



Hardness of 4-Colouring G -Colourable Graphs

Sergey Avvakumov
savvakumov@tauex.tau.ac.il
Tel Aviv University
Tel Aviv, Israel

Marek Filakovský
filakovsky@fi.muni.cz
Masaryk University
Brno, Czech Republic

Jakub Opršal
j.oprsal@bham.ac.uk
University of Birmingham
Birmingham, UK

Gianluca Tasinato
gianluca.tasinato@ist.ac.at
ISTA
Klosterneuburg, Austria

Uli Wagner
uli@ist.ac.at
ISTA
Klosterneuburg, Austria

Abstract

We study the complexity of a class of promise graph homomorphism problems. For a fixed graph H , the H -colouring problem is to decide whether a given graph has a homomorphism to H . By a result of Hell and Nešetřil, this problem is NP-hard for any non-bipartite loop-less graph H . Brakensiek and Guruswami [SODA 2018] conjectured the hardness extends to promise graph homomorphism problems as follows: fix a pair of non-bipartite loop-less graphs G, H such that there is a homomorphism from G to H , it is NP-hard to distinguish between graphs that are G -colourable and those that are not H -colourable. We confirm this conjecture in the cases when both G and H are 4-colourable. This is a common generalisation of previous results of Khanna, Linial, and Safra [Comb. 20(3): 393-415 (2000)] and of Krokhn and Opršal [FOCS 2019]. The result is obtained by combining the algebraic approach to promise constraint satisfaction with methods of topological combinatorics and equivariant obstruction theory.

CCS Concepts

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

Keywords

graph colouring, constraint satisfaction problems, promise problems, topological methods, equivariant cohomology

ACM Reference Format:

Sergey Avvakumov, Marek Filakovský, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. 2025. Hardness of 4-Colouring G -Colourable Graphs. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25)*, June 23–27, 2025, Prague, Czechia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3717823.3718154>

Acknowledgments

This research was supported by the Austrian Science Fund (FWF project P31312-N35) and by project MSCA fellow5_MUNI (CZ.02.01.01/00/22_010/0003229) financed by the Ministry of Education, Youth

and Sports of the Czech Republic. This project has also received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 101034413.

1 Introduction

Deciding whether a given finite graph is 3-colourable (or, in general, k -colourable, for a fixed $k \geq 3$) was one of the first problems shown to be NP-complete (Karp [17]). Since then, the complexity of *approximating* the chromatic number of a graph has been studied extensively; in particular, it is known that the chromatic number of an n -vertex graph cannot be approximated in polynomial time within a factor of $n^{1-\varepsilon}$, for any fixed $\varepsilon > 0$, unless $P = NP$ (Zuckerman [34]).

However, this inapproximability result only applies to graphs whose chromatic number grows with the number of vertices; by contrast, the case of graphs with bounded chromatic number is much less well understood. For instance, given an input graph that is promised to be 3-colourable, what is the complexity of finding a colouring of G with some larger number $k > 3$ of colours? Khanna, Linial, and Safra [19] proved that this is NP-hard for $k = 4$ (see also [3; 13]), and only quite recently Bulín, Krokhn, and Opršal [8] showed NP-hardness for $k = 5$. On the other hand, the currently best polynomial-time algorithm for colouring 3-colourable graphs, due to Kawarabayashi, Thorup, and Yoneda [18], uses $k = \tilde{O}(n^{0.19747})$ colours, where n is the number of vertices of the input graph.

In general, it is believed that colouring c -colourable graphs with k colours is NP-hard for all constants $k \geq c \geq 3$. However, the best results known to date (apart from the above) are NP-hardness for $c = 4$ and $k = 7$ (Bulín, Krokhn, and Opršal [8]), and for $c \geq 5$ and $k = \binom{c}{\lfloor c/2 \rfloor} - 1$ (Wrochna and Živný [32]). Moreover, conditional hardness results — assuming different variants of Khot's *Unique Games Conjecture* — have been obtained for all $k \geq c \geq 3$ by Dinur, Mossel, and Regev [10], Guruswami and Sandeep [14], and Braverman, Khot, Lifshitz, and Minzer [5].

In the present paper, we study a generalisation of this question. A *graph homomorphism* $f: G \rightarrow H$ between two graphs is a map $f: V(G) \rightarrow V(H)$ between the vertex sets that preserves edges, i.e., $(u, v) \in E(G)$ implies $(f(u), f(v)) \in E(H)$; we write $G \rightarrow H$ if such a homomorphism exists. Throughout this paper, we assume all graphs to be finite and undirected and we treat them as symmetric binary relational structures, i.e., we view the edge set $E(G)$ as a



This work is licensed under a Creative Commons Attribution 4.0 International License. STOC '25, Prague, Czechia

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1510-5/25/06
<https://doi.org/10.1145/3717823.3718154>

subset of $V(G) \times V(G)$ that satisfies $(u, v) \in E(G)$ if and only if $(v, u) \in E(G)$, and we allow loops, i.e., edges of the form (v, v) . A graph homomorphism $f: G \rightarrow H$ is also called an H -colouring of G since a k -colouring of G is the same as a homomorphism from G to the complete (loopless) graph K_k on k vertices. Vastly generalising the fact that k -colouring is NP-hard if $k \geq 3$, and in P if $k \leq 2$, Hell and Nešetřil [16] proved the following dichotomy: For every fixed graph H , deciding whether a given input graph admits an H -colouring is NP-complete, unless H is bipartite or has a loop, in which case the problem is in P. Analogously to approximate graph colouring, it is natural to consider the complexity of the following *promise graph homomorphism problem*: Fix two graphs G and H such that $G \rightarrow H$. What is the complexity of H -colouring graphs that are promised to be G -colourable? More precisely, we consider the decision version of this problem, denoted by $\text{PCSP}(G, H)$: Given an input graph I , output YES if $I \rightarrow G$ and NO if $I \not\rightarrow H$ (no output is required if neither is the case). Brakensiek and Guruswami conjectured [4, Conjecture 1.2] that $\text{PCSP}(G, H)$ is NP-hard for all non-bipartite, loopless graphs G and H (i.e., unless the problem is guaranteed to lie in P by the Hell–Nešetřil dichotomy). This problem fits into the much broader framework of *promise constraint satisfaction problems* (PCSPs), from which the notation is adopted; see the survey [22] for more background.

As a first step towards the Brakensiek–Guruswami conjecture, Krokhn and Opršal [21] showed that $\text{PCSP}(G, K_3)$ is NP-hard for every 3-colourable non-bipartite graph G . Their proof was based on ideas from algebraic topology; this topological intuition was formalised by Wrochna and Živný [32] (and in the joint journal version [23]). We extend this to 4-colouring:

THEOREM 1.1. *Let G be a non-bipartite 4-colourable graph. Then $\text{PCSP}(G, K_4)$ is NP-hard.*

The proof of Theorem 1.1, whose structure we present in detail in Section 3 below, builds on and significantly extends the topological approach used in [23], bringing to bear more powerful tools from algebraic topology as well as more refined combinatorial arguments. By a simple reduction, Theorem 1.1 reduces to the special case where $G = C_\ell$ is a cycle of arbitrary odd length $\ell \geq 3$ (Theorem 3.1).

We rely on a general algebraic theory of PCSPs by Barto, Bulín, Krokhn, and Opršal [2], which guarantees that the complexity of the problem only depends on its *polymorphisms*. In our case, a polymorphism is a graph homomorphism $f: C_\ell^n \rightarrow K_4$ where n is a natural number and $C_\ell^n = C_\ell \times \dots \times C_\ell$ is the n -fold power of C_ℓ (see Section 2.1).

The proof of Theorem 3.1 has two parts. In the first part of the proof, we use topological methods (*homomorphism complexes* and *equivariant obstruction theory*) to show that with every polymorphism f , we can associate a map $\phi(f): \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ of the form $\phi(f)(x_1, \dots, x_n) = \sum_{i=1}^n \alpha_i x_i$, for some $\alpha_i \in \mathbb{Z}_2$ such that $\sum_{i=1}^n \alpha_i$ is odd (Lemma 3.3). Moreover, the map $f \mapsto \phi(f)$ preserves natural *minor relations* between polymorphisms that arise from substituting and permuting variables; in technical terms, ϕ is a *minion homomorphism* (see Definition 2.2). The second part of the proof uses combinatorial arguments to show (Theorem 5.1) that the affine maps $\phi(f): \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ arising from polymorphisms are of *bounded essential arity*: the number of non-zero coefficients $\alpha_i \in \mathbb{Z}_2$ describing $\phi(f)$ is at most $O(\ell^2)$, independently of n . Hardness of

$\text{PCSP}(C_\ell, K_4)$ then follows from a criterion (Theorem 2.3) obtained as part of the general algebraic theory developed in [2].

Related work. Graph colouring and H -colouring are examples of *constraint satisfaction problems* (CSPs), a general framework that encompasses many other fundamental problems including 3SAT, HornSAT, solving systems of linear equations, and linear programming. CSPs can be formulated in several equivalent ways; the one most relevant for us is in terms of homomorphisms between relational structures: fix a relational structure A (e.g., a graph, a digraph, or a uniform hypergraph), the CSP with template A , denoted by $\text{CSP}(A)$, is the problem of deciding whether a given input structure I allows a homomorphism $I \rightarrow A$. Thus, H -colouring is the same as $\text{CSP}(H)$. A key result in the complexity theory of CSPs (which is a culmination of decades of research and subsumes various previous results, including the Hell–Nešetřil dichotomy and an earlier one by Schaefer [29] for Boolean CSPs) is a general Dichotomy Theorem of Bulatov [7] and Zhuk [33], which asserts that for every finite relational structure A , $\text{CSP}(A)$ is either NP-complete, or solvable in polynomial time. *Promise constraint satisfaction problems* (PCSPs) are a natural extension of CSPs, analogous to how the promise graph homomorphism problem extends the H -colouring problem. This notion of PCSPs was introduced by Austrin, Guruswami, and Håstad [1], and the general theory was further developed by Brakensiek and Guruswami [4], and by Barto, Bulín, Krokhn, and Opršal [2].

In addition to the aforementioned work [21; 23; 32], several other recent papers explore the connection between topology and the computational complexity of (P)CSPs. Schnider and Weber [30] showed that Schaefer’s dichotomy for the complexity of Boolean CSPs is reflected by the “topological complexity” of their solution spaces: the solution spaces of NP-complete problems are topologically arbitrarily complicated (in a precise technical sense), whereas the solution spaces of polynomial-time solvable problems are homotopy equivalent to a discrete set; this was further generalised to arbitrary CSPs by Meyer [27]. Meyer and Opršal [28] gave a new, topological proof of the Hell–Nešetřil dichotomy, and Filakovský, Nakajima, Opršal, Tasinato, and Wagner [11] use topological methods related to the ones in the present paper to show that a certain hypergraph PCSP is NP-hard. We believe that these results, and the ones presented here, are just the starting point of a promising line of research and that topological methods have the potential to yield further complexity-theoretic insights in the future.

Acknowledgement. We are grateful to Andrei Krokhn, Marcin Wrochna, Standa Živný, and Libor Barto for valuable discussions during the early stages of this project. In particular, we would like to thank Marcin Wrochna and Libor Barto for sharing with us key observations that played an important role in the construction of a minion homomorphism to affine \mathbb{Z}_2 -maps.

2 Preliminaries

We denote the identity function on a set A by 1_A , we use the notation $[n] = \{1, \dots, n\}$, and the symbol \sqcup for disjoint union.

2.1 Polymorphisms and a Hardness Criterion

We outline the fundamentals of the algebraic theory of PCSPs, in particular the core concept of *polymorphisms* and the hardness

criterion (Theorem 2.3) used in the proof of Theorem 1.1, focusing on special the case of graphs; see a survey by Krokhin and Opršal [22] for a detailed treatment.

Given two graphs G_1 and G_2 , their product $G_1 \times G_2$ is defined by $V(G_1 \times G_2) = V(G_1) \times V(G_2)$ and $((u_1, u_2), (v_1, v_2)) \in E(G_1 \times G_2)$ if and only if $(u_1, v_1) \in E(G_1)$ and $(u_2, v_2) \in E(G_2)$; moreover, we denote by $G^n = G \times \dots \times G$ the product of n copies of G .

Definition 2.1. An n -ary *polymorphism* from a graph G to a graph H is homomorphism $f: G^n \rightarrow H$, in other words, a map $f: V(G)^n \rightarrow V(H)$ such that $(f(u_1, \dots, u_n), f(v_1, \dots, v_n)) \in E(H)$ whenever $(u_1, v_1), \dots, (u_n, v_n) \in E(G)$. We denote¹ the set of all polymorphisms from G to H by $\text{pol}(G, H)$, and the set of n -ary polymorphisms by $\text{pol}^{(n)}(G, H)$.

Polymorphisms are enough to describe the complexity of a promise CSP up to certain log-space reductions [22, Theorem 2.20]. Loosely speaking, the more complex the polymorphisms are, the easier the problem is. To formalise this, we define the notions of *minor*, *minion*, *minion homomorphism*, and *essential arity* which are necessary to formulate the hardness criterion.

Let $\pi: [n] \rightarrow [m]$, and let A and B be sets. The π -*minor* of a function $f: A^n \rightarrow B$ is the function $f^\pi: A^m \rightarrow B$ given by $f^\pi(x_1, \dots, x_m) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$ for all $x_1, \dots, x_m \in A$ (equivalently, if we view elements of $x \in A^n$ as functions $x: [n] \rightarrow A$, then $f^\pi(x) = f(x \circ \pi)$). A subset of the set of all functions $\{f: A^n \rightarrow B, n > 0\}$ that is non-empty and closed under taking minors is called a *function minion*. For example, it is easy to see that $\text{pol}(G, H)$ has this property whenever G and H are graphs such that $G \rightarrow H$. Abstracting from this, we arrive at the following notion:

Definition 2.2. An (abstract) *minion* \mathcal{M} is a collection of non-empty sets $\mathcal{M}^{(n)}$, where $n > 0$ is an integer, and mappings

$$\pi^{\mathcal{M}}: \mathcal{M}^{(n)} \rightarrow \mathcal{M}^{(m)},$$

for $\pi: [n] \rightarrow [m]$, which satisfy $\pi^{\mathcal{M}} \circ \sigma^{\mathcal{M}} = (\pi \circ \sigma)^{\mathcal{M}}$ whenever $\pi \circ \sigma$ is defined, and $(1_{[n]})^{\mathcal{M}} = 1_{\mathcal{M}^{(n)}}$. We will often write f^π instead of $\pi^{\mathcal{M}}(f)$, and call this element the π -minor of f .

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

Given a minion \mathcal{M} , an element $f \in \mathcal{M}^{(n)}$ is said to have *essential arity at most k* if it is a minor of an element $g \in \mathcal{M}^{(k)}$. If there is a bound N , such that every element of \mathcal{M} has essential arity at most N , \mathcal{M} is said to have *bounded essential arity*. An element $f \in \mathcal{M}^{(n)}$ is *constant* if all its minors coincide, i.e., $f^\pi = f^\sigma$ for all $m > 0$ and $\pi, \sigma: [n] \rightarrow [m]$. For example, in function minions, being constant coincides with the usual notion of being a constant function, and if a function $f: A^n \rightarrow B$ depends only on a subset of variables with indices $\{i_1, \dots, i_k\}$, then $f(x_1, \dots, x_n) = g(x_{i_1}, \dots, x_{i_k})$, so f is of arity at most k . Our proof uses the following hardness criterion.

¹ Somewhat unconventionally, we use lower-case notation for polymorphisms to highlight that we are not considering any topology on them, in contrast to the homomorphism complexes introduced below.

THEOREM 2.3 ([2, PROPOSITION 5.14]). *Let G and H be two graphs such that $G \rightarrow H$. If there exists a minion homomorphism*

$$\xi: \text{pol}(G, H) \rightarrow \mathcal{B}$$

for some minion \mathcal{B} of bounded essential arity which does not contain a constant, then $\text{PCSP}(G, H)$ is NP-complete.

2.2 Topology and Homomorphism Complexes

We review a number of topological notions that we will need in what follows, in particular the notion of *homomorphism complexes*, a well-known construction in topological combinatorics that goes back to the work of Lovász [24]. We refer the reader to Hatcher [15] and Matoušek [25] for accessible general introductions to algebraic topology and topological combinatorics, respectively, and to Kozlov [20] for an in-depth treatment of homomorphism complexes.

Simplicial sets. In applications of topological methods in combinatorics and theoretical computer science, topological spaces are often specified combinatorially as simplicial complexes. For our purposes, it will be convenient to work instead with *simplicial sets*, which generalize simplicial complexes in a way analogous to how directed multigraphs generalize simple graphs. Simplicial sets are a somewhat less common notion in topological combinatorics, but play an important role in homotopy theory, see Friedman [12] for a gentle combinatorial introduction.

Similarly to a simplicial complex, a simplicial set is a combinatorial description of how to build a space from vertices, edges, triangles, and higher-dimensional simplices. Informally speaking, we view the vertex set of each n -dimensional simplex as totally ordered (equivalently, labelled by $\{0, 1, \dots, n\}$) and we are allowed to glue simplices together by linear maps between them that are given by (not necessarily strictly) monotone maps between their vertex sets. On the one hand, this permits more general glueings than in simplicial complexes (which allows constructing spaces using fewer simplices): for instance, we may glue both endpoints of an edge to the same vertex (creating a loop), or glue the endpoints of multiple edges to the same pair of vertices, or we may glue the the boundary of a triangle to a single vertex, forming a 2-dimensional sphere S^2 . On the other hand, the description is still purely combinatorial and, moreover, retains the information about the ordering of the vertices of each simplex before the glueing. This yields a natural notion of products of simplicial sets and will play an important role in the combinatorial arguments below.

Definition 2.4. A *simplicial set* X is given by the following data: First, a collection of pairwise disjoint sets X_0, X_1, X_2, \dots ; the elements of X_n are called the n -*simplices* of X . Second, for every pair of integers $m, n \geq 0$ and every (not necessarily strictly) monotone map $\alpha: \{0, \dots, m\} \rightarrow \{0, \dots, n\}$, there is a map $\alpha^X: X_n \rightarrow X_m$, such that $1_{\{0, \dots, n\}}^X = 1_{X_n}$ and such that $(\alpha \circ \beta)^X = \beta^X \circ \alpha^X$ whenever the composition is defined.

Every simplicial set X defines a topological space $|X|$, the *geometric realization* of X , which is obtained by glueing geometric simplices together according to the combinatorial data in X ; we refer to [12, Section 4] for a precise definition. We say that a simplicial set X is a *triangulation* of a topological space T if $|X|$ is homeomorphic to T . A k -simplex $\sigma \in X_k$ is called *degenerate* if $\sigma = \alpha^X(\tau)$

for some $\tau \in X_m$ and $\alpha: \{0, \dots, k\} \rightarrow \{0, \dots, m\}$ with $m < k$. In the geometric realization, degenerate simplices are collapsed down to lower-dimensional simplices, and $|X|$ is the disjoint union of the interiors of non-degenerate simplices;² however, degenerate simplices play an important role in specifying the glueings and the combinatorial data keeps track of them. All simplicial sets used in this paper have only finitely many non-degenerate simplices; this is equivalent to $|X|$ being a compact space. The *dimension* of a simplicial set X is defined as the maximum dimension of a non-degenerate simplex of X .

A *simplicial map* $f: X \rightarrow Y$ between simplicial sets is a collection of maps $f_n: X_n \rightarrow Y_n$, $n \geq 0$, such that $f_m \circ \alpha^X = \alpha^Y \circ f_n$ for all monotone maps $\alpha: \{0, \dots, m\} \rightarrow \{0, \dots, n\}$. Every simplicial map $f: X \rightarrow Y$ defines a continuous map $|f|: |X| \rightarrow |Y|$. An *isomorphism* of simplicial sets X and Y is a simplicial map $f: X \rightarrow Y$ with a simplicial inverse $g: Y \rightarrow X$ (f_n is inverse to g_n for all $n \geq 0$).

Products. The product $X \times Y$ of two simplicial sets X and Y is the simplicial set whose n -simplices of $X \times Y$ are ordered pairs (σ, τ) , i.e., $(X \times Y)_n = X_n \times Y_n$, and $\alpha^{X \times Y}(\sigma, \tau) = (\alpha^X(\sigma), \alpha^Y(\tau))$. On the level of geometric realizations, this corresponds to the usual product of topological spaces, i.e., $|X \times Y| \cong |X| \times |Y|$, under some mild conditions on X and Y that are satisfied for all simplicial sets we work with (e.g., if both X and Y are countable, see [12, Theorem 5.2] for a general statement). The n th power of a simplicial set X is $X^n = X \times \dots \times X$ (the product of n copies of X).

Group actions. Various objects we work with in this paper (graphs, simplicial sets, topological spaces, etc.) have a natural symmetry given by an action of the cyclic group \mathbb{Z}_2 , which is described by a structure-preserving involution. For instance, a \mathbb{Z}_2 -action on a simplicial set X is given by a simplicial map $\nu: X \rightarrow X$ that satisfies $\nu^2 := \nu \circ \nu = 1_X$ (thus, ν is necessarily a simplicial automorphism). We mainly work with actions that are *free*, which for \mathbb{Z}_2 -actions simply means that ν has no fixed points. If (X, ν_X) and (Y, ν_Y) are simplicial sets with \mathbb{Z}_2 -actions, then a simplicial map $f: X \rightarrow Y$ is called *equivariant* if it preserves the \mathbb{Z}_2 -symmetry, i.e., $f \circ \nu_X = \nu_Y \circ f$.

\mathbb{Z}_2 -actions on graphs (by isomorphisms) or on spaces (by homeomorphisms), and the notions of equivariant graph homomorphisms and equivariant continuous maps, etc., are defined analogously.

Relational simplicial sets. Most simplicial sets in this paper are of the following special form, which we call *relational* (a non-standard term): The set X_0 of vertices (0-simplices) is a finite set, and $X_n \subseteq (X_0)^{n+1}$ is an $(n+1)$ -ary relation, i.e., every n -simplex of X is an ordered $(n+1)$ -tuple $[u_0, \dots, u_n]$ of vertices (we use square brackets as a reminder that we view these $(n+1)$ -tuples as simplices, and we identify each element $u \in X_0$ with the singleton tuple $[u]$). Moreover, for every monotone map $\alpha: \{0, \dots, m\} \rightarrow \{0, \dots, n\}$, the map α^X is defined by $\alpha^X([u_0, \dots, u_n]) = [u_{\alpha(0)}, \dots, u_{\alpha(m)}]$. To get a simplicial set this way, the collection of relations X_n , $n \geq 0$, needs to be closed under the operations α^X , i.e., if σ is a simplex of X , then any tuple obtained from σ by omitting and/or repeating vertices without changing their order is a simplex as well.

²In more technical terms, $|X|$ is a CW complex with one k -cell for each non-degenerate k -simplex of X .




Figure 1: The simplicial set Σ^2

Example 2.5 (\mathbb{Z}_2 -symmetric triangulations of spheres). We define a relational simplicial set Σ^2 that defines a triangulation of the 2-dimensional sphere S^2 , together with a natural \mathbb{Z}_2 -action that corresponds to the antipodal map $x \mapsto -x$ on S^2 . The vertex set of Σ^2 is $\Sigma_0^2 = \{\bullet, \circ\}$ (which we think of as a pair of antipodal points in S^2), and Σ_n^2 is the set of all $(n+1)$ -tuples of \bullet and \circ 's with at most 2 alternations. Thus, e.g., $[\bullet, \circ, \bullet, \circ]$ is a 3-simplex of Σ^2 , but $[\bullet, \circ, \bullet, \bullet]$ is not. The \mathbb{Z}_2 -action on Σ^2 is given by the simplicial map that swaps the two vertices.

This construction naturally generalises to yield a sequence of simplicial sets $\Sigma^0 \subseteq \Sigma^1 \subseteq \Sigma^2 \subseteq \dots$, such that Σ^k (whose simplices are tuples with entries in $\{\bullet, \circ\}$ and at most k alternations) is a triangulation of S^k . A simplex of Σ^k is degenerate if and only if it contains two consecutive vertices of the same color. Thus, the only non-degenerate simplices of Σ^0 are the two vertices \bullet, \circ ; Σ^1 additionally has two non-degenerate 1-simplices $[\bullet, \circ]$ and $[\circ, \bullet]$ connecting these two vertices (geometrically, this corresponds to two distinct paths between a pair of antipodal points, each following half of an equatorial circle clockwise); Σ^2 adds two non-degenerate triangles $[\bullet, \circ, \bullet]$ and $[\circ, \bullet, \circ]$ which corresponds to glueing the northern and southern hemisphere, respectively (see Figure 1); Σ^3 adds two non-degenerate 3-simplices; etc.

OBSERVATION 2.6. If X is a (relational)³ simplicial set then a simplicial map $X \rightarrow \Sigma^2$ is completely described by a 2-colouring of the vertex set X_0 with colours yellow or blue. Conversely, a 2-colouring f of X_0 defines a simplicial map if and only if there is no 3-simplex $[u_0, u_1, u_2, u_3]$ of X such that $[f(u_0), f(u_1), f(u_2), f(u_3)]$ has three alternations (is equal to either $[\bullet, \circ, \bullet, \circ]$ or $[\circ, \bullet, \circ, \bullet]$). Moreover, if \mathbb{Z}_2 -acts on X by a simplicial involution ν , then such a 2-colouring defines an equivariant map if and only if u and $\nu(u)$ have different colours for every vertex u of X .

Order complexes of posets. Another important example of relational simplicial sets are order complexes: Given a finite partially ordered set (poset) P , the *order complex* $\Delta(P)$ is the simplicial set whose n -simplices are weakly monotone chains, i.e., $(n+1)$ -tuples $[u_0, \dots, u_n] \in P^{n+1}$ with $u_0 \leq \dots \leq u_n$; moreover, for every monotone map $\alpha: \{0, \dots, m\} \rightarrow \{0, \dots, n\}$, $\alpha^{\Delta(P)}[u_0, \dots, u_n] =$

³The claim is true for arbitrary simplicial set X ; it is only required that Σ^2 is relational.

$[u_{\alpha(0)}, \dots, u_{\alpha(m)}]$ as above. Note that monotonicity of α is crucial here to ensure that chains are mapped to chains. An n -simplex $[u_0, \dots, u_n]$ of $\Delta(P)$ is non-degenerate if and only if $u_0 < \dots < u_n$.

Any monotone map $f: P \rightarrow Q$ between posets naturally extends componentwise to chains and hence to a simplicial map $f: \Delta(P) \rightarrow \Delta(Q)$ between order complexes.

Example 2.7. Let L be a positive integer divisible by 4. Define a partial order \preccurlyeq on $\mathbb{Z}_L = \{0, 1, \dots, L-1\}$ by $a \prec b$ if and only if a is even, b is odd, and $a - b = \pm 1 \pmod L$. We define the simplicial set Γ_L as the order complex of this poset,

$$\Gamma_L := \Delta(\mathbb{Z}_L, \preccurlyeq)$$

The simplicial set Γ_L is a triangulation of S^1 , see Figure 3 (as a digraph, it is a cycle of length L with edges oriented alternatingly). Moreover, the map $\mathbb{Z}_L \rightarrow \mathbb{Z}_L$, $x \mapsto x + L/2$ defines a simplicial involution $\Gamma_L \rightarrow \Gamma_L$ that corresponds to the antipodal involution on S^1 .

If P and Q are posets and if we consider the product $P \times Q$ with the componentwise partial order $(p, q) \leq (p', q')$ if and only if $p \leq p'$ and $q \leq q'$, then $\Delta(P \times Q)$ and $\Delta(P) \times \Delta(Q)$ are isomorphic simplicial sets. In particular, $\Gamma_L^n = \Gamma_L \times \dots \times \Gamma_L$ is a triangulation of the n -dimensional torus $T^n = S^1 \times \dots \times S^1$. Note that the vertices Γ_L^n are n -tuples $\mathbf{u} \in \mathbb{Z}_L^n$, and k -simplices are $(k+1)$ -tuples of vertices $[\mathbf{u}_0, \dots, \mathbf{u}_k]$ such that \mathbf{u}_{i+1} is obtained from \mathbf{u}_i by choosing a subset of coordinates of \mathbf{u}_i all that are even and changing each of them by ± 1 modulo L .

Homomorphism complexes. Given two graphs F and G , the homomorphism complex $\text{Hom}(F, G)$ is a simplicial set capturing the structure of all homomorphisms $F \rightarrow G$. Following Matoušek [25, Section 5.9], we define homomorphism complexes as order complexes of the poset of *multihomomorphisms* from F to G .⁴ By definition, a *multihomomorphism* is a function $f: V(F) \rightarrow 2^{V(G)} \setminus \{\emptyset\}$ such that, for all edges $(u, v) \in E(F)$, we have that

$$f(u) \times f(v) \subseteq E(G).$$

We denote the set of all multihomomorphisms by $\text{mhom}(F, G)$. Multihomomorphisms are partially ordered by component-wise inclusion: $f \leq g$ if and only if $f(u) \subseteq g(u)$ for all $u \in V(F)$.

Definition 2.8. Let F and G be graphs. The *homomorphism complex* $\text{Hom}(F, G)$ is the order complex $\Delta(\text{mhom}(F, G), \leq)$ of the poset of multihomomorphisms.

Multimorphisms can be composed in a natural way: if $f \in \text{mhom}(F, G)$ and $g \in \text{mhom}(G, H)$, then $(g \circ f)(a) = \bigcup_{b \in f(a)} g(b)$ is a multihomomorphism from F to H . In particular, every homomorphism $f: G \rightarrow H$ induces a simplicial map $f_*: \text{Hom}(F, G) \rightarrow \text{Hom}(F, H)$ defined on vertices by mapping a multihomomorphism $m \in \text{mhom}(F, G)$ to the composition $f \circ m$.

In what follows, we will focus on the special case of $\text{Hom}(K_2, G)$, a common tool in the study of graph colourings. Note that a multihomomorphism m from K_2 to a graph G corresponds to an ordered pair of subsets $m(1), m(2) \subseteq V(G)$ such that any pair of vertices $v_1 \in m(1)$ and $v_2 \in m(2)$ are connected by an edge. If G has no loops, then

⁴We remark that in [25] order complexes are defined as simplicial complexes, but the two definitions are equivalent. There are several other alternative definitions of homomorphism complexes that lead to topologically equivalent spaces.




Figure 2: The simplicial set $\Gamma_{12} \cong \text{Hom}(K_2, K_3)$; see also Examples 2.7 and 2.9.

$m(1)$ and $m(2)$ are disjoint and induce a complete bipartite subgraph of G . The natural \mathbb{Z}_2 -action on K_2 that swaps the two vertices induces an action on multihomomorphisms $m: K_2 \rightarrow G$, namely swapping the two sets $m(1)$ and $m(2)$, which in turn induces a \mathbb{Z}_2 -action on the simplicial set $\text{Hom}(K_2, G)$; this action is free provided G has no loops. Moreover, it is easy to check that for every graph homomorphism $f: G \rightarrow H$, the induced simplicial map $f_*: \text{Hom}(K_2, G) \rightarrow \text{Hom}(K_2, H)$ is equivariant.

The following two examples will play an important role in this paper.

Example 2.9. For every odd integer $\ell \geq 3$, $\text{Hom}(K_2, C_\ell)$ is isomorphic to the simplicial set $\Gamma_{4\ell}$ defined above; moreover, this isomorphism is equivariant, i.e., it preserves the \mathbb{Z}_2 -action.

Example 2.10. The simplicial set $\text{Hom}(K_2, K_4)$ is a triangulation of a sphere S^2 ; it is depicted in Figure 3 which shows two hemispheres of this sphere that are glued together along their boundary. The homomorphisms/edges are explicitly labelled (the edge (u, v) is labelled by uv) to highlight the global structure, and a few multihomomorphisms are labelled (where $3|01$ denotes the multihomomorphism $0 \mapsto 3$ and $1 \mapsto \{0, 1\}$, etc.) to explain how the triangles are constructed.

LEMMA 2.11. *There exists an equivariant simplicial map*

$$t: \text{Hom}(K_2, K_4) \rightarrow \Sigma^2.$$

PROOF. By Observation 2.6, such a map is given by a suitable 2-colouring of the vertices of $\text{Hom}(K_2, K_4)$. One suitable 2-colouring is depicted in Figure 3. \square

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




Figure 3: The simplicial set $\text{Hom}(K_2, K_4)$; see also Example 2.10.

These notions naturally generalize to the setting of spaces with group actions: Two equivariant maps $f, g: X \rightarrow Y$ between spaces with \mathbb{Z}_2 -actions are *equivariantly homotopic*, denoted by $f \sim_{\mathbb{Z}_2} g$, if there exists an *equivariant homotopy* between them, i.e., a homotopy $h: X \times [0, 1] \rightarrow Y$ such that all maps $h(\cdot, t): X \rightarrow Y$ are equivariant. We denote by $[X, Y]_{\mathbb{Z}_2}$ the set of all equivariant maps $X \rightarrow Y$ up to equivariant homotopy, i.e.,

$$[X, Y]_{\mathbb{Z}_2} = \{[f] \mid f: X \rightarrow Y \text{ is equivariant}\},$$

where $[f]$ denotes the set of all equivariant maps g s.t. $f \sim_{\mathbb{Z}_2} g$.

3 Overview of the Proof

We present a detailed overview of the proof of Theorem 1.1. Every non-bipartite, loopless graph G contains a cycle C_ℓ of odd length $\ell \geq 3$. In particular, there exists a homomorphism $C_\ell \rightarrow G$, hence every C_ℓ -colourable graph is G -colourable. This yields a trivial reduction from $\text{PCSP}(C_\ell, K_4)$ to $\text{PCSP}(G, K_4)$. Thus, Theorem 1.1 follows from the following:

THEOREM 3.1. *For all odd integers $\ell \geq 3$, the decision problem $\text{PCSP}(C_\ell, K_4)$ is NP-hard.*

We will prove this using Theorem 2.3; to this end, we need to construct a minion homomorphism from $\text{pol}(C_\ell, K_4)$ to a minion \mathcal{B} that contains no constant and is of bounded essential arity.

Informally speaking, as mentioned in Section 2.1, the general philosophy of the algebraic approach to PCSPs is that in order to understand the complexity of a problem, we need to get a good-enough understanding of the structure of its polymorphisms, in our case, the structure of all graph homomorphisms $C_\ell^n \rightarrow K_4$, $n > 0$, i.e., 4-colourings of powers of an odd cycle. Prima facie, such colourings do not seem to have any apparent structure, so we use topology to simplify the problem and reveal more information. In the first step, using homomorphism complexes, we pass from the problem of understanding graph homomorphisms to the problem of understanding equivariant homotopy classes of equivariant continuous maps $T^n \rightarrow S^2$. This provides an approximation of the structure of polymorphisms, nevertheless classifying such continuous maps is still difficult (this is connected to the fact S^2 has many non-trivial higher homotopy groups $\pi_k(S^2)$, $k \geq 3$). Thus, in a second step, we replace S^2 by a “topologically simpler” space Y . We

can then quite explicitly describe, in a third step, the set of $[X, Y]_{\mathbb{Z}_2}$ in terms of a suitable (equivariant) cohomology group (using *equivariant obstruction theory*); this yields a minion homomorphism ϕ from $\text{pol}(C_\ell, K_4)$ to a minion \mathcal{Z}_2 (defined precisely below). The fact that all maps and homotopies are equivariant ensures that the minion \mathcal{Z}_2 does not contain any constants; however, it is still not of bounded essential arity. In a fourth step, we then argue that the image of ϕ actually is of bounded essential arity, for which we use some of the previously neglected combinatorial structure. We now describe these steps in more detail:

Step 1. If X and Y are simplicial sets with \mathbb{Z}_2 -actions, then the set of all equivariant simplicial maps $X^n \rightarrow Y$, $n > 0$, is closed under taking minors, i.e., it forms a minion, which we denote by $\text{spol}(X, Y)$ (this follows easily from the definition of products of simplicial sets).

In the first step of the construction, we use homomorphism complexes to associate with every graph homomorphism $f: C_\ell^n \rightarrow K_4$ an equivariant simplicial map $\mu(f): \Gamma_{4\ell}^n \rightarrow \Sigma^2$, where $\Gamma_{4\ell}^n$ and Σ^2 are the simplicial sets described in Examples 2.7 and 2.5, respectively. The simplicial map $\mu(f)$ is defined as a composition $t \circ f_* \circ \iota_n$:

$$\Gamma_{4\ell}^n \cong \text{Hom}(K_2, C_\ell^n) \xrightarrow{\iota} \text{Hom}(K_2, C_\ell^n) \xrightarrow{f_*} \text{Hom}(K_2, K_4) \xrightarrow{t} \Sigma^2,$$

where $f_*: \text{Hom}(K_2, C_\ell^n) \rightarrow \text{Hom}(K_2, K_4)$ is the simplicial map induced by f , $t: \text{Hom}(K_2, K_4) \rightarrow \Sigma^2$ is the simplicial map from Lemma 2.11, the isomorphism $\Gamma_{4\ell}^n \cong \text{Hom}(K_2, C_\ell^n)$ is given by the isomorphism from Example 2.9, and the simplicial map ι_n is given by the special case $G = C_\ell$ of the following fact:

LEMMA 3.2. *For every graph G and $n \geq 1$, there is an equivariant simplicial map*

$$\iota_n: \text{Hom}(K_2, G)^n \rightarrow \text{Hom}(K_2, G^n).$$

PROOF. Given an n -tuple $m = (m_1, \dots, m_n)$ of multihomomorphisms $m_i: K_2 \rightarrow G$, we can view m as a multihomomorphism $\iota_n(m): K_2 \rightarrow G^n$ by setting $\iota_n(m)(u) = m_1(u) \times \dots \times m_n(u)$ for each vertex u of K_2 . This yields a map $\iota_n: \text{mhom}(K_2, G)^n \rightarrow$

$\text{mhom}(K_2, G^n)$ that is monotone and equivariant and hence extends to the desired simplicial map.⁵ \square

The assignment $f \mapsto \mu(f)$ defines a map $\mu: \text{pol}(C_\ell, K_4) \rightarrow \text{spol}(\Gamma_{4\ell}, \Sigma^2)$ that preserves arity. The map μ does *not* strictly speaking preserve minors, i.e., for a general function $\pi: [n] \rightarrow [m]$, the simplicial maps $\mu(f)^\pi$ and $\mu(f^\pi)$ need not be equal, but it is not hard to see that the induced continuous maps are equivariantly homotopic. Thus, if we denote by $[\mu(f)] \in [T^n, S^2]_{\mathbb{Z}_2}$ the equivariant homotopy class of the map $|\mu(f)|: T^n \cong |\Gamma_{4\ell}^n| \rightarrow |\Sigma^2| \cong S^2$, then $[\mu(f)^\pi] = [\mu(f^\pi)]$.

Step 2. Determining the set of equivariant homotopy classes of maps $[T^n, S^2]_{\mathbb{Z}_2}$ is a difficult problem (and closely related homotopy-theoretic questions regarding maps $X \rightarrow S^2$ for spaces of dimension $\dim X \geq 4$ are algorithmically undecidable [9]). We circumvent this difficulty by enlarging $|\Sigma^2| \cong S^2$ to a larger \mathbb{Z}_2 -space Y that is “homotopically simpler” (in technical terms, Y is an *Eilenberg–MacLane space*), which makes $[T^n, Y]_{\mathbb{Z}_2}$ much easier to compute.

Given a simplicial map $g: \Gamma_{4\ell}^n \rightarrow \Sigma^2$, we define $\eta(g) \in [T^n, Y]_{\mathbb{Z}_2}$ as the equivariant homotopy class of the composition of the geometric realization $|g|: T^n \rightarrow S^2$ with the inclusion map $j: S^2 \hookrightarrow Y$. It is easy to show that η preserves minors, and hence defines a minion homomorphism from $\text{spol}(\Gamma_{4\ell}, \Sigma^2)$ to the minion $\text{hpol}(S^1, Y)$ of equivariant homotopy classes of equivariant maps, i.e., the minion with $\text{hpol}^{(n)}(S^1, P) = [T^n, P]_{\mathbb{Z}_2}$, where $T^n = (S^1)^n$, and minors defined in the natural way.

By considering the composition $\phi := \eta \circ \mu$ with the map constructed in Step 1, we get the following:

LEMMA 3.3. *There are minion homomorphisms $\phi: \text{pol}(C_\ell, K_4) \rightarrow \text{hpol}(S^1, Y)$ and $\eta: \text{spol}(\Gamma_{4\ell}, \Sigma^2) \rightarrow \text{hpol}(S^1, Y)$ such that $\text{im } \phi \subseteq \text{im } \eta$, i.e., for each polymorphism $f: C_\ell^n \rightarrow K_4$, there is a simplicial map $g: \Gamma_{4\ell}^n \rightarrow \Sigma^2$ with $\phi(f) = \eta(g)$.*

Step 3. Next, we give an explicit description of the sets $[T^n, Y]_{\mathbb{Z}_2}$. This description is by the means of functions $f_\alpha: \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ of the form $f_\alpha(x_1, \dots, x_n) = \sum_{i=1}^n \alpha_i x_i$, where $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_2^n$ and $\sum \alpha_i \equiv 1 \pmod{2}$. For a fixed n , the set of such functions forms an affine space, which we denote by $\mathcal{X}_2^{(n)}$, and together, these sets form a function minion \mathcal{X}_2 . Below, we will often identify an affine function f_α with the corresponding n -tuple $\alpha \in \mathbb{Z}_2^n$ of coefficients, i.e., we will often view \mathcal{X}_2 as an abstract minion, with $\mathcal{X}_2^{(n)} = \{\alpha \in \mathbb{Z}_2^n: \sum_i \alpha_i \equiv 1 \pmod{2}\}$.

PROPOSITION 3.4. *For each $n > 0$, there is a bijection*

$$\gamma_n: [T^n, Y]_{\mathbb{Z}_2} \rightarrow \mathcal{X}_2^{(n)}.$$

Moreover, these bijections preserve minors, hence they form a minion isomorphism $\gamma: \text{hpol}(S^1, P) \rightarrow \mathcal{X}_2$.

The proof of this proposition has two parts. One the one hand, using *equivariant obstruction theory*, we can prove the following:

LEMMA 3.5. *The set $[T^n, Y]_{\mathbb{Z}_2}$ has cardinality 2^{n-1} .*

⁵It is easy to see that ι_n is injective, though generally not surjective, and it is known [20, Proposition 18.17] that ι_n defines an equivariant homotopy equivalence between the spaces $|\text{Hom}(K_2, G)|^n$ and $|\text{Hom}(K_2, G^n)|$, but we will not need this fact in what follows.

One the other hand, every $\alpha \in \mathcal{X}_2^{(n)}$ corresponds to a square-free monomial $\prod_{i \in I} z_i$ of odd degree in the variables z_1, \dots, z_n , where $I = \{i \in [n]: \alpha_i = 1\}$. If we view $S^1 = \{z \in \mathbb{C}: |z| = 1\}$ as the unit circle in the complex plane then each such monomial gives rise to an equivariant map $T^n = (S^1)^n \rightarrow S^1$ given by $(z_1, \dots, z_n) \mapsto \prod_{i \in I} z_i$. By composing first with a fixed equivariant inclusion $S^1 \hookrightarrow S^2$ (e.g., the one given by the inclusion $\Sigma^1 \subseteq \Sigma^2$) and then with the inclusion $j: S^2 \hookrightarrow Y$, we can also view each such monomial $\prod_{i \in I} z_i$ as an equivariant map $m_\alpha: T^n \rightarrow Y$. Using a geometrically defined set of \mathbb{Z}_2 -valued invariants $\deg_i, 1 \leq i \leq n$, we will show in Section 4 that these maps are pairwise non-homotopic; in fact, we will see that the map $\gamma_n: [T^n, Y]_{\mathbb{Z}_2} \rightarrow \mathcal{X}_2^{(n)}$ defined by $\gamma_n([f]) = (\deg_1(f), \dots, \deg_n(f))$ satisfies $\gamma_n(m_\alpha) = \alpha$. Thus, γ_n is surjective, and hence bijective, by Lemma 3.5; therefore, every equivariant map $T^n \rightarrow Y$ is equivariantly homotopic to a unique monomial map m_α with $\alpha \in \mathcal{X}_2^{(n)}$. Moreover, we will show that the maps γ_n preserve minors, hence they form a minion isomorphism.

Step 4. Finally, we show (Theorem 5.1) that for every equivariant simplicial map $f: \Gamma_{4\ell}^n \rightarrow \Sigma^2$, the equivariant homotopy class $\eta(f) \in [T^n, Y]_{\mathbb{Z}_2}$ corresponds to an odd monomial map $\prod_{i \in I} z_i$ with $|I| = O(\ell^2)$. This is proved by a combinatorial averaging argument, using the structure of the triangulation $\Gamma_{4\ell}^n$, the fact that simplicial maps to Σ^2 correspond to vertex 2-colourings without alternating 3-simplices, and the geometric definition of the invariants \deg_i . Thus, the image of $\text{pol}(\Gamma_{4\ell}, \Sigma^2)$ under η , and hence the image of $\text{pol}(C_\ell, K_4)$ under ϕ , has bounded essential arity. This concludes the proof of Theorem 3.1.

Comparison with earlier work. The topological approach in [23] for proving hardness of $\text{PCSP}(C_\ell, K_3)$, on which our work builds, required understanding the structure of the set equivariant maps from T^n to S^1 up to equivariant homotopy. Such maps can be classified by much more elementary arguments using fundamental groups and winding numbers, which show that $[T^n, S^1]_{\mathbb{Z}_2}$ is isomorphic to the affine space of maps $\mathbb{Z}^n \rightarrow \mathbb{Z}$ of the form $(x_1, \dots, x_n) \mapsto \sum_i \alpha_i x_i$, where $\alpha_i \in \mathbb{Z}$ and $\sum_i \alpha_i \equiv 1 \pmod{2}$ (this implicitly exploits the fact that S^1 is already an Eilenberg–MacLane space, i.e., has trivial higher homotopy groups). Moreover, bounding the essential arity of such maps that arise from graph homomorphisms is also relatively simple: by considering suitable simplicial embeddings of $|\Gamma_{4\ell}| \cong S^1$ into T^n , the sum $\sum_i |\alpha_i|$ of absolute values of coefficients in such a map can be read of as the winding number of a simplicial map $\Gamma_{4\ell} \rightarrow S^1$, hence $O(\ell)$. By contrast, the more careful counting argument required in our case, although elementary in hindsight, was elusive for several years.

Moreover, although the method based on equivariant obstruction theory was developed to address approximate graph colouring, it first found an application [11] in a hardness proof for promise *linearly-ordered colouring* of hypergraphs, which uses a simpler hardness criterion, and allows for an easier combinatorial argument bounding the arity.

4 Monomial Maps, Degrees, and $[T^n, Y]_{\mathbb{Z}_2}$

The goal of this section is to prove Proposition 3.4.

To this end, we will define, for every equivariant continuous map $f: T^n \rightarrow Y$, a sequence of numbers $\deg_i(f) \in \mathbb{Z}_2, 1 \leq i \leq n$, that

are invariant under equivariant homotopy. As we will see below, these numbers satisfy $\sum_{i=1}^n \deg_i(f) \equiv 1 \pmod{2}$. Thus, by assigning to every equivariant homotopy class $[f] \in [T^n, Y]_{\mathbb{Z}_2}$ the sequence $\gamma_n([f]) = (\deg_1(f), \dots, \deg_n(f)) \in \mathbb{Z}_2^n$, we get a well-defined map $\gamma_n: [T^n, Y]_{\mathbb{Z}_2} \rightarrow \mathcal{X}_2^{(n)}$.

To define the invariants \deg_i and throughout this section, we assume some familiarity with fundamental notions of algebraic topology, including homotopy, CW complexes, simplicial and cellular approximation theorems, and simplicial and cellular homology and cohomology; we refer to Hatcher [15] for general background, and to May et al. [26, Chapters I and II], tom Dieck [31], and Bredon [6] for more details on the equivariant setting.

We will use the fact that the space Y is a CW complex constructed from Σ^2 by attaching higher-dimensional cells (see the proof in the full version of this paper); in particular, the 1-dimensional and 2-dimensional skeleta of Y are Σ^1 and Σ^2 , respectively.

For a CW complex X , let $C_\bullet(X)$ and $C^\bullet(X)$ denote the cellular chain and cochain complexes of X with \mathbb{Z}_2 -coefficients, respectively (since we work with \mathbb{Z}_2 -coefficients, i -dimensional cochains correspond to subsets of i -dimensional cells of X , and i -dimensional chains correspond to finite subsets); in the special case that X is a simplicial complex or simplicial set, $C_\bullet(X)$ and $C^\bullet(X)$ are isomorphic to the simplicial chain and cochain complex of X , respectively.

To motivate the following definition, consider the torus $T^2 \cong |\Gamma_L^2|$ and an equivariant map $f: T^2 \rightarrow Y$. Consider a loop in T^2 that wraps around the first coordinate direction, say the circle $x_1 = \{(z_1, 1) : z_1 \in S^1\} \cong S^1 \subseteq T^2$. Note that the circle x_1 is triangulated by a subcomplex of Γ_L^2 , but it is not fixed under the \mathbb{Z}_2 -action on T^2 . If $f = m_{(1,0)}$ is the monomial map given by $(z_1, z_2) \mapsto j(z_1)$, then f maps the circle x_1 to the 1-skeleton Σ^1 of Y , which is a circle as well, and as a map between circles, f has degree 1; thus, there is an odd number of edges $[v, w]$ in the triangulation of x_1 that satisfy $f(v) = \bullet$ and $f(w) = \circ$. If f is merely equivariantly homotopic to $m_{(1,0)}$, however, then this need no longer be the case: Intuitively, we can visualize the homotopy as “moving” the image of the edges of the torus through the discs in Σ^2 , therefore potentially changing the parity of the degree we are interested in. The homotopy has to be equivariant, however, and thus has to modify each antipodal edge in the opposite way. As a consequence, a 2-dimensional band connecting the circle x_1 and its antipodal circle $v \cdot x_1$ has to be “dragged” around over the discs of Σ^2 and thus, while the degree might change along equivariant homotopies, this difference will be registered in the behaviour of the connecting band. We will now formalize this geometric intuition.

Definition 4.1. Let L, L' be two positive integers divisible by 4, and consider the 2-dimensional torus $T^2 = |\Gamma_L \times \Gamma_{L'}|$.

Let $e^0 \in C^1(Y)$ and $d^0 \in C^2(Y)$ be the dual of $e_0 = [\bullet, \circ] \in C_1(\Sigma^2) = C_1(Y)$ and $d_0 = [\bullet, \circ, \bullet] \in C_2(\Sigma^2) = C_2(Y)$ respectively (i.e., $e^0([\bullet, \circ]) = 0$ and $e^0([\bullet, \circ]) = 1$, similarly for d^0). Moreover, let $x_1 \in Z_1(\Gamma_L \times \Gamma_{L'})$ be the “first coordinate cycle” in $|\Gamma_L \times \Gamma_{L'}| \cong T^2$ (i.e., $x_1 = \sum_{k=0}^{L-1} [(k, 0)(k+1, 0)] \in Z_1(\Gamma_L \times \Gamma_{L'})$), and let $b_1 \in C_2(\Gamma_L \times \Gamma_{L'})$ be the “band” connecting x_1 with $v \cdot x_1$ (i.e., $\partial b_1 = x_1 + v \cdot x_1$; see Figure 4).

Let $f: |\Gamma_L \times \Gamma_{L'}| \rightarrow Y$ be an equivariant map. By the (equivariant) cellular approximation theorem, f induces an equivariant cochain map $f^*: C^\bullet(Y) \rightarrow C^\bullet(\Gamma_L \times \Gamma_{L'})$ (i.e., equivariant homomorphisms




Figure 4: Coordinate cycle and band in T^2 and T^3 .

$f^*: C^i(Y) \rightarrow C^i(X)$ that commute with the coboundary map). We define

$$\deg_1(f) = (f^*(e^0)(x_1) + f^*(d^0)(b_1)) \pmod{2}$$

Crucially, this notion of degree is invariant under equivariant homotopies:

LEMMA 4.2. Fix positive integers L, L_0 and L_1 divisible by 4. Let $f_0: |\Gamma_L \times \Gamma_{L_0}| \rightarrow Y$ and $f_1: |\Gamma_L \times \Gamma_{L_1}| \rightarrow Y$ be equivariant maps that are equivariantly homotopic. Then $\deg_1(f_0) = \deg_1(f_1)$.

PROOF. Assume first that $L_0 = L_1$. By the cellular approximation theorem again, there is an equivariant cochain homotopy between the induced cochain maps $f_0^*, f_1^*: C^\bullet(Y) \rightarrow C^\bullet(X)$, i.e., there exist equivariant homomorphisms $h: C^i(Y) \rightarrow C^{i-1}(X)$ satisfying

$$f_0^* + f_1^* = \delta h + h \delta.$$

Therefore, on the cochains of dimension 1 we have:

$$\begin{aligned} f_0^*(e^0)(x_1) + f_1^*(e^0)(x_1) &= (\delta h(e^0))(x_1) + (h \delta(e^0))(x_1) \\ &= 0 + h(d^0 + d^1)(x_1) \end{aligned}$$

where the second equality is obtained by using the fact that $\partial x_1 = 0$ and $\delta e^0 = d^0 + d^1$.

On the 2-cochains we have:

$$\begin{aligned} f_0^*(d^0)(b_1) + f_1^*(d^0)(b_1) &= (\delta h(d^0))(b_1) + (h \delta(d^0))(b_1) \\ &= h(d^0)(x_1 + v \cdot x_1) + 0 \end{aligned}$$

where we use that $\partial b_1 = x_1 + v \cdot x_1$ and $\delta d^0 = 0$.

Moreover, h is equivariant, hence $h(d^0)(v \cdot x_1) = h(v \cdot d^0)(x_1)$. Summing everything together, using this fact and that $v \cdot d^0 = d^1$, we obtain that

$$\begin{aligned} \deg_1(f_0) + \deg_1(f_1) &= (f_0^*(e^0)(x_1) + f_1^*(e^0)(x_1)) \\ &\quad + (f_0^*(d^0)(b_1) + f_1^*(d^0)(b_1)) \\ &= h(d^0 + d^1)(x_1) + h(d^0)(x_1 + v \cdot x_1) \\ &= h(d^0)(x_1 + v \cdot x_1 + x_1 + v \cdot x_1) \\ &= h(d^0)(2x_1) = 2h(d^0)(x_1) = 0 \pmod{2}. \end{aligned}$$

If $L_0 \neq L_1$, suppose without loss of generality that $L_0 < L_1$. Then $\Gamma_L \times \Gamma_{L_1}$ is a subdivision of $\Gamma_L \times \Gamma_{L_0}$, and the equivariant chain map $\iota: C_\bullet(\Gamma_L \times \Gamma_{L_0}) \rightarrow C_\bullet(\Gamma_L \times \Gamma_{L_1})$ that maps every i -cell σ of $\Gamma_L \times \Gamma_{L_0}$ to the sum of i -cells of $\Gamma_L \times \Gamma_{L_1}$ that are contained in σ is a chain homotopy equivalence. Thus, by the previous case $\deg_1(f_0) = \deg_1(f_1 \circ \iota)$ and from the definition of degree, $\deg_1(f_1 \circ \iota) = \deg_1(f_1)$. \square

We can now define $\deg_i(f)$ of an equivariant map $f: T^n \rightarrow Y$ as $\deg_i(f^{\sigma_i})$ for a suitable 2-minor $f^{\sigma_i}: T^2 \rightarrow Y$ (see Figure 4):

Definition 4.3. Let L a positive integer divisible by 4, and let $f: |\Gamma_L^n| \rightarrow Y$ be a \mathbb{Z}_2 -equivariant map. For $i \in [n]$, we define $\sigma_i: [n] \rightarrow [2]$ by $\sigma_i(i) = 1$ and $\sigma_i(j) = 2$ for $j \neq i$. Then the i -degree of f is defined as

$$\begin{aligned} \deg_i(f) &= \deg_i(f^{\sigma_i}) \\ &= ((f \circ \sigma_i)^*(e^0)(x_1) + (f \circ \sigma_i)^*(d^0)(b_1)) \bmod 2 \end{aligned}$$

An immediate consequence of Lemma 4.2 is the invariance of the i -degree under equivariant homotopies:

COROLLARY 4.4. Let $f_0, f_1: |\Gamma_L^n| \rightarrow Y$ be equivariant maps that are equivariantly homotopic. Then $\deg_i(f_0) = \deg_i(f_1)$ for all $i \in [n]$.

PROOF. Since f_0 and f_1 are equivariantly homotopic, so are their minors $f_0^{\sigma_i}$ and $f_1^{\sigma_i}$. \square

It will be convenient to extend the notation for monomial maps to general integer coefficients. As before, let us view $S^1 = \{z \in \mathbb{C} : |z| = 1\}$ as the unit circle in the complex plane. Given an n -tuple of integers $\alpha = (\alpha_1, \dots, \alpha_n)$ with $\sum_i \alpha_i \equiv 1 \pmod 2$, we get an equivariant map from T^n to S^1 defined by $(z_1, \dots, z_n) \mapsto z_1^{\alpha_1} \dots z_n^{\alpha_n}$. By composing this map first with a fixed equivariant inclusion $S^1 \hookrightarrow S^2$ and then with the inclusion $j: S^2 \rightarrow Y$, we get an equivariant monomial map $m_\alpha: T^n \rightarrow Y$ given by

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

Remark 4.5. It is not hard to observe that the assignment $\alpha \mapsto m_\alpha$ preserves minors when α is interpreted as a function $f: \mathbb{Z}^n \rightarrow \mathbb{Z}$. While we implicitly use this minion homomorphism, this is not the minion homomorphism we are looking for – importantly, \mathcal{Z}_2 is not a subminion of the minion of tuples $\alpha \in \mathbb{Z}^n$ with $\sum_i \alpha_i \equiv 1 \pmod 2$ since, e.g., the unary minor of $(1, 1, 1)$ disagrees in the two minions.

Since monomial maps form a minion, we can easily compute the degree of any of them.

LEMMA 4.6. Let $\alpha \in \mathbb{Z}^n$ such that $\sum_i \alpha_i \equiv 1 \pmod 2$. Then

$$\deg_i(m_\alpha) = \alpha_i \bmod 2$$

PROOF. Let σ_i the minor used to define \deg_i . Then $m_\alpha^{\sigma_i} = m_\beta$ with $\beta = (\alpha_i, \sum_{j \neq i} \alpha_j) \in \mathbb{Z}^2$. Since the image of m_β is contained in the 1-skeleton, $m_\beta^*(d^0) \equiv 0$. Moreover, $(m_\beta)_*(x_1) = \alpha_i e_0 + \alpha_i e_1$, hence $e^0((m_\beta)_*(x_1)) = \alpha_i$ and thus $\deg_i(m_\alpha) = \deg(m_\beta) = \alpha_i + 0 \bmod 2$. \square

COROLLARY 4.7. Let $\alpha, \beta \in \mathcal{Z}_2^{(n)}$. Then m_α and m_β are equivariantly homotopic if and only if $\alpha = \beta$.

PROOF. If $\alpha = \beta$ then m_α and m_β are identical as maps. Conversely, if m_α and m_β are equivariantly homotopic, then $\deg_i(m_\alpha) = \deg_i(m_\beta)$ for all $i \in [n]$, by Corollary 4.4. Thus, by Lemma 4.6, $\alpha_i = \beta_i$, for all $i \in [n]$. \square

We are now ready to prove Proposition 3.4:

PROOF OF PROPOSITION 3.4. For every $n \geq 1$, consider the map $\gamma_n: [T^n, Y]_{\mathbb{Z}_2} \rightarrow \mathbb{Z}_2^n$ given by

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

By Corollary 4.4, this mapping is well-defined. Moreover, by Lemma 4.6, if $\alpha \in \mathcal{Z}_2^{(n)}$, then the homotopy class $[m_\alpha] \in [T^n, Y]_{\mathbb{Z}_2}$ of the corresponding monomial map satisfies $\gamma_n([m_\alpha]) = \alpha$, i.e., the homotopy classes $[m_\alpha]$, $\alpha \in \mathcal{Z}_2^{(n)}$, are pairwise distinct, and by Lemma 3.5, they account for all elements of $[T^n, Y]_{\mathbb{Z}_2}$, i.e., every equivariant map $f: T^n \rightarrow Y$ is equivariantly homotopic to m_α for a unique $\alpha \in \mathcal{Z}_2^{(n)}$. It follows that $\gamma_n([f]) \in \mathcal{Z}_2^{(n)}$ and that γ_n is a bijection.

Furthermore, if $\alpha \in \mathbb{Z}^n$ with $\sum_i \alpha_i = 1$, and $\pi: [n] \rightarrow [m]$ then

$$\gamma_m([m_\alpha]) = (\alpha_1 \bmod 2, \dots, \alpha_n \bmod 2)$$

by Lemma 4.6, hence $\gamma_n([m_\alpha])^\pi = \beta$, where $\beta_j = (\sum_{i \in \pi^{-1}(j)} \alpha_i) \bmod 2$. Furthermore, $m_\alpha^\pi = m_{\beta'}$ where $\beta'_j = \sum_{i \in \pi^{-1}(j)} \alpha_i$. Consequently,

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

hence

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

Thus, the maps γ_n preserve minors for homotopy classes of monomial maps. Since these account for all homotopy classes, the maps γ_n define a minion isomorphism $\gamma: \text{hpol}(S^1, Y) \rightarrow \mathcal{Z}_2$. \square

Finally, we show non zero degree guarantees a colour swapping edge.

LEMMA 4.8. Let $f: \Gamma_L \times \Gamma_{L'} \rightarrow \Sigma^2$ a simplicial equivariant map such that $\deg_1(f) = 1$. Then there is an horizontal color swapping edge, that is there is a vertex $(v_1, v_2) \in \Gamma_L \times \Gamma_{L'}$ such that $f(v_1, v_2) = \bullet$ and $f(v_1 + 1, v_2) = \bullet$.

PROOF. Suppose, by contradiction, that every horizontal edge is monochrome. Therefore, the image of the horizontal coordinate cycle is constant so that $f^*(e_0)(x_1) = 0$. Additionally, the image of a triangle is non degenerate if and only if it is alternating (i.e., $f([u, v, w]) = [\bullet, \bullet, \bullet]$ or $[\bullet, \bullet, \bullet]$); since we are assuming that every horizontal edge is monochrome, there are no alternating triangles and therefore $f^*(d^0)(b_1) = 0$. The total degree is then $\deg_1(f) = f^*(e_0)(x_1) + f^*(d^0)(b_1) = 0$. \square

5 Bounding Essential Arity

We prove the key technical result that bounds the essential arity of simplicial maps from Γ_L^n to Σ^2 .

THEOREM 5.1. Let $L \geq 4$ be an integer divisible by 4, let $f: \Gamma_L^n \rightarrow \Sigma^2$ be an equivariant simplicial map such that the composition with the map $\Sigma^2 \rightarrow Y$ is equivariantly homotopic to the map given by the monomial $\prod_{i \in [I]} z_i$; equivalently, $\deg_i(f) = 1$ if and only if $i \in I$. Then $|I| \leq O(L^2)$.

We recall (Observation 2.6) that equivariant simplicial maps $f: \Gamma_L^n \rightarrow \Sigma^2$ correspond bijectively to 2-colourings of the vertices of Γ_L^n with the following two properties: The colouring is equivariant (i.e., every pair of antipodal vertices of Γ_L^n have distinct colours), and no 3-simplex $[u_0, u_1, u_2, u_3]$ is coloured with alternating colours. We will show that this is impossible if $|I|$ is large; more precisely, we will show that if $i \in I$, then there are many edges $[u, v]$ such

that the colours of \mathbf{u} and \mathbf{v} are different and \mathbf{u} and \mathbf{v} differ only in the i th coordinate (note that the difference in this coordinate is 1 by the definition of Γ_L^n). This is then used to show that we need to have an alternating simplex of dimension proportional to the size of I .

To present the details of the argument, we need a number of definitions. We recall the description of Γ_L^n : Its vertices are the n -tuples $\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{Z}_L^n$; edges (1-simplices) are pairs $[\mathbf{u}, \mathbf{v}]$ of vertices such that \mathbf{v} is obtained from \mathbf{u} by choosing a non-empty subset of coordinates of \mathbf{u} that are all even, and changing each of them by ± 1 modulo L ; and the k -simplices are $(k+1)$ -tuples $[\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_k]$ such that $[\mathbf{u}_{j-1}, \mathbf{u}_j]$ is an edge for $1 \leq j \leq k$. We define the *height* $\text{ht}(\mathbf{u})$ of a vertex $\mathbf{u} = (u_1, \dots, u_n)$ as the number of coordinates $i \in [n]$ such that u_i is odd; moreover, we define the height of an edge $[\mathbf{u}, \mathbf{v}]$ as the height of \mathbf{u} . Note that every edge $[\mathbf{u}, \mathbf{v}]$, we have $\text{ht}(\mathbf{u}) < \text{ht}(\mathbf{v})$. A special role will be played by edges $[\mathbf{u}, \mathbf{v}]$ such that $\text{ht}(\mathbf{v}) = \text{ht}(\mathbf{u}) + 1$, or equivalently, such that \mathbf{u} and \mathbf{v} differ in exactly one coordinate; we call such edges *coordinate edges*. More precisely, we say that an edge $[\mathbf{u}, \mathbf{v}]$ is in *coordinate direction* i if \mathbf{u} and \mathbf{v} differ exactly in the i th coordinate. For $i \in [n]$, we denote the set of all edges in coordinate direction i by E_i , and denote by $E := E_1 \sqcup \dots \sqcup E_n$ the set of all coordinate edges. We will also need the following more refined classification: For $i \in [n]$ and $0 \leq h \leq n-1$, let $E_i(h)$ denote the set of all edges in E_i of height h , and let $E(h) = E_1(h) \sqcup \dots \sqcup E_n(h)$ denote the set of all coordinate edges of height h (note that the height h of a coordinate edge determines the heights h and $h+1$ of both endpoints).

Given a 2-colouring of the vertices of Γ_L^n , we say that edge $[\mathbf{u}, \mathbf{v}]$ of Γ_L^n is *colour-swapping* if \mathbf{u} and \mathbf{v} have different colours. We now state a key lemma used in the proof of Theorem 5.1. The lemma shows that, if f depends on the coordinate i (up to homotopy), then some fraction (independent from the arity of f) of edges in coordinate direction i is colour-swapping.

LEMMA 5.2. *Let $f: \Gamma_L^n \rightarrow \Sigma^2$ be an equivariant simplicial map such that $\deg_i(f) = 1$ and let $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$. Then a fraction of at least $\frac{1}{CL^2}$ of the edges in $E_i(h) \sqcup E_i(n-1-h)$ are colour-swapping, where $C > 0$ is a suitable constant.*

We postpone the proof of the lemma, and first show how it implies Theorem 5.1.

PROOF OF THEOREM 5.1 ASSUMING LEMMA 5.2. We first observe that the theorem reduces to the case that n is odd and $I = [n]$. To see this, let $m = |I|$, and choose any function $\pi: [n] \rightarrow [m]$ that is injective on I . Then the minor f^π is an equivariant simplicial map $f^\pi: \Gamma_L^m \rightarrow \Sigma^2$ that is equivariantly homotopic to the monomial map $\prod_{i \in [m]} z_i$ (see the full version for a detailed proof).

Thus (by replacing f by f^π and n by m), we may assume without loss of generality that n is odd and $I = [n]$, i.e., $\deg_i(f) = 1$ for all $i \in [n]$. Now, consider a non-degenerate n -simplex $\sigma = [\mathbf{u}_0, \dots, \mathbf{u}_n]$ of Γ_L^n chosen uniformly at random among all such n -simplices of Γ_L^n . For $0 \leq h \leq n-1$, define the random variable $X_h(\sigma)$ as 1 or 0 depending on whether the edge $[\mathbf{u}_h, \mathbf{u}_{h+1}]$ is colour-swapping or not. Then $X(\sigma) := \sum_{h=0}^{n-1} X_h(\sigma)$ equals the total number of times the colour of $f(\mathbf{u}_i)$ changes as we traverse the vertices of σ in their given order. Observe that, for every $0 \leq h \leq n-1$, the edge $[\mathbf{u}_h, \mathbf{u}_{h+1}]$ of the random simplex σ is distributed uniformly among

all edges of $E(h)$ (this is since the simplicial automorphisms of Γ_L^n act transitively on $E(h)$). Thus, the expected value $\mathbb{E}[X_h(\sigma)]$ is the probability that a uniformly random edge in $E(h)$ is colour-swapping. Moreover, by Lemma 5.2 and summing over $1 \leq i \leq n$, we get that for every $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$, the fraction of edges in $E(h) \sqcup E(n-1-h)$ that are colour-swapping is at least $\frac{1}{CL^2}$. Hence, by linearity of expectation, $\mathbb{E}[X_h(\sigma)] + \mathbb{E}[X_{n-1-h}(\sigma)] \geq \frac{1}{CL^2}$ for $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$. Consequently,

$$\mathbb{E}[X(\sigma)] = \sum_{h=0}^{n-1} \mathbb{E}[X_{n-1-h}(\sigma)] \geq \lfloor \frac{n-1}{3} \rfloor \cdot \frac{1}{CL^2}.$$

Thus, there exists some n -simplex $\sigma = [\mathbf{u}_0, \dots, \mathbf{u}_n]$ of Γ_L^n such that the colour of $f(\mathbf{u}_i)$ changes at least k times, where $k = \lfloor \frac{n-1}{3} \rfloor \cdot \frac{1}{CL^2}$, i.e., σ contains some k -simplex $[\mathbf{u}_{i_0}, \mathbf{u}_{i_1}, \dots, \mathbf{u}_{i_k}]$ whose colours alternate. Since f is a simplicial map to Σ^2 , this implies that $k \leq 2$ as noted above, and therefore $\lfloor \frac{n-1}{3} \rfloor \leq 2CL^2$, hence $|I| = n = O(L^2)$. \square

The rest of this section is dedicated to proving Lemma 5.2.

In this proof, we will use Lemma 4.8 in combination with another averaging argument over a special family of triangulated 2-dimensional tori $\Gamma_L \times \Gamma_{L'}$, which we call *slices*, that are simplicially (and equivariantly) embedded in the triangulation Γ_L^n .

To simplify notation, let us fix a coordinate direction, say $i = 1$, and write $\Gamma_L \times \Gamma_{L'}^{n-1}$. The archetype of a slice is the following *standard slice*: Consider the *diagonal embedding* $\text{diag}: \Gamma_L \hookrightarrow \Gamma_L^{n-1}$ given by $\text{diag}(y) = (y, \dots, y)$. This is an equivariant simplicial map, which induces an equivariant simplicial embedding $s_{\text{diag}}: \Gamma_L \times \Gamma_L \hookrightarrow \Gamma_L^n$ given by $s_{\text{diag}} := 1_{\Gamma_L} \times \text{diag}$, i.e., $s_{\text{diag}}(x, y) = (x, y, \dots, y)$.

More generally, let L' be an integer divisible by 4, and let $\zeta: \Gamma_{L'} \rightarrow \Gamma_L^{n-1}$ be an equivariant simplicial map; we call ζ a *generalized diagonal* if its geometric realization $|\zeta|$, seen as an equivariant embedding $S^1 \rightarrow T^{n-1}$, is equivariantly homotopic to the diagonal embedding $S^1 \rightarrow T^{n-1}$ (here, we implicitly fix equivariant homeomorphisms $|\Gamma_{L'}| \cong S^1 \cong |\Gamma_L|$). Given a generalized diagonal ζ , we call the induced equivariant simplicial embedding $s_\zeta: \Gamma_L \times \Gamma_{L'} \rightarrow \Gamma_L^n$ given by $s_\zeta = 1_{\Gamma_L} \times \zeta$ a *slice*. Moreover, we say that s_ζ is an h -*slice* if every vertex of Γ_L^{n-1} in the image of ζ is at height h or $n-1-h$, or equivalently, if every edge of Γ_L^n that lies in both E_1 and the image of s_ζ belongs to $E_1(h) \sqcup E_1(n-1-h)$.

LEMMA 5.3. *Let $f: \Gamma_L^n \rightarrow \Sigma^2$ be an equivariant simplicial map such that $\deg_1(f) = 1$, and let $s_\zeta: \Gamma_L \times \Gamma_{L'} \rightarrow \Gamma_L^n$ be a slice (respectively, an h -slice, $0 \leq h \leq n-1$). Then the image of s_ζ contains at least one edge in E_1 (respectively, in $E_1(h) \sqcup E_1(n-1-h)$) that is colour-swapping.*

PROOF. The composition $f \circ s_{\text{diag}}$ is the same as the 2-minor f^π of f given by the map $\pi: [n] \rightarrow [2]$, $\pi(1) = 1$ and $\pi(j) = 2$ for $2 \leq j \leq n$. Thus, $\deg_1(f \circ s_{\text{diag}}) = \deg_1(f^\pi) = \deg_1(f) = 1$ by Definition 4.3. Moreover, by definition of generalized diagonals, it follows that $|f \circ s_\zeta|$ and $|f \circ s_{\text{diag}}|$ are equivariantly homotopic as maps $T^2 = S^1 \times S^1 \rightarrow S^2$, hence $\deg_1(f \circ s_\zeta) = \deg_1(f \circ s_{\text{diag}}) = 1$ (here, we use that the equivariant homeomorphism $|\Gamma_L \times \Gamma_{L'}| \cong |\Gamma_L \times \Gamma_L|$ fixes the two coordinate copies of S^1 in T^2). Thus, the existence of the desired colour-swapping edge follows from Lemma 4.8. \square




Figure 5: A path starting with the point $u_0 = (1, 0, 0, 0)$ shown as projection on the first two (left) and last two coordinates (right).

The last puzzle piece we need to prove Lemma 5.2 (and thus to complete the proof of Theorem 5.1) is the following lemma which constructs a generalised diagonal of a special shape.

LEMMA 5.4. *Let $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$. Then there exists a generalised diagonal $\zeta_0: \Gamma_{3L} \rightarrow \Gamma_L^{n-1}$ whose image contains only vertices of height h or $n-1-h$; moreover, the vertices of height h and $n-1-h$ alternate.*

PROOF. We start with constructing a simplicial map $\zeta_0: \Gamma_{3L} \rightarrow \Gamma_L^{n-1}$, i.e., a cyclic path in Γ_L^{n-1} , that contains only vertices of height h or $n-1-h$.

We start with the vertex u_0 of the form $u_0 = (1, \dots, 1, 0, \dots, 0)$ where the first h coordinates are 1, and construct a path from u_0 to its antipode in pieces of length 3. The first three steps of the path have the following form (where the first three blocks are of length h and the last block is of length $n-1-3h$).

$$\begin{aligned} u_0 &= (\underbrace{1, \dots, 1}_h, \underbrace{0, \dots, 0}_{n-1-h}, \dots, 0) \\ u_1 &= (\underbrace{1, \dots, 1}_h, \underbrace{0, \dots, 0}_h, \underbrace{1, \dots, 1}_{n-1-2h}, \dots, 1) \\ u_2 &= (\underbrace{2, \dots, 2}_h, \underbrace{0, \dots, 0}_h, \underbrace{1, \dots, 1}_h, \underbrace{0, \dots, 0}_{n-1-3h}, \dots, 0) \\ u_3 &= (\underbrace{2, \dots, 2}_h, \underbrace{1, \dots, 1}_{n-1-h}, \dots, 1) \end{aligned}$$

In the first step, we increase the values in the last two blocks changing $h + (n-1-3h) = n-1-2h$ values. In the second step, we increase the value in the first block and decrease the value in the last block, again changing the same number of values. And in the third step, we increase the values in the second and the last block. See also Fig. 5 for a visual representation of the case $n=4$ and $h=1$. Note that the height of u_0 and u_2 is h and the height of u_1 and u_3 is $n-1-h$, and that u_3 is u_0 shifted along the diagonal by 1.

We then repeat this pattern (until we return to u_0) by adding 1 to all coordinates in each subsequent sequence of three steps, i.e.,

$$u_4 = (\underbrace{2, \dots, 2}_h, \underbrace{1, \dots, 1}_h, \underbrace{2, \dots, 2}_{n-1-2h}, \dots, 2),$$

etc. It is easy to check that the height of u_{2k} is h and the height of u_{2k+1} is $n-1-h$ for all k , and that subsequent vertices are connected by an edge in Γ_L^{n-1} . Furthermore, observe that $u_{k+3L/2} = u_k + \frac{L}{2}\mathbb{1}$

is the antipode of u_k , hence $\zeta_0: \Gamma_{3L} \rightarrow \Gamma_L^{n-1}$ defined by $\zeta_0(k) = u_k$ is an equivariant simplicial map.

Next, we prove that ζ_0 is a generalized diagonal. We view the geometric realization of Γ_L as $\mathbb{R}/L\mathbb{Z} \cong S^1$. Observe that every point $x = (x_1, \dots, x_{n-1}) \in T^{n-1}$ on the (geometric realization of the) path from u_0 to $u_3 = u_0 + \mathbb{1}$ satisfies $x_i \in [1, 2]$ if $i \leq h$ and $x_i \in [0, 1]$ if $i > h$; thus, $x \in [1, 2]^h \times [0, 1]^{n-1-h}$, i.e., x lies inside a unit box. Since this box is convex, we can homotope the path to the “straight” path from u to $u + \mathbb{1}$ inside the box, keeping the endpoints fixed, by linear interpolation. By an analogous argument applied to each path segment corresponding to a sequence of three steps from u_{3k} to u_{3k+3} , we get a homotopy between the embedding ζ_0 and a translated copy of the diagonal that passes through u_0 . Moreover, this translated copy to the diagonal is homotopic to the diagonal itself, hence ζ_0 is a generalized diagonal (note that translated copies of the diagonal are not simplicial embeddings in general, which is why we use the more complicated construction). \square

We may now finish the proof of Lemma 5.2 and, consequently, of Theorem 5.1.

PROOF OF LEMMA 5.2. Let us fix a coordinate direction, without loss of generality $i = 1$, and let $f: \Gamma_L^n \rightarrow \Sigma^2$ be an equivariant simplicial map such that $\deg_1(f) = 1$. Let $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$.

First, we prove that there exists a collection Z of generalised diagonals that contain only vertices of heights h and $n-1-h$ such that each vertex of such a height appears in the same number of diagonals accounting for multiplicity. This collection is constructed by shifting the diagonal ζ_0 obtained in Lemma 5.4 by some automorphisms of Γ_L^{n-1} . We consider only those automorphisms that respect the winding direction in each coordinate, which consequently the homotopy class of the diagonal. More precisely, consider the subgroup A of automorphism group of Γ_L^{n-1} generated by automorphisms of one of the following two types:

- a_π , where $\pi: [n-1] \rightarrow [n-1]$ is permutation, which permutes the coordinates of each vertex, i.e.,

$$a_\pi(u_1, \dots, u_{n-1}) = (u_{\pi(1)}, \dots, u_{\pi(n-1)});$$

- b_i , where $i \in [n-1]$, which shifts the coordinate i by 2, i.e.,

$$b_i(u_1, \dots, u_{n-1}) = (u_1, \dots, u_{i-1}, (u_i + 2) \bmod L, u_{i+1}, \dots, u_{n-1}).$$

Observe that A acts transitively on vertices of height h : for example, first use b_i 's to make all coordinates 0 or 1, and then use a_π to permute them in the first h positions. In fact, the orbits of A are exactly sets of vertices of the same height. Now, we let $Z = \{g \circ \zeta_0 \mid g \in A\}$. Since this family is invariant under the action of A which, as we said, is transitive on vertices of height h and of height $n-1-h$, respectively, each such vertex appears in the same number of generalised diagonals in Z . Since the vertices of height h and $n-1-h$ alternate in ζ_0 , and consequently, they alternate in each of the shifts, we also get the number of times a vertex of height h appears is the same as the number of times a vertex of height $n-1-h$ appears.

For each $\zeta \in Z$, the image of the corresponding h -slice $s_\zeta: \Gamma_L \times \Gamma_{3L} \rightarrow \Gamma_L^n$ contains $3L^2$ edges in $E_1(h) \sqcup E_1(n-1-h)$, and at least one of these edges is colour-swapping, by Lemma 5.3. Moreover, the number of slices $\zeta \in Z$ whose image contain a given edge in $E_i(h) \sqcup E_i(n-1-h)$ does not depend on the edge. Thus, we

can choose a uniformly random element of $E_i(h) \sqcup E_i(n-1-h)$ by first choosing a uniformly random element $\zeta \in Z$, and then choosing uniformly at random a vertex $v \in \Gamma_{3L}$ and an edge in coordinate direction i which projects to $\zeta(v)$. Since the probability that we selected a colour-swapping edge in the last choice is at least $\frac{1}{3L^2}$, the overall probability that an uniformly random edge from $E_i(h) \sqcup E_i(n-1-h)$ is colour-swapping is also at least $\frac{1}{3L^2}$. \square

References

- [1] Per Austrin, Venkatesan Guruswami, and Johan Håstad. 2017. $(2+\epsilon)$ -Sat Is NP-hard. *SIAM Journal on Computing* 46, 5 (2017), 1554–1573. doi:10.1137/15M1006507
- [2] 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:10.1145/3457606
- [3] Joshua Brakensiek and Venkatesan Guruswami. 2016. New Hardness Results for Graph and Hypergraph Colorings. In *Conference on Computational Complexity (CCC 2016) (LIPIcs, Vol. 50)*, Ran Raz (Ed.), Schloss Dagstuhl–LZI, Dagstuhl, Germany, 14:1–14:27. doi:10.4230/LIPIcs.CCC.2016.14
- [4] Joshua Brakensiek and Venkatesan Guruswami. 2021. Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy. *SIAM Journal on Computing* 50, 6 (2021), 1663–1700. doi:10.1137/19M128212X
- [5] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. 2022. An invariance principle for the multi-slice, with applications. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021*. IEEE Computer Soc., Los Alamitos, CA, 228–236. doi:10.1109/FOCS52979.2021.00030
- [6] Glen E. Bredon. 1967. *Equivariant Cohomology Theories*. Lecture Notes in Mathematics, Vol. 34. Springer Berlin Heidelberg. doi:10.1007/BFb0082690
- [7] Andrei A. Bulatov. 2017. A Dichotomy Theorem for Nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, Berkeley, CA, USA, 319–330. doi:10.1109/FOCS.2017.37
- [8] Jakub Bulín, Andrei Krokhin, and Jakub Opršal. 2019. Algebraic Approach to Promise Constraint Satisfaction. In *Proc. of the 51st Annual ACM-SIGACT Symposium on the Theory of Computing, STOC 2019*. ACM, New York, USA, 602–613. doi:10.1145/3313276.3316300
- [9] Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner. 2013. Extending continuous maps: polynomiality and undecidability. In *Proceedings of the 2013 ACM Symposium on Theory of Computing, STOC 2013*. ACM, New York, 595–604. doi:10.1145/2488608.2488683
- [10] Irit Dinur, Elchanan Mossel, and Oded Regev. 2009. Conditional hardness for approximate coloring. *SIAM J. Comput.* 39, 3 (2009), 843–873. doi:10.1137/07068062X
- [11] Marek Filakovský, Tamio-Vesa Nakajima, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. 2024. Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs. In *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024 (Clermont-Ferrand, France, March 12–14, 2024) (LIPIcs, Vol. 289)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 34:1–34:19. doi:10.4230/LIPIcs.STACS.2024.34 arXiv:2406.00357
- [12] Greg Friedman. 2012. Survey article: An elementary illustrated introduction to simplicial sets. *Rocky Mountain J. Math.* 42, 2 (2012), 353–423. doi:10.1216/RMJ-2012-42-2-353
- [13] Venkatesan Guruswami and Sanjeev Khanna. 2004. On the Hardness of 4-Coloring a 3-Colorable Graph. *SIAM J. Discret. Math.* 18, 1 (2004), 30–40. doi:10.1137/S0895480100376794
- [14] Venkatesan Guruswami and Sai Sandeep. 2020. d -to-1 hardness of coloring 3-colorable graphs with $O(1)$ colors. In *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.
- [15] Allen Hatcher. 2002. *Algebraic topology*. Cambridge University Press, Cambridge. xii+544 pages.
- [16] Pavol Hell and Jaroslav Nešetřil. 1990. On the complexity of H -coloring. *J. Comb. Theory B* 48, 1 (1990), 92–110. doi:10.1016/0095-8956(90)90132-J
- [17] Richard M. Karp. 1972. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, Mar 20–22, 1972, Yorktown Heights, New York*, Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (Eds.). Springer US, Boston, MA, 85–103. doi:10.1007/978-1-4684-2001-2_9
- [18] 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 (Vancouver, BC, Canada, June 24–28, 2024)*. ACM, 331–339. doi:10.1145/3618260.3649768 arXiv:2406.00357
- [19] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. 2000. On the Hardness of Approximating the Chromatic Number. *Combinatorica* 20, 3 (2000), 393–415. doi:10.1007/s004930070013
- [20] Dmitry Kozlov. 2008. *Combinatorial algebraic topology*. Algorithms and Computation in Mathematics, Vol. 21. Springer. XX, 390 pages. doi:10.1007/978-3-540-71962-5
- [21] Andrei Krokhin and Jakub Opršal. 2019. The complexity of 3-colouring H -colourable graphs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS'19)*. IEEE, Baltimore, MD, USA, 1227–1239. doi:10.1109/FOCS.2019.00076
- [22] Andrei Krokhin and Jakub Opršal. 2022. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News* 9, 3 (2022), 30–59. doi:10.1145/3559736.3559740 arXiv:2208.13538
- [23] 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. doi:10.1137/20M1378223 arXiv:2003.11351
- [24] László Lovász. 1978. Kneser’s conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A* 25, 3 (1978), 319–324. doi:10.1016/0097-3165(78)90022-5
- [25] Jiří Matoušek. 2003. *Using the Borsuk-Ulam Theorem* (first ed.). Springer-Verlag Berlin Heidelberg. xii+214 pages. doi:10.1007/978-3-540-76649-0
- [26] 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:10.1090/cbms/091
- [27] Sebastian Meyer. 2024. A Dichotomy for Finite Abstract Simplicial Complexes. (2024). doi:10.48550/arXiv.2408.08199 arXiv:2408.08199 preprint.
- [28] Sebastian Meyer and Jakub Opršal. 2025. A topological proof of the Hell–Nešetřil dichotomy. In *Proceedings of the 2025 Annual ACM–SIAM Symposium on Discrete Algorithms, SODA 2025 (New Orleans, LA, USA, January 12–15, 2025)*, Yossi Azar and Debmalya Panigrahi (Eds.). SIAM, 4507–4519. doi:10.1137/1.9781611978322.154 arXiv:2409.12627
- [29] Thomas J. Schaefer. 1978. The complexity of satisfiability problems. In *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, Calif., 1978)*. ACM, New York, 216–226.
- [30] Patrick Schnider and Simon Weber. 2024. A Topological Version of Schaefer’s Dichotomy Theorem. In *40th International Symposium on Computational Geometry, SoCG 2024 (June 11–14, 2024, Athens, Greece) (LIPIcs, Vol. 293)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 77:1–77:16. doi:10.4230/LIPIcs.SoCG.2024.77
- [31] Tammo tom Dieck. 1987. *Transformation Groups*. De Gruyter, Berlin, New York. doi:10.1515/9783110858372
- [32] Marcin Wrochna and Stanislav Živný. 2020. Improved hardness for H -colourings of G -colourable graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*. SIAM, Salt Lake City, UT, USA, 1426–1435. doi:10.1137/1.9781611975994.86
- [33] Dmitriy Zhuk. 2020. A Proof of the CSP Dichotomy Conjecture. *J. ACM* 67, 5 (2020), 30:1–30:78. doi:10.1145/3402029
- [34] David Zuckerman. 2007. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.* 3 (2007), 103–128. doi:10.4086/toc.2007.v003a006

Received 2024-11-04; accepted 2025-02-01