderings

In this section we will investigate the structure of $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$, and we provide a description of the space P which we use in place of $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$.

Luckily the first homomorphism complex $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ is simple-enough so that its structure is apparent from a drawing; Figure 2. 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 $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ is given by rotation.

Lemma 5.1. *There exists a $\mathbb{Z}_3$-equivariant map $S^1 \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$.*

Proof. Consider the cycle in $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ given by the vertices[14]

$$(2,1,1) \to (2,1,3) \to (1,1,3) \to (1,1,2) \to (1,3,2) \to (1,3,1) \to (1,2,1) \to (3,2,1) \to (3,1,1) \to (2,1,1).$$

This cycle is highlighted in Figure 2. Observe that it is invariant under the $\mathbb{Z}_3$ action. Thus we see that there is a $\mathbb{Z}_3$-equivariant map $S^1 \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$. □

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

Now, we turn to $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$. 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 $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ to some sufficiently homotopically simple $\mathbb{Z}_3$ space P (an Eilenberg–MacLane space with a free $\mathbb{Z}_3$-action). 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 $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ into P is not entirely straightforward; instead, it is much easier to define an intermediate $\mathbb{Z}_3$-space L_4 and show that $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \to L_4$ and $L_4 \to P$ equivariantly.

Definition 5.2. Let $n \geq 1$, and consider the natural linear order on $[n]$, and an antichain order on $\mathbb{Z}_3$ (i.e., the three elements are incomparable). We define L_{n+1} as the order complex of the *lexicographically ordered* poset $[n] \times \mathbb{Z}_3$, i.e., the poset whose elements are pairs (i, ω^α), where ω is the generator of $\mathbb{Z}_3$, and $(i, \omega^\alpha) < (j, \omega^\beta)$ if $i < j$. The group $\mathbb{Z}_3$ acts on the poset $[n] \times \mathbb{Z}_3$, and consequently on L_n, by the natural action on the second coordinate.

Lemma 5.3. *There is a $\mathbb{Z}_3$-map $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_n) \to L_n$.*

Proof. In order to provide the required equivariant map, we will provide an equivariant simplicial map to the barycentric subdivision of L_n. By definition (see [27, Definition 1.7.2]), this is just the order complex of F, the poset of nonempty simplices of L_n ordered by inclusion. Hence it suffices to give a monotone map from $\mathrm{mhom}(\mathbf{R}_3, \mathbf{LO}_n)$ to F.

First, we define an auxiliary function $h\colon\ \mathrm{hom}(\mathbf{R}_3, \mathbf{LO}_n) \to [n-1] \times \mathbb{Z}_3$ by

$$h(f) = (\max_j f(j) - 1, \omega^{\arg\max_j f(j)}).$$

Note that h respects the $\mathbb{Z}_3$-action.

We extend h to a map $\phi\colon\ \mathrm{mhom}(\mathbf{R}_3, \mathbf{LO}_n) \to F$ as follows:

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

We need to check that ϕ is well-defined, i.e., that $\phi(m)$ is a chain in F for each $m \in \mathrm{mhom}(\mathbf{R}_3, \mathbf{LO}_n)$. For a contradiction assume that $f, g \leq m$ are homomorphisms such that $h(f)$ and $h(g)$ are incomparable. This means that $\max f = \max g$ but the maximum is attained at distinct points; say $\max f = f(1) = g(2) = \max g$. The function $f' \leq m$ defined by $f'(1) = f(1)$, $f'(2) = g(2)$, $f'(3) = f(3)$ is thus not a homomorphism — since $(f'(1), f'(2), f'(3))$ does not have a unique maximum — which yields a contradiction with the fact that m is a multihomomorphism.

It is straightforward to check that ϕ is monotone, and equivariant. □

5.1 Constructing the Replacement Space P

We now present an explicit construction of the space P and an equivariant map $L_4 \to P$. Our requirement on P is that it is topologically "particularly nice", in technical terms, P is constructed as a *CW-complex* (we formally introduce this notion below), it is an Eilenberg–MacLane space, and has a free $\mathbb{Z}_3$ action. In plain words, we construct P by first taking a 2-dimensional space which is an approximation of L_4 (see Figure 3), and then change this space by attaching higher-dimensional discs to ensure that $\pi_n(P) = 0$ for all $n > 2$.

5.1.1 CW-complexes. The notion of a *CW-complex* can be viewed as a generalisation of a simplicial complex; similarly to simplicial complexes, they allow for a discretisation of topological spaces, but they allow more freedom in constructing the space from *cells* of a given dimension. 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 Hatcher [18, p. 5].

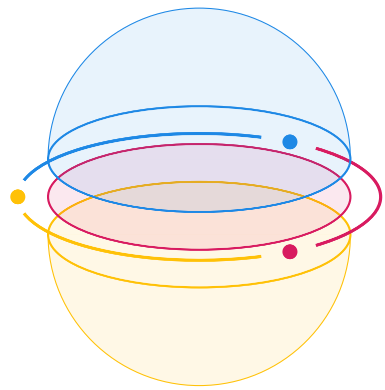

Fig. 3. The 2-skeleton of P.

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 $\phi\colon S^n \longrightarrow Y$ (called attaching map) to define a new space

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

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.[15]

Definition 5.4. A *CW-complex (of finite type)* is the union X of spaces

$$\mathrm{sk}_0 X \subseteq \mathrm{sk}_1 X \subseteq \cdots \subseteq \mathrm{sk}_d X \subseteq \cdots$$

constructed inductively as follows:

- $\mathrm{sk}_0 X = \{p_0, \dots, p_k\}$ is a disjoint union of points.
- Given $\mathrm{sk}_d X$, we can attach a finite number of $(d+1)$-discs $D_1^{d+1}, \dots, D_\ell^{d+1}$ along some attaching maps $\phi_i\colon \partial D_i^{d+1} = S^d \longrightarrow \mathrm{sk}_d X$ to obtain the new space $\mathrm{sk}_{d+1} X$ that contains $\mathrm{sk}_d X$.

Finally, we let

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

For any $d \geq 0$, the space $\mathrm{sk}_d X$ 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 $\mathrm{sk}_i X = \mathrm{sk}_{i+1} X$ for any $i \geq n$ (i.e., no cells of dimension higher than n were used in "building" X).

We note that all concrete spaces mentioned in this article are in fact CW-complexes. In particular, 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. We also note that if X and Y are CW-complexes, then $X \times Y$ has a natural structure of a CW-complex: for example, 2-cells of $X \times Y$ are either products of a 1-cell of X and a 1-cell of Y or products of a 0-cell and 2-cell of the two spaces. Maps between CW-complexes that preserve the structure are called *cellular*.

Definition 5.5. A map $f\colon X \longrightarrow Y$ between CW-complexes is *cellular* if for all $d \geq 0$

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

A CW-complex with an action of $\mathbb{Z}_3$ by cellular maps is called a *$\mathbb{Z}_3$-CW-complex*.

[15]In the language of category theory: Z is the pushout of the diagram $Y \longleftarrow S^n \hookrightarrow D^{n+1}$ where the first arrow is ϕ.

CW-complexes are the natural way to "discretize" homotopy theory: They are much more flexible than simplicial complexes while still giving an manageable description of the available maps thanks to the Cellular Approximation Theorem [18, Theorem 4.8].

Theorem 5.6 (Cellular Approximation Theorem). *Let X, Y be CW-complex and $f\colon X \longrightarrow Y$ a continuous map. Then there is a cellular map $\tilde{f}\colon X \longrightarrow Y$ 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 $x \in A$, $0 \le t \le 1$.*

This means that, as far as homotopy theory is concerned, every map between CW-complex is a cellular map. We remark that the theorem above can be generalised to (free) Z_3-equivariant spaces in a straightforward way.

Remark 5.7. 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 $d \ge 0$ and let $\iota_d\colon \operatorname{sk}_d X \hookrightarrow X$. Then $(\iota_d)_*\colon \pi_d(\operatorname{sk}_d X, x_0) \longrightarrow \pi_d(X, x_0)$ is surjective while $(\iota_{d+1})_*\colon \pi_d(\operatorname{sk}_{d+1} X, x_0) \longrightarrow \pi_d(X, x_0)$ is an isomorphism.

5.1.2 The 2-skeleton of P. We first construct the 2-skeleton of P which, on one hand, is easy to describe geometrically, and on the other hand, determines several key properties of the space. Since P is a CW-complex, the 2-skeleton is also defined inductively (the addition in the indices is considered modulo 3):

- $\operatorname{sk}_0 P$ is the discrete space $\{v_0, v_1, v_2\}$.
- $\operatorname{sk}_1 P$ has three edges (1-cells) $\{e_0, e_1, e_2\}$ glued along to the following maps:

$$\phi_i\colon \partial e_i \longrightarrow \operatorname{sk}_0 P, \quad \phi_i(\partial e_i) = v_{i+1} - v_i.$$

- $\operatorname{sk}_2 P$ has three 2-cells $\{d_0, d_1, d_2\}$ glued along the following maps:

$$\psi_i\colon \partial d_i \longrightarrow \operatorname{sk}_1 P, \quad \psi_1(\partial d_i) = e_i + e_{i+1} + e_{i+2}.$$

Geometrically (see Figure 3), $\operatorname{sk}_1 P$ is a circle S^1 which is decomposed into three arcs so that the $\mathbb{Z}_3$ action via rotation by $2\pi/3$ is cellular. Moreover, $\operatorname{sk}_2 P$ is a sphere S^2 with a disc glued along an equator $\operatorname{sk}_1 P$ where the action of $\mathbb{Z}_3$ cyclically exchanges the three discs around while rotating them by $2\pi/3$ (to maintain consistency on the equator). Formally, the generator ω acts cellularly as

$$\begin{aligned}\omega \cdot v_i &= v_{i+1}\\ \omega \cdot e_i &= e_{i+1}\\ \omega \cdot d_i &= d_{i+1}\end{aligned}$$

The next lemma describes a few properties of the space $\operatorname{sk}_2 P$ that will play role in our argument. These three properties are also shared by the space P.

Lemma 5.8. *The space $\operatorname{sk}_2 P$ has the following properties:*

(1) *the $\mathbb{Z}_3$ action on $\operatorname{sk}_2 P$ is free;*
(2) *$\operatorname{sk}_2 P$ is simply connected; and*
(3) *$\pi_2(\operatorname{sk}_2 P) \simeq \mathbb{Z} \oplus \mathbb{Z}$ where the induced action of the generator ω of $\mathbb{Z}_3$ is multiplication by the matrix*

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

Proof. (1) The action is cellular and no cell is mapped to itself, therefore it is free.

(2) The space $\mathrm{sk}_2 P$ is homotopy equivalent to a wedge of 2-spheres $S^2 \vee_{v_0} S^2$, i.e., the space obtained by glueing two 2-disks D^2 to a point v_0 along their boundary. By the Cellular Approximation Theorem, every map $\gamma\colon (S^1, p) \longrightarrow (\mathrm{sk}_2 P, v_0) \sim (S^2 \vee_{v_0} S^2, v_0)$ can be assumed to map in $\mathrm{sk}_1(S^2 \vee_{v_0} S^2) = \{v_0\}$ and thus nullhomotopic.

(3) Since $\mathrm{sk}_2 P$ is simply connected, by Hurewicz theorem [18, see Section 4.2], there is an isomorphism between $\pi_2(\mathrm{sk}_2 P)$ and $H_2(\mathrm{sk}_2 P)$ that is an isomorphism of $\mathbb{Z}[\mathbb{Z}_3]$-modules. By direct computation, it is easy to show that the chains $\beta_i = d_i - d_{i+1} \in C_2(\mathrm{sk}_2 P)$ are cycles and any two out of $\{[\beta_i]\}_{0\le i\le 2}$ form a basis for $H_2(\mathrm{sk}_2 P) \cong \mathbb{Z} \oplus \mathbb{Z}$. Fix $[\beta_0]$ and $[\beta_1]$, then the multiplication by ω is

$$\omega \cdot [\beta_0] = [\omega \cdot \beta_0] = [\beta_1]$$
$$\omega \cdot [\beta_1] = [\omega \cdot \beta_1] = [\beta_2] = -[\beta_0] - [\beta_1],$$

as we wanted to show. □

Remark 5.9. Although it is not required for our proof, we note that the third homotopy group of $\mathrm{sk}_2 P$ is non-trivial. This could be shown by observing that the composition of the Hopf fibration $S^3 \longrightarrow S^2$ with an inclusion $S^2 \longrightarrow \mathrm{sk}_2 P$, that avoids one of the discs, is not nulhomotopic. Similar parallel could be drawn between higher homotopy groups of $\mathrm{sk}_2 P$ and S^2, and these observations make classifying maps into $\mathrm{sk}_2 P$ complicated. The final space P does not share this property.

Finally, the following lemma will ensure that we have an equivariant map $L_4 \longrightarrow P$.

Lemma 5.10. *There is an equivariant map $L_4 \longrightarrow \mathrm{sk}_2 P$.*

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 the standard obstruction theoretic arguments well.

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 $V \in \mathrm{sk}_0 L_4$, 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 $\mathrm{sk}_0 L_4$ 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 $\mathrm{sk}_2 P$ 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 $\mathrm{sk}_1 L_4$: if $\omega^\alpha \cdot E_j$ 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 $f_1|_{\partial D_k} \colon S^1 \longrightarrow \mathrm{sk}_2 P$ is a loop in $\mathrm{sk}_2 P$. Since $\mathrm{sk}_2 P$ is simply connected, $f_1|_{\partial D_k}$ extends to a map defined on the whole disc $f_2\colon D_k \cong D^2 \longrightarrow \mathrm{sk}_2 P$. Once again, completing the definition of f_2 along the orbits in the same way as before we obtain an equivariant map $f = f_2\colon L_4 \longrightarrow \mathrm{sk}_2 P$, as desired. □

5.1.3 *Higher Dimensional Cells of P.*

In this subsection, we described cells of P of dimensions $d > 2$. The purpose of this construction is ensure that $\pi_d(P) = 0$ for all $d > 2$, i.e., that P is an Eilenberg–MacLane space. We achieve that by inductively glueing $(i+1)$-discs to the i-skeleton $\mathrm{sk}_i P$ obtained at the previous step. We give an explicit construction, although it should be noted it is based on a standard method of algebraic topology.

First, we let $\operatorname{sk}_3 P = \operatorname{sk}_2 P$, i.e., we don't glue any 3-discs. This is because we want to keep the second homotopy group unchanged.

Next, we glue cells to representatives of generators of $\pi_3(\operatorname{sk}_3 P)$ so that they become nullhomotopic in the new space. By a result of Serre [31], 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 $\pi_3(\operatorname{sk}_3 P)$. We glue three 4-dimensional discs $\{D^4_{0,i}, D^4_{1,i}, D^4_{2,i}\}$ 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

$$\operatorname{sk}_4 P = \bigcup_{0 \le i \le d_4} \left(\operatorname{sk}_3 P \cup_{\gamma_i} D^4_{0,i} \cup_{\omega \cdot \gamma_i} D^4_{1,i} \cup_{\omega^2 \cdot \gamma_i} D^4_{2,i} \right)$$

Since we glued disc of dimension 4, $\operatorname{sk}_3 \operatorname{sk}_4 P = \operatorname{sk}_3 P$; by cellular approximation, $\pi_i(\operatorname{sk}_4 P) \cong \pi_i(\operatorname{sk}_3 P)$ for $i \le 2$, while $\iota_* \colon \pi_3(\operatorname{sk}_3 P) \longrightarrow \pi_3(\operatorname{sk}_4 P)$ is surjective. For any $[\gamma] \in \pi_3(\operatorname{sk}_3 P)$, $[\gamma] = \sum n_i [\gamma_i]$ hence $\iota_*([\gamma]) = \sum n_i \iota_*([\gamma_i]) = 0$. Therefore, $\pi_3(\operatorname{sk}_4 P) = 0$. What is more, $\operatorname{sk}_4 P$ is still a finite CW-complex hence Serre theorem still applies and $\pi_\ell(\operatorname{sk}_4 P)$ remains finitely generated for any ℓ. Finally, the inclusion $\iota\colon\ \operatorname{sk}_3 P \hookrightarrow \operatorname{sk}_4 P$ is $\mathbb{Z}_3$-equivariant by construction and the action on $\operatorname{sk}_4 P$ remains consistent with the action on the 3-skeleton.

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

$$\operatorname{sk}_{k+1} P = \bigcup_{0 \le i \le d_k} \left(\operatorname{sk}_k P \cup_{\gamma} D^{k+1} \cup_{\omega \cdot \gamma} D^{k+1} \cup_{\omega^2 \cdot \gamma} D^{k+1} \right)$$

By Cellular Approximation, we have not changed anything in the homotopy groups in dimension less than k and $\iota_* \colon \pi_k(\operatorname{sk}_k P) \longrightarrow \pi_k(\operatorname{sk}_{k+1} P)$ is surjective. This again implies that $\pi_k(\operatorname{sk}_{k+1} P) = 0$ and by Serre's Theorem every homotopy group remains finitely generated.

Recall, that the final space P is defined by

$$P = \bigcup_{k \ge 2} \operatorname{sk}_k P.$$

By construction, P is a CW-complex and since $\operatorname{sk}_{k+1} P$ is its $(k+1)$-skeleton, we have, for all $k \ge 1$,

$$\pi_k(P) = \pi_k(\operatorname{sk}_{k+1} P) \cong \begin{cases} \pi_2(\operatorname{sk}_3 P) & \text{if } k = 2, \\ 0 & \text{otherwise.} \end{cases}$$

These properties imply that P is the Eilenberg–MacLane space $K(\mathbb{Z} \oplus \mathbb{Z}, 2)$. However, due to our careful construction, P has a free cellular $\mathbb{Z}_3$-action that acts transitively on vertices, edges, and 2-cells.

Finally, we have the required equivariant map as phrased in the following lemma.

Lemma 5.11. *There is an equivariant map* $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \to P$.

Proof. We have $\operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \longrightarrow L_4 \longrightarrow \operatorname{sk}_2 P \longrightarrow P$ where the first two maps follow from Lemmas 5.3 and 5.10, and the last one from the definition of P. □

5.2 The Minion Homomorphism

We conclude the section with a proof of the existence of a minion homomorphism $\operatorname{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \longrightarrow \operatorname{hpol}(S^1, P)$ which is one of the two pieces that constitutes Lemma 3.2. The second piece, that the latter minion is isomorphic to $\mathscr{Z}_3$, is proved in the following section.

Lemma 5.12. *There is a minion homomorphism* $\operatorname{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \operatorname{hpol}(S^1, P)$.

Proof. The required minion homomorphism is the composition of the following two minion homomorphisms

$$\mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathrm{hpol}(\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3), \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)) \to \mathrm{hpol}(S^1, P),$$

where the first arrow follows from Corollary 4.20, and the second arrow follows from Lemmata 4.16, 5.1, and 5.11. □

6 Classification of Continuous Maps up to Equivariant Homotopy

In this section, we use the fact that P is the Eilenberg–MacLane space $K(\mathbb{Z} \oplus \mathbb{Z}, 2)$ with a free $\mathbb{Z}_3$ action as shown in the previous section. In particular, the following is a direct consequence of Theorem 2.11:

$$[T^n, P]_{\mathbb{Z}_3} \simeq H^2_{\mathbb{Z}_3}(T^n, G).$$

The goal of this section is to classify equivariant maps $[T^n, P]_{\mathbb{Z}_3}$. Consequently, this classification will provide the following lemma which is a key part of our proof.

Lemma 6.1. *There is a minion isomorphism* $\mu\colon \mathscr{Z}_3 \to \mathrm{hpol}(S^1, P)$.

Recall that $\mathscr{Z}_3$ is the minion of affine maps $f\colon \mathbb{Z}_3^n \to \mathbb{Z}_3$ such $f(1, \dots, 1) = 1$. 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 section, we view $\mathscr{Z}_3$ as an abstract minion whose elements are tuples of the form above, in particular, we have $\mathscr{Z}_3^{(n)} \subseteq \mathbb{Z}_3^n$.

The proof of Lemma 6.1, 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 $S^1 \to S^1$ (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 maps. These maps are then composed with a fixed embedding $S^1 \to P$ (the choice of a concrete embedding is irrelevant since P is simply connected) to obtain an equivariant map $T^n = (S^1)^n \to P$.

Definition 6.2. Fix an embedding $\iota\colon S^1 \hookrightarrow P$ (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* $m_\alpha\colon T^n \to P$, 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

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

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

Remark 6.3. It is not hard to observe that the assignment $\alpha \mapsto m_\alpha$ preserves minors when α is interpreted as a function $f\colon \mathbb{Z}^n \to \mathbb{Z}$. While we implicitly use this minion homomorphism, this is not the minion homomorphism we are looking for — importantly, $\mathscr{Z}_3$ 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 (in Section 6.1) that, 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 allows us to define an injective mapping $\mu\colon \mathscr{Z}_3 \to \mathrm{hpol}(S^1, P)$ — we view each element $\alpha \in \mathscr{Z}_3^{(n)}$ as a tuple in $\mathbb{Z}$ with $\alpha_i \in \{0, 1, 2\}$.

In Section 6.2, we show that any equivariant map $T^n \to P$ is homotopic to some monomial map m_α for $\alpha \in \mathscr{Z}_3^{(n)}$. This is done by showing that $\#[T^n, P]_{\mathbb{Z}_3} = 3^{n-1}$ which we show via obstruction theory — the bijection is given by a bijection between $[T^n, P]_{\mathbb{Z}_3}$ and the Bredon cohomology $H^2_{\mathbb{Z}_3}(T^n, G)$ (here use that P is an Eilenberg–MacLane space), and a computation that gives that $H^2_{\mathbb{Z}_3}(T^n, G) \simeq \mathbb{Z}_3^{n-1}$. Consequently, we get that μ is a bijection (or more precisely, that $\mu_n \colon \mathscr{Z}^{(n)} \to [T^n, P]_{\mathbb{Z}_3}$ 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 6.1) preserves minors. Similarly to homomorphisms of algebraic structures (e.g., groups), an inverse of a bijective minion homomorphism is a homomorphism as well.

6.1 Monomial Maps are all Homotopically Distinct

Here, we prove the following:

Lemma 6.4. *Let $\alpha, \beta \in \mathscr{Z}_3^{(n)}$. If m_α is equivariantly homotopic to m_β, then $\alpha = \beta$.*

The proof of the lemma is built on the following informal geometric intuition. By cellular approximation, a homotopy between two maps $f_0, f_1 \colon T^n \to P$ "drags" around the image of any of the coordinate circles (i.e., the image of $c_i \colon S^1 \hookrightarrow T^n$ 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. 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 formalisation 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 6.5. Let $e^j \in C^1(P)$ and $d^j \in C^2(P)$ 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 ith 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 $f \colon T^n \to P$ is

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

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

Lemma 6.6. *Let $f_0, f_1 \colon T^n \to P$ be 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

$$f_0^*(d^0)(b_i) - f_1^*(d^0)(b_i) = (\delta h^*(d^0))(b_i) + (h^*\delta(d^0))(b_i) = h^*(d^0)(x_i - \omega \cdot x_i) + 0$$
$$f_0^*(d^0)(B_i) - f_1^*(d^0)(B_i) = (\delta h^*(d^0))(B_i) + (h^*\delta(d^0))(B_i) = h^*(d^0)(x_i - \omega^2 \cdot x_i) + 0$$

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

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

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 6.7. *Let $\alpha \in \mathbb{Z}^n$ such that $\sum_i \alpha_i = 1 \pmod 3$. Then $\deg_i(m_\alpha) = \alpha_i \bmod 3$.*

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 $\deg_i(m_\alpha) = \alpha_i \bmod 3$. □

We can now conclude the proof of Lemma 6.4 which follows immediately from the above.

Proof of Lemma 6.4. 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 6.7, and therefore m_α and m_β are not equivariantly homotopic by Lemma 6.6. □

Before, we progress further, let us discuss one more consequence of Lemma 6.4, namely, the following.

Lemma 6.8. *The mapping $\gamma\colon \operatorname{hpol}(S^1, P) \to \mathscr{Z}_3$ defined by*

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

satisfies $\gamma([m_\alpha]^\pi) = \gamma([m_\alpha])^\pi$ for each $\alpha \in \mathbb{Z}^n$ with $\sum_i \alpha_i = 1$ and $\pi\colon [n] \to [m]$, 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 6.6). Further, using Lemma 6.7, we have that

$$\gamma([m_\alpha]) = (\alpha_1 \bmod 3, \dots, \alpha_n \bmod 3)$$

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

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

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 $\operatorname{hpol}(S^1, P)$, i.e., γ (and consequently μ) preserves all minors.

6.2 Monomial Maps Cover all Possible Maps up to Homotopy

To show that there are no other maps besides the monomial maps m_α, $\alpha \in \mathscr{Z}^3$ up to equivariant homotopy, we will rely on Theorem 2.11 in the particular case of $X = T^n$ and P playing the role of $K(\pi_2(P), 2)$, which asserts that there is a bijection $[T^n, P]_{\mathbb{Z}_3} \simeq H^2_{\mathbb{Z}_3}(T^n; \pi_2(P))$. The rest of this section will be dedicated to precisely defining the objects involved in Theorem 2.11 and computing the Bredon cohomology of the torus, i.e., proving the following proposition.

Proposition 6.9. *For all $n, d \geq 1$, $H^d_{\mathbb{Z}_3}(T^n; \pi_2(P)) \cong \mathbb{Z}_3^{\binom{n-1}{d-1}}$.*

6.2.1 Equivariant Cohomology: A Primer. We will compute equivariant cohomology of certain CW-complexes with a free cellular $\mathbb{Z}_3$-action. Similarly to the regular cohomology, it is enough to consider the chains that are obtained as a linear combinations of cells. More precisely, we consider the ring $\Lambda = \mathbb{Z}[\mathbb{Z}_3]$, and let $C_d(X)$ be the free Λ-module generated by the d-dimensional cells of X. Since the action is cellular, it commutes with the boundary maps $\partial\colon C_d(X) \to C_{d-1}(X)$. The equivariant cohomology of X with coefficients in a Λ-module N is then equivalently defined as the cohomology of the equivariant cochain:

$$C^i_{\mathbb{Z}_3}(X;N) = \hom_\Lambda\left(C_i^{\mathbb{Z}_3}(X), N\right)$$

We start with a few technical lemmas that will be useful in our computation of cohomology of tori.

Lemma 6.10. *If the action on X is free and cellular, then $C_\bullet^{\mathbb{Z}_3}(X)$ is a chain complex of free Λ-modules.*

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

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

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. ◻

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

$$0 \to N \to M \to Q \to 0,$$

there is a corresponding short exact sequence of cochain complexes

$$0 \to C^\bullet_{\mathbb{Z}_3}(X;N) \to C^\bullet_{\mathbb{Z}_3}(X;M) \to C^\bullet_{\mathbb{Z}_3}(X;Q) \to 0,$$

and thus a long exact sequence in cohomology

$$\cdots \to H^{i-1}_{\mathbb{Z}_3}(X;Q) \to H^i_{\mathbb{Z}_3}(X;N) \to H^i_{\mathbb{Z}_3}(X;M) \to H^i_{\mathbb{Z}_3}(X;Q) \to \cdots$$

We consider $\mathbb{Z}$ as a Λ-module by taking the trivial action of $\mathbb{Z}_3$ on $\mathbb{Z}$; note that this is the only such action. We have the following diagram of Λ-modules, where $I = (1+\omega+\omega^2)\Lambda$ is the ideal generated by $1+\omega+\omega^2$, and Z_{cyc} is the module obtained by considering the action on $\mathbb{Z}^3$ which permutes the three coordinates cyclically:

$$\begin{tikzcd} \mathbb{Z} \arrow[r, "\phi_1"] \arrow[d, hook, "d"'] & I \arrow[d, hook, "\iota"] \\ Z_{\mathrm{cyc}} \arrow[r, "\phi_2"] & \Lambda \end{tikzcd}$$

The maps in the diagram are defined as: $d(n) = (n,n,n)$, ι is the inclusion, $\phi_1(n) = (1+\omega+\omega^2)n$, and $\phi_2(n_0,n_1,n_2) = n_0 + n_1\omega + n_2\omega^2$. Moreover, observe that ϕ_1 and ϕ_2 are isomorphisms.

Furthermore, assume that X is a CW-complex with a free cellular $\mathbb{Z}_3$-action, and let $p\colon X \to X/\mathbb{Z}_3$ be the projection map that maps each element of X to its orbit. We have two isomorphisms of chains of Abelian groups:

$$h_1\colon C^\bullet(X/\mathbb{Z}_3,\mathbb{Z}) \to C^\bullet_{\mathbb{Z}_3}(X;\mathbb{Z}) \qquad h_1(\alpha)\colon \sigma \mapsto \alpha(p(\sigma))$$
$$h_2\colon C^\bullet(X,\mathbb{Z}) \to C^\bullet_{\mathbb{Z}_3}(X;Z_{\mathrm{cyc}}) \qquad h_2(\alpha)\colon \sigma \mapsto \big(\alpha(\sigma), \alpha(\omega\cdot\sigma), \alpha(\omega^2\cdot\sigma)\big).$$

Lemma 6.11. *Let X be a CW complex with a free $\mathbb{Z}_3$-action, then the following diagram commutes*

$$\begin{tikzcd}
C^\bullet(X/\mathbb{Z}_3;\mathbb{Z}) \arrow[r, "h_1"] \arrow[d, "p^*"] & C^\bullet_{\mathbb{Z}_3}(X;\mathbb{Z}) \arrow[r, "(\phi_1)_*"] \arrow[d, "d_*"] & C^\bullet_{\mathbb{Z}_3}(X;I) \arrow[d, "\iota_*"] \\
C^\bullet(X;\mathbb{Z}) \arrow[r, "h_2"] & C^\bullet_{\mathbb{Z}_3}(X;Z_{\mathrm{cyc}}) \arrow[r, "(\phi_2)_*"] & C^\bullet_{\mathbb{Z}_3}(X;\Lambda)
\end{tikzcd}$$

Proof. The right square commutes by functoriality of $C^\bullet_{\mathbb{Z}_3}(X;-)$. It is enough to show that the left square commutes. Let $d \geq 0$ and $\alpha \in C^d(X/\mathbb{Z}_3;\mathbb{Z})$. Then we have, for all $\sigma \in C_d^{\mathbb{Z}_3}(X)$,

$$\begin{aligned}
h_2(p^*(\alpha))(\sigma) &= \big(p^*(\alpha)(\sigma), p^*(\alpha)(\omega\cdot\sigma), p^*(\alpha)(\omega^2\cdot\sigma)\big) = \big(\alpha(p(\sigma)), \alpha(p(\omega\cdot\sigma)), \alpha(p(\omega^2\cdot\sigma))\big) \\
&= \big(\alpha(p(\sigma)), \alpha(p(\sigma)), \alpha(p(\sigma))\big) = \big(h_1(\alpha)(\sigma), h_1(\alpha)(\sigma), h_1(\alpha)(\sigma)\big) = d_*(h_1(\alpha))(\sigma)
\end{aligned}$$

Hence $h_2 \circ p^* = d_* \circ h_1$ as claimed. □

6.2.2 Computing the Cohomology of the Torus. We are now ready to apply the technical observation above in our case $X = T^n$ with the diagonal action of $\mathbb{Z}_3$. Note that this action can be made cellular by splitting S^1 into three arcs, and considering the induced product CW-structure on T^n.

For the coefficient module, we will use $\pi_2(P)$ with the structure of Λ-module as described in Lemma 5.8(iii). We denote this module by M. The unique structure of the module M will allow us to use Lemma 6.11. In particular, we have the following lemma.

Lemma 6.12. *The module M is generated by a single element and* $\mathrm{Ann}(M) = I$*, the ideal generated by* $1 + \omega + \omega^2$.

Proof. M is generated over Λ by $m = \binom{1}{0}$. As a consequence, $\lambda \in \mathrm{Ann}(M)$ 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 the action on the torus a little to simplify the calculations:

Lemma 6.13. *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., $\omega\cdot(z_1,\ldots,z_n) = (\omega z_1,\ldots,\omega z_n)$. Let Y be the same torus but with $\mathbb{Z}_3$ acting only on the first coordinate, i.e., $\omega\cdot(z_1,\ldots,z_n) = (\omega z_1, z_2,\ldots,z_n)$. Then there is a $\mathbb{Z}_3$-equivariant homeomorphism $X \to Y$.*

Proof. The maps $h\colon X \to Y$ and $h'\colon Y \to X$ defined by

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

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

$$h(\omega\cdot_X(z_1,\ldots,z_n)) = h(\omega z_1,\ldots,\omega z_n) = \big(\omega z_1, z_1^{-1}z_2,\ldots,z_nz_1^{-1}\big) = \omega\cdot_Y h(z_1,\ldots,z_n)$$ □

Remark 6.14. If we see a torus T^n as a quotient of $\mathbb{R}^n$ over the standard lattice, Lemma 6.13 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 $p^*\colon H^\bullet(T^n/\mathbb{Z}_3;\mathbb{Z}) \to H^\bullet(T^n;\mathbb{Z})$. 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 [18, see Section 3.1] shows 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 $H^1(T^n) \cong \hom(H_1(T^n),\mathbb{Z})$; 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 [18, 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 $x^I = x^{i_1} \smile \cdots \smile x^{i_d}$, where $I = (i_1, \dots, i_d)$ and $i_1 < \cdots < i_d$, 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,

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

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 6.9 which claims that $H^d_{\mathbb{Z}_3}(T^n;\pi_2(P)) \cong \mathbb{Z}_3^{\binom{n-1}{d-1}}$.

Proof of Proposition 6.9. Fix $n \geq 2$. By Lemma 6.12, we have a short exact sequence

$$0 \longrightarrow I \longrightarrow \Lambda \longrightarrow M \longrightarrow 0$$

which induces short exact sequence of cochain complexes

$$0 \longrightarrow C^\bullet_{\mathbb{Z}_3}(T^n;I) \longrightarrow C^\bullet_{\mathbb{Z}_3}(T^n;\Lambda) \longrightarrow C^\bullet_{\mathbb{Z}_3}(T^n;M) \longrightarrow 0$$

Using Lemma 6.11, we get that the following short sequence is also exact

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

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

$$\cdots \longrightarrow H^d(T^n/\mathbb{Z}_3;\mathbb{Z}) \xrightarrow{p_d^*} H^d(T^n;\mathbb{Z}) \longrightarrow H^d_{\mathbb{Z}_3}(T^n;M) \longrightarrow H^{d+1}(T^n/\mathbb{Z}_3;\mathbb{Z}) \xrightarrow{p_{d+1}^*} H^{d+1}(T^n;\mathbb{Z}) \longrightarrow \cdots$$

Since p_d^* is injective for any $d \geq 1$, by exactness we have that $H^d_{\mathbb{Z}_3}(T^n;M) \cong \operatorname{coker} p_d^*$. Furthermore,

$$\operatorname{coker} p_d^* = \mathbb{Z}^{\binom{n}{d}}/\operatorname{im} p_d^* \simeq \mathbb{Z}_3^{\binom{n-1}{d-1}}$$

which yields the desired. □

6.3 The Minion Isomorphism

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

Proof of Lemma 6.1. We show that μ and γ are mutually inverse minion homomorphisms. Recall that μ is injective by Lemma 6.4. We first show that it is surjective. Indeed, we have that, for each $n \geq 1$,

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

by a combination of Proposition 6.9 and Theorem 2.11 since $H^2(T^n, \pi_2(P)) \simeq \mathbb{Z}_3^{n-1}$, and the latter has 3^{n-1} elements.

The above, in particular, means that $\mu_n\colon \mathscr{Z}_3^{(n)} \to \mathrm{hpol}^{(n)}(S^1, P)$ is onto, and hence a bijection. Consequently, γ is a minion homomorphism by Lemma 6.8 since every class in $[T^n, P]_{\mathbb{Z}_3}$ contains a monomial map. Observe that $\mu \circ \gamma$ is the identity map by Lemma 6.7, 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 $\alpha \in \mathscr{Z}_3$,

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

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

Together with Lemma 5.12, the above lemma concludes the proof of Lemma 3.2.

7 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 $\xi\colon\ \mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathscr{Z}_3$ acts on binary polymorphisms. This is the key to the argument that the image of ξ is the projection minion and not the whole of $\mathscr{Z}_3$.

We say that two polymorphisms $f, g \in \mathrm{pol}^{(n)}(\mathbf{LO}_3, \mathbf{LO}_4)$ are *mutually reconfigurable* if a path between f and g exists within the homomorphism complex $\mathrm{Hom}(\mathbf{LO}_3^n, \mathbf{LO}_4)$. (Note that every polymorphism is a homomorphism $\mathbf{LO}_3^n \to \mathbf{LO}_4$, 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 7.1. *Let $\mathbf{A}$ be a symmetric relational structure. If $f, g\colon \mathbf{A} \to \mathbf{LO}_4$ 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 $m\colon A \to 2^{[4]}$ 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 $\mathbf{A} \to \mathbf{LO}_4$, and thus $\mathbf{A}$ has no constant tuples.

Thus m is a multihomomorphism in all cases.

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

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 $f \in \mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$ has, as its domain, the set $[3]^2$, and thus it can naturally be represented as a matrix:

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

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 $f\colon A^n \to B$ is essentially unary if it depends on at most one input coordinate.) The lemma is an analogue of the *Trash Colour Lemma* [5, Lemma 3.4] for polymorphisms from K_d to K_{2d-2}.

Lemma 7.2. *For each $f \in \mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$ there exists an increasing function $h \in \mathrm{pol}^{(1)}(\mathbf{LO}_3, \mathbf{LO}_4)$, 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 $((a,c),(a,c),(b,d)) \in R^{\mathbf{LO}_3^2}$.

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 $((x,y),(x',y'),(x,y')) \in R^{\mathbf{LO}_3^2}$, 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 $((1,1),(x,y),(x',y')) = ((1,1),(1,y),(x',1)) \in R^{\mathbf{LO}_3^2}$ 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. Let us write the matrix of values of f which has been set so far (in the form as above):

$$\begin{array}{ccc} 2 & \star & 1 \\ (1?) & 3 & \\ (1?) & & 4 \end{array}$$

where $(1?)$ marks the possible positions of $f(x', y') = 1$, and the $\star$ denotes the discussed value $f(1, 2)$.

$f(1, 2) \neq 1$ since $((1, 2), (1, 3), (x', y')) \in R^{\mathbf{LO}_3^2}$ is mapped by f to $(f(1, 2), 1, 1)$.
$f(1, 2) \neq 2$ since $((1, 2), (1, 1), (x', y')) \in R^{\mathbf{LO}_3^2}$ is mapped by f to $(f(1, 2), 2, 1)$.
$f(1, 2) \neq 3$ since $((1, 2), (2, 2), (1, 3)) \in R^{\mathbf{LO}_3^2}$ is mapped by f to $(f(1, 2), 3, 1)$.
$f(1, 2) \neq 4$ since $f(1, 2) < f(3, 3) = 4$.

$x' = 3$. Here the contradiction follows analogously to the previous case.

$x' = y = 2$. We consider the pair of values $f(1, 3)$ and $f(3, 1)$. Again, the matrix of values of f which has been set so far is

$$\begin{array}{ccc} 2 & 1 & \star \\ 1 & 3 & \\ \star & & 4 \end{array}$$

where $\star$'s denote the discussed values. First, we claim that $f(1, 3) = 4$ — the other values are ruled out by tuples $((1, 2), (2, 1), (1, 3)) \in R^{\mathbf{LO}_3^2}$, $((1, 1), (2, 1), (1, 3)) \in R^{\mathbf{LO}_3^2}$, and $((1, 1), (2, 2), (1, 3)) \in R^{\mathbf{LO}_3^2}$, respectively. Symmetrically, we get $f(3, 1) = 4$. However, then $(f(1, 1), f(1, 3), f(3, 1)) = (2, 4, 4) \notin R^{\mathbf{LO}_4}$, which is not possible, as $((1, 1), (1, 3), (3, 1)) \in R^{\mathbf{LO}_3^2}$, 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 *row* type if $f(x, y) = c$ implies $x = a_c$ for some fixed $a_c \in [3]$, and is of *column* 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, which implies $h \in \mathrm{pol}^{(1)}(\mathbf{LO}_3, \mathbf{LO}_4)$. This is since there are $y < y'$ with $y \neq b_t$ and $y' \neq b_t$, and consequently

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

This concludes the proof of the lemma. □

Lemma 7.3. *Every binary polymorphism $f \in \mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$ is reconfigurable to an essentially unary polymorphism.*

Proof. The proof relies on Lemma 7.2. 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 $h \in \mathrm{pol}^{(1)}(\mathbf{LO}_3, \mathbf{LO}_4)$, and that in particular $f(x_0, y_0) = t$. Furthermore, suppose that among all such pairs, (x_0, y_0) is the one that maximises x_0. We claim that $f'(x, y)$, which is equal to $f(x, y)$ everywhere except that $f'(x_0, y_0) = h(x_0)$ is also a polymorphism, which gives us our inductive step.

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

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

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

In Figure 4, we can see the reconfiguration graph of $\mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$, which is in fact the 1-dimensional skeleton of the Hom complex $\mathrm{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$. 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 $\mathrm{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$ contains a homomorphism, and hence a unary one, we can derive from these observation that $\mathrm{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$ has at most two connected components. Figure 4 shows that there are exactly two of them.

We have also shown in Section 6.1 that $\mathrm{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$ has at least two connected components: The two projections give distinct monomial maps (see Definition 6.2) and, by Lemma 6.4, they are not equivariantly homotopic.

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

Lemma 7.4. *Let* $\mathbf{A}$, $\mathbf{B}$, *and* $\mathbf{C}$ *be three structures,* G *a group acting on* $\mathbf{A}$, *and assume that* $f, g \in \hom(\mathbf{B}, \mathbf{C})$ *are mutually reconfigurable. Then the induced maps* $f_*, g_* \colon \mathrm{Hom}(\mathbf{A}, \mathbf{B}) \to \mathrm{Hom}(\mathbf{A}, \mathbf{C})$ *are* G*-homotopic.*

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

$$c\colon \mathrm{Hom}(\mathbf{B}, \mathbf{C}) \times \mathrm{Hom}(\mathbf{A}, \mathbf{B}) \to \mathrm{Hom}(\mathbf{A}, \mathbf{C})$$

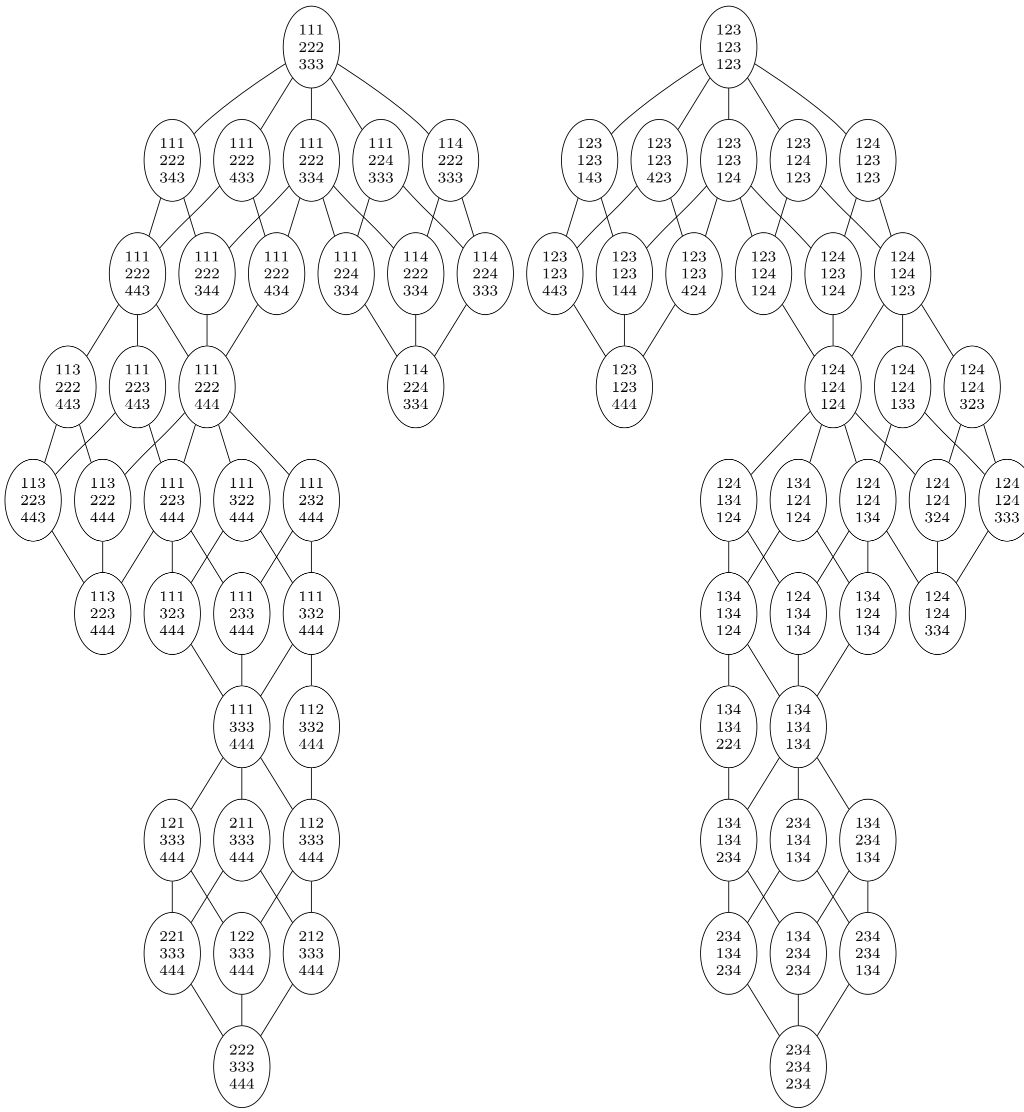

Fig. 4. Graph of reconfigurations of $\mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$.

(see also [23, 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 $h\colon [0,1] \to \mathrm{Hom}(\mathbf{B}, \mathbf{C})$ is an arc connecting f and g, i.e., such that $h(0) = f$ and $h(1) = g$, then the map $H\colon [0,1] \times \mathrm{Hom}(\mathbf{A}, \mathbf{B}) \to \mathrm{Hom}(\mathbf{A}, \mathbf{C})$ 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. $\square$

The following corollary then follows directly from the above and Lemma 7.3.

Corollary 7.5. *For every binary polymorphism $f \in \mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$, the induced map*

$$f_*\colon \operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_3)^2 \longrightarrow \operatorname{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$$

is equivariantly homotopic either to the map $(x, y) \mapsto i_*(x)$, or to the map $(x, y) \mapsto i_*(y)$ where $i\colon \mathbf{LO}_3 \to \mathbf{LO}_4$ is the inclusion.

This concludes the proof of Lemma 3.3 and the main theorem.

References

[1] Per Austrin, Venkatesan Guruswami, and Johan Håstad. 2017. $(2+\varepsilon)$-Sat is NP-hard. *SIAM J. Comput.* 46, 5 (2017), 1554–1573. DOI: https://doi.org/10.1137/15M1006507

[2] Libor Barto, Diego Battistelli, and Kevin M. Berg. 2021. Symmetric promise constraint satisfaction problems: Beyond the boolean case. In *Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021* . Markus Bläser and Benjamin Monmege (Eds.), LIPIcs, Vol. 187, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 10:1–10:16. DOI: https://doi.org/10.4230/LIPIcs.STACS.2021.10

[3] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. 2021. Algebraic approach to promise constraint satisfaction. *J. ACM* 68, 4 (2021), 28:1–66. DOI: https://doi.org/10.1145/3457606

[4] Avrim Blum. 1994. New approximation algorithms for graph coloring. *J. ACM* 41, 3 (1994), 470–516. DOI: https://doi.org/10.1145/176584.176586

[5] Joshua Brakensiek and Venkatesan Guruswami. 2016. New hardness results for graph and hypergraph colorings. In *Proceedings of the Conference on Computational Complexity (CCC 2016)* . Ran Raz (Ed.), LIPIcs, Vol. 50, Schloss Dagstuhl–LZI, Dagstuhl, Germany, 14:1–14:27. DOI: https://doi.org/10.4230/LIPIcs.CCC.2016.14

[6] Joshua Brakensiek and Venkatesan Guruswami. 2021. Promise constraint satisfaction: Algebraic structure and a symmetric boolean dichotomy. *SIAM J. Comput.* 50, 6 (2021), 1663–1700. DOI: https://doi.org/10.1137/19M128212X

[7] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. 2022. An invariance principle for the multi-slice, with applications. In *Proceedings of the 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021.* IEEE Computer Soc., Los Alamitos, CA, 228–236. DOI: https://doi.org/10.1109/FOCS52979.2021.00030

[8] Glen E. Bredon. 1967. *Equivariant Cohomology Theories.* Lecture Notes in Mathematics, Vol. 34. Springer Berlin. DOI: https://doi.org/10.1007/BFb0082690

[9] Andrei A. Bulatov. 2017. A dichotomy theorem for nonuniform CSPs. In *Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS).* IEEE, Berkeley, CA, USA, 319–330. DOI: https://doi.org/10.1109/FOCS.2017.37

[10] Panagiotis Cheilaris, Balázs Keszegh, and Dömötör Pálvölgyi. 2013. Unique-maximum and conflict-free coloring for hypergraphs and tree graphs. *SIAM J. Discret. Math.* 27, 4 (2013), 1775–1787. DOI: https://doi.org/10.1137/120880471

[11] Lorenzo Ciardo, Marcin Kozik, Andrei A. Krokhin, Tamio-Vesa Nakajima, and Stanislav Živný. 2025. 1-in-3 vs. Not-All-Equal: Dichotomy of a broken promise. *ACM Trans. Comput. Log.* 26, 2 (2025), 10:1–10:22. DOI: 10.1145/3719007arXiv:2302.03456

[12] Víctor Dalmau and Jakub Opršal. 2024. Local consistency as a reduction between constraint satisfaction problems. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2024 (Tallinn, Estonia, July 8-11, 2024),* Pawel Sobocinski, Ugo Dal Lago, and Javier Esparza (Eds.). ACM, 29:1–29:15. DOI: 10.1145/3661814.3662068

[13] Irit Dinur, Elchanan Mossel, and Oded Regev. 2009. Conditional hardness for approximate coloring. *SIAM J. Comput.* 39, 3 (2009), 843–873. DOI: https://doi.org/10.1137/07068062X

[14] Irit Dinur, Oded Regev, and Clifford Smyth. 2005. The hardness of 3-uniform hypergraph coloring. *Combinatorica* 25, 5 (2005), 519–535. DOI: https://doi.org/10.1007/s00493-005-0032-4

[15] Venkatesan Guruswami and Sai Sandeep. 2020. d-to-1 hardness of coloring 3-colorable graphs with $O(1)$ colors. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming.* LIPIcs. Leibniz Int. Proc. Inform., Vol. 168. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, Art. No. 62, 12.

[16] Magnús M. Halldórsson. 1993. A still better performance guarantee for approximate graph coloring. *Inform. Process. Lett.* 45, 1 (1993), 19–23. DOI: https://doi.org/10.1016/0020-0190(93)90246-6

[17] Johan Håstad, Björn Martinsson, Tamio-Vesa Nakajima, and Stanislav Živný. 2024. A logarithmic approximation of linearly-ordered colourings. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2024* . Amit Kumar and Noga Ron-Zewi (Eds.), LIPIcs, Vol. 317, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:6. DOI: https://doi.org/10.4230/LIPICS.APPROX/RANDOM.2024.7

[18] Allen Hatcher. 2002. *Algebraic Topology.* Cambridge University Press, Cambridge. xii+544 pages.

[19] Richard M. Karp. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations: Proceedings of a Symposium on the Complexity of Computer Computations.* Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (Eds.), Springer US, Boston, MA, 85–103. DOI: https://doi.org/10.1007/978-1-4684-2001-2_9

[20] Ken-ichi Kawarabayashi and Mikkel Thorup. 2017. Coloring 3-colorable graphs with less than $n^{1/5}$ colors. *J. ACM* 64, 1 (2017), 4:1–4:23. DOI: https://doi.org/10.1145/3001582

[21] Ken-ichi Kawarabayashi, Mikkel Thorup, and Hirotaka Yoneda. 2024. Better coloring of 3-colorable graphs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024* . Bojan Mohar, Igor Shinkar, and Ryan O'Donnell (Eds.), ACM, 331–339. DOI: https://doi.org/10.1145/3618260.3649768

[22] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. 2000. On the hardness of approximating the chromatic number. *Combinatorica* 20, 3 (2000), 393–415. DOI: https://doi.org/10.1007/s004930070013

[23] Dmitry Kozlov. 2008. *Combinatorial Algebraic Topology.* Algorithms and Computation in Mathematics, Vol. 21. Springer. XX, 390 pages. DOI: https://doi.org/10.1007/978-3-540-71962-5

[24] Andrei Krokhin and Jakub Opršal. 2022. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News* 9, 3 (2022), 30–59. arXiv:2208.13538 DOI: https://doi.org/10.1145/3559736.3559740

[25] Andrei Krokhin and Danny Vagnozzi. 2025. Approximating 1-in-3 SAT by linearly ordered hypergraph 3-colouring is NP-hard. arXiv:2508.14606. Retrieved from https://arxiv.org/abs/2508.14606. DOI: https://doi.org/10.48550/ARXIV.2508.14606

[26] Andrei A. Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. 2023. Topology and adjunction in promise constraint satisfaction. *SIAM J. Comput.* 52, 1 (2023), 38–79. arXiv:2003.11351 DOI: https://doi.org/10.1137/20M1378223

[27] Jiří Matoušek. 2003. *Using the Borsuk-Ulam Theorem* (1st ed.). Springer-Verlag Berlin. xii+214 pages. DOI: https://doi.org/10.1007/978-3-540-76649-0

[28] J. Peter May, M. Cole, G. Comezana, S. Costenoble, A. Elmendorf, J. Greenlees, L. Lewis, R. Piacenza, G. Triantafillou, and S. Waner. 1996. *Equivariant Homotopy and Cohomology Theory.* American Mathematical Society. DOI: https://doi.org/10.1090/cbms/091

[29] Tamio-Vesa Nakajima and Stanislav Živný. 2023. Linearly ordered colourings of hypergraphs. *ACM Trans. Comput. Theory* 14, 3–4, Article 12 (Feb. 2023), 19 pages. DOI: https://doi.org/10.1145/3570909

[30] Tamio-Vesa Nakajima and Stanislav Živný. 2023. Boolean symmetric vs. functional PCSP dichotomy. In *Proceedings of the 2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS).* 1–12. DOI: https://doi.org/10.1109/LICS56636.2023.10175746

[31] Jean-Pierre Serre. 1953. Groupes D'Homotopie Et Classes De Groupes Abeliens. *Ann. Math.* 58, 2 (1953), 258–294. Retrieved from http://www.jstor.org/stable/1969789

[32] Tammo tom Dieck. 1987. *Transformation Groups.* De Gruyter, Berlin, New York. DOI: https://doi.org/10.1515/9783110858372

[33] Marcin Wrochna. 2020. Homomorphism reconfiguration via homotopy. *SIAM J. Discret. Math.* 34, 1 (2020), 328–350. DOI: https://doi.org/10.1137/17M1122578

[34] Marcin Wrochna and Stanislav Živný. 2019. Improved hardness for H-colourings of G-colourable graphs. arXiv:1907.00872v3. Retrieved from https://arxiv.org/abs/1907.00872. DOI: https://doi.org/10.48550/arXiv.1907.00872 full version of SODA'20 paper.

[35] Dmitriy Zhuk. 2020. A proof of the CSP dichotomy conjecture. *J. ACM* 67, 5 (2020), 30:1–30:78. DOI: https://doi.org/10.1145/3402029

[36] David Zuckerman. 2007. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.* 3 (2007), 103–128. DOI: https://doi.org/10.4086/toc.2007.v003a006

Received 11 March 2025; revised 8 September 2025; accepted 11 November 2025