Topological Methods in Discrete Geometry and Theoretical Computer Science

Measure Partitioning and Constraint Satisfaction Problems

by

Gianluca Tasinato

September, 2025

A thesis submitted to the
Graduate School
of the
Institute of Science and Technology Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Committee in charge:

Johannes Fink, Chair
Uli Wagner
Herbert Edelsbrunner
Eric Babson

Institute of Science and Technology Austria

The thesis of Gianluca Tasinato, titled *Topological Methods in Discrete Geometry and Theoretical Computer Science*, is approved by:

Supervisor: Uli Wagner, ISTA, Klosterneuburg, Austria

Signature: ________________________________

Committee Member: Herbert Edelsbrunner, ISTA, Klosterneuburg, Austria

Signature: ________________________________

Committee Member: Eric Babson, UC Davis, Davis, United States

Signature: ________________________________

Defense Chair: Johannes Fink, ISTA, Klosterneuburg, Austria

Signature: ________________________________

Signed page is on file

ISTA Thesis, ISSN: 2663-337X

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I accept full responsibility for the content and factual accuracy of this work, including the data and their analysis and presentation, and the text and citation of other work.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: ______________________________

Gianluca Tasinato
September, 2025

Signed page is on file

Abstract

This thesis investigates the interplay between algebraic and topological methods and combinatorial problems, focusing on approximate graph colourings and mass partitioning. The unifying theme throughout the dissertation is the use of continuous maps and symmetry constraints to extract combinatorial insights.

We first explore approximate graph colouring problems and more generally promise constraint satisfaction problems. Using tools from equivariant topology in combination with the general theory of polymorphism of a promise constraint satisfaction problem, we establish hardness for specific types of approximations.

In the second part, we address mass partitioning problems, where one seeks to divide geometric objects or measures in Euclidean space into parts of equal size using hyperplanes. Employing techniques from topological combinatorics (configuration space/test map setup and Borsuk–Ulam type theorems), we both obtain a new equipartitioning result in the and provide a fast algorithm for computing equipartitioning of point sets in 3D.

About the Author

Gianluca Tasinato completed his undergraduate studies at Pisa University in Italy where he obtained his BSc in Mathematics in 2017 and MSc in Mathematics in 2019 before joining ISTA in September 2019.

Gianluca joined Uli Wagner's research group in May 2021 and his main research area is topological combinatorics, with a focus on the use of (equivariant) topological methods in problems from discrete mathematics and theoretical computer science. During his PhD studies, Gianluca presented his research results at the Symposium on Theoretical Aspects of Computer Science (STACS) in Clermont-Ferrand in March 2024 and at the Symposium of Computational Geometry (SoCG) in Athens, in June 2024 and at the Symposium on Theory of Computing (STOC) in Prague in June 2025.

List of Collaborators and Publications

Sergey Avvakumov, Marek Filakovský, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. Hardness of 4-colouring G-colourable graphs. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, Stoc '25, pages 72–83, Prague, Czechia and New York, NY, USA, 2025. Association for Computing Machinery

(Reused in Chapters 2 to 5).

Boris Aronov, Abdul Basit, Indu Ramesh, Gianluca Tasinato, and Uli Wagner. Eight-partitioning points in 3D, and efficiently too. In *40th International Symposium on Computational Geometry*, volume 293 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 8, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024

(Reused in Chapters 6 to 8).

Marek Filakovský, Tamio-Vesa Nakajima, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs. *LIPIcs, Volume 289, STACS 2024*, 289:34:1–34:19, 2024

(Reused in Chapters 2 to 5).

Table of Contents

List of Figures

CHAPTER 1

Introduction

Since its inception, topology has been used as a powerful tool to approach problems in a variety of different fields; ranging from analysis to algebraic geometry and anything in between. Furthermore, algebraic topology has developed into a vast and rich research field in its own right, with powerful techniques suitable to study a wide breath of problems throughout mathematics. Despite the early influence of combinatorics in the development of algebraic topology, not many ideas during the early days of algebraic topology were employed in discrete mathematics.

The focus of this work will be on two different applications of tools from topology to different areas of discrete mathematics and combinatorics: graph theory and partitioning of geometric objects (point sets, convex bodies or mass distributions).

Arguably, the catalyst that jump-started the use of algebraic topology in graph theory was the seminal work of Lovász on the chromatic number of Kneser graphs [Lov78]. Since then, a wide variety of results have been established via topology,including topological lower bounds for the chromatic number of graphs and hypergraphs (see , e.g. [MZ03, BK07]), evasiveness of monotone graph properties [KSS84, Yao88], topological lower bounds and impossibility results for distributed computing (see e.g. [HKR14]).

Additionally, the construction of the *neighbourhood complex* at the underlies Lovász' proof has been generalized and studied extensively (see e.g [BK07, Mat17, DS23]). Here, our focus will be on how to use topological tools, in particular the concept of homomorphism complex, in the study of approximate graph colouring problem and other promise constraint satisfaction problems; see chapter 2, for a more detailed introduction and further background on this class of problems.

Furthermore, topology has been a foundational tool for the study of mass partitioning problems, see, e.g. the survey [RS21]. In particular, the well known *Ham-Sandwich* theorem was proven (by Banach in 1938 in the three dimensional case and Stone and Tuckey in the general case) by a direct application of the Borsuk-Ulam theorem (see [RS21] for more details on the early history of the Ham-Sandwich theorem).

Since then, a lot of effort has been invested in studying how partitions of the standard Euclidian space $\mathbb{R}^d$ split families of measures (in the continuous setting) or points (in the discrete setting). In Chapter 6, we will give a more thorough introduction to this more general class of questions and present our contributions.

1.1 Overview

We briefly describe the content of the remaining chapters. This work is roughly split in two general topics; focussing first on topological methods in promise constraint satisfaction problems, then on the mass equipartitioning problem both from a theoretical and an algorithmic point of view.

In Chapter 2, we briefly discuss approximate graph colouring problems and the more general context of promise constraint satisfaction problem and give a short presentation of the needed notions.

In Chapter 3, we will give an high level presentation of the proof structure for the main complexity results claimed in Chapter 2. Chapters 4 and 5 will be dedicated to filling in the more technical details needed for the argument; the presentation of this topic is based on [AFO+25, FNO+24].

Chapter 6 is dedicated to introducing mass partitioning problems and presenting our contribution in this context. Chapter 7 will discuss the existence result we obtained while Chapter 8 will focus on presenting the algorithm for finding a eight partition of points in the 3D euclidian space. The presentation is based on [ABR+24].

CHAPTER 2

Promise Constraint Satisfaction Problems

2.1 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 [Kar72]. 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 $\mathsf{P} = \mathsf{NP}$ ([Zuc07]).

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 [KLS00] proved that this is NP-hard for $k = 4$ (see also [GK04, BG16]), and only quite recently Bulín, Krokhin, and Opršal [BKO19] 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 [KTY24], 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, Krokhin, and Opršal [BKO19]), and for $c \geq 5$ and $k = \binom{c}{\lfloor c/2 \rfloor} - 1$ (Wrochna and Živný [WŽ20]). 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 [DMR09], Guruswami and Sandeep [GS20], and Braverman, Khot, Lifshitz, and Minzer [BKLM22].

One of the main question studied in this thesis is a generalisation of this question. A *graph homomorphism* $f : G \to H$ between two graphs is a map

$f\colon V(G) \to V(H)$ between the vertex sets that preserves edges, i.e., $(u, v) \in E(G)$ implies $(f(u), f(v)) \in E(H)$; we write $G \to H$ if such a homomorphism exists. Throughout this work, 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 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 $f\colon G \to H$ 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 [HN90] 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 $G \to H$. 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 PCSP(G, H): Given an input graph I, output YES if $I \to G$ and NO if $I \not\to H$ (no output is required if neither is the case). Brakensiek and Guruswami conjectured [BG21, Conjecture 1.2] that 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).

As a first step towards the Brakensiek–Guruswami conjecture, Krokhin and Opršal [KO19] showed that 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 [WŽ20] (and in the joint journal version [KOWŽ23]). We extend this to 4-colouring ([AFO+25]):

Theorem 2.1.1. *Let G be a non-bipartite 4-colourable graph. Then* PCSP(G, K_4) *is NP-hard.*

More generally, 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 general formulation of the constraint satisfaction problem is then as follows (see Section 2.2.1 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 [Bul17] and Zhuk [Zhu20], which asserts that for every finite relational structure $\mathbf{A}$, CSP$(\mathbf{A})$ is either NP-complete, or solvable in polynomial time.

The framework of CSPs can be extended to *promise constraint satisfaction problems* (*PCSPs*), which include approximate graph colouring. PCSPs were first introduced by Austrin, Guruswami, and Håstad [AGH17], and the general theory of these problems was further developed by Brakensiek and Guruswami [BG21], and by Barto, Bulín, Krokhin, and Opršal [BBKO21].

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 $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$.

PCSPs encapsulate a wide variety of problems, including versions of hypergraph colouring studied by Dinur, Regev, and Smyth [DRS05] and Brakensiek and Guruswami [BG16]. A variant of hypergraph colouring that is closely connected to approximate graph colouring and generalises (monotone[1]) *1-in-3-SAT* is *linearly ordered (LO) hypergraph colouring*. A linearly ordered k-colouring of a hypergraph H is an assignment of the colours $[k] = \{1, \ldots, k\}$ to the vertices of H such that, for every hyperedge, the maximal colour assigned to elements of that hyperedge occurs exactly once. Note that for graphs, linearly ordered colouring is the same as 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 work, 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 (see Section 2.2.1). In particular, 1-in-3-SAT corresponds to $\mathrm{CSP}(\mathbf{LO}_2)$.[2]

The promise version of LO hypergraph colouring was introduced by Barto, Battistelli, and Berg [BBB21], 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. [BBB21] 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 $k \geq c \geq 2$ (see Reference [BBB21, Conjecture 27]). Subsequently, the following conjecture emerged and circulated as folklore (first formally stated by Nakajima and Živný [NŽ23a]): $\mathrm{PCSP}(\mathbf{LO}_2, \mathbf{B})$ is either solved by the *affine integer programming relaxation*, or NP-hard. (See Ciardo, Kozik, Krokhin, Nakajima, and Živný [CKK+23] for recent progress in this direction.)

Promise LO hypergraph colouring was further studied by Nakajima and Živný [NŽ23b], 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 hyper-

[1]In the present work, 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]Observe that this notion corresponds exactly to the notion of *Unique Maximum Colouring* [CKP13] — however, in the context of promise CSPs these were first identified by [BBB21], and thus we follow their terminology. We thank Dömötör Pálvölgyi for informing us that LO colourings were also studied under this name.

graphs with a superconstant number of colours, by adapting methods used for similar algorithms for approximate graph colouring, e.g., [Blu94, KT17]. The number of colours used was then reduced by Håstad, Martinsson, Nakajima and Živný [HMNŽ24] to $\log_2(n)$. In the other direction, the NP-hardness of $\mathrm{PCSP}(\mathbf{LO}_k, \mathbf{LO}_c)$ for $4 \le k \le c$ follows relatively easily from the NP-hardness of the approximate graph colouring $\mathrm{PCSP}(K_{k-1}, K_{c-1})$, as was observed by Nakajima and Živný and by Austrin (personal communications).[3]

Our result within this body of work is the following, which cannot be obtained using these arguments.

Theorem 2.1.2. $\mathrm{PCSP}(\mathbf{LO}_3, \mathbf{LO}_4)$ *is NP-complete.*

The proof of both Theorems 2.1.1 and 2.1.2 are structurally very similar and mirror the argument in [KOWŽ23], however they require both more sophisticated tools from algebraic topology and significantly different combinatorial arguments to obtain the claimed results. The outline of the proofs, as well as a more precise comparison with the topological approach of [KOWŽ23] will be presented in detail in Chapter 3.

2.2 Preliminaries

We use the notation $[n]$ for the n-element set $\{1, \dots, n\}$. 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.

2.2.1 Promise CSPs

We start by recalling some fundamental notions from the theory of promise constraint satisfaction problems, following the presentation of [BBKO21] and [KO22].

A *relational structure* is a tuple $\mathbf{A} = (A; R_1^{\mathbf{A}}, \dots, R_k^{\mathbf{A}})$, where A is a set, and $R_i^{\mathbf{A}} \subseteq A^{\mathrm{ar}(R_i)}$ is a relation of arity $\mathrm{ar}(R_i)$. The *signature* of $\mathbf{A}$ is the tuple $(\mathrm{ar}(R_1), \dots, \mathrm{ar}(R_k))$. For two relational structures $\mathbf{A} = (A; R_1^{\mathbf{A}}, \dots, R_k^{\mathbf{A}})$ and $\mathbf{B} = (B; R_1^{\mathbf{B}}, \dots, R_k^{\mathbf{B}})$ with the same signature, a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$, denoted $h\colon \mathbf{A} \to \mathbf{B}$, 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 $\hom(\mathbf{A}, \mathbf{B})$.

Note that graphs fit in the general framework of relational structures. In fact, a graph can be interpreted, in this language, as a relational structure with a single relation of arity 2 (that is symmetric since we focus only on unoriented graphs).

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

The other type of relational structures we will focus on is structures with a single ternary relation R: pairs $(A; R^{\mathbf{A}})$ with $R^{\mathbf{A}} \subseteq A^3$. Moreover, most structures in this context 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 2.2.1 (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 $\mathbf{I}$ with the same signature as $\mathbf{A}$ and $\mathbf{B}$, we are promised that $\mathbf{I} \to \mathbf{A}$, and we are tasked with finding a homomorphism $h\colon \mathbf{I} \to \mathbf{B}$.
- In the *decision* version of the problem, we are given a relational structure $\mathbf{I}$, and we must answer *Yes* if $\mathbf{I} \to \mathbf{A}$, and *No* if $\mathbf{I} \not\to \mathbf{B}$. (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 $(2, 2, 1)$.

2.2.2 Polymorphisms and a hardness condition

Our proof of Theorem 2.1.1 uses a hardness criterion (Theorem 2.2.1 below) obtained as part of a general algebraic theory of PCSPs developed by [BBKO21], which we will briefly review.

Definition 2.2.2. 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 $\mathrm{pol}^{(n)}(\mathbf{A}, \mathbf{B})$.[4]

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

Concretely, in the context of graphs, a n-ary polymorphism from G to H is a map $f\colon V(G)^n \to V(H)$ such that $(f(u_1,\dots,u_n), f(v_1,\dots,v_n)) \in E(H)$ whenever $(u_1,v_1),\dots,(u_n,v_n) \in E(G)$.

Similarly, 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

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

Let $\pi\colon [n] \to [m]$, and let A and B be sets. The *π-minor* of a function $f\colon A^n \to B$ is the function $f^\pi\colon A^m \to B$ 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 $x\colon [n] \to A$, then $f^\pi(x) = f(x \circ \pi)$). A subset of the set of all functions $\{f\colon A^n \to B, n > 0\}$ that is non-empty and closed under taking minors is called a *function minion*. For example, it is easy to see that $\mathrm{pol}(\mathbf{A},\mathbf{B})$ has this property whenever $\mathbf{A}$ and $\mathbf{B}$ are relational structures such that $\mathbf{A} \to \mathbf{B}$. Abstracting from this, we arrive at the following notion:

Definition 2.2.3. An *(abstract) minion* $\mathscr{M}$ is a collection of non-empty sets $\mathscr{M}^{(n)}$, where $n > 0$ is an integer, and mappings

$$\pi^{\mathscr{M}}\colon \mathscr{M}^{(n)} \to \mathscr{M}^{(m)},$$

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

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.

Given a minion $\mathscr{M}$, an element $f \in \mathscr{M}^{(n)}$ is said to have *essential arity at most k* if it is a minor of an element $g \in \mathscr{M}^{(k)}$. If there is a bound N, such that every element of $\mathscr{M}$ has essential arity at most N, $\mathscr{M}$ is said to have *bounded essential arity*. An element $f \in \mathscr{M}^{(n)}$ is *constant* if all its minors coincide, i.e., $f^\pi = f^\sigma$ for all $m > 0$ and $\pi,\sigma\colon [n] \to [m]$. For example, in function minions, being constant coincides with the usual notion of being a constant function, and if a function $f\colon A^n \to B$ 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. In the proof of hardness for $\mathrm{PCSP}(C_\ell, K_4)$ we rely on the following hardness criterion.[5]

Theorem 2.2.1 ([BBKO21, Proposition 5.14]). *Let G and H be two graphs such that $G \to H$. If there exists a minion homomorphism*

$$\xi\colon \mathrm{pol}(G,H) \to \mathscr{B}$$

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

[5] Here stated in the graph case, but the result holds for general relational structures.

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 $i \leq n$, the i-th n-ary *projection* on A is the function $p_i : A^n \to A$ 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 $\operatorname{pol}(\mathbf{LO}_2, \mathbf{LO}_2)$.

It is clear that any function in a function minion isomorphic to $\mathscr{P}$ has essential arity 1, therefore we immediately get the following corollary:

Theorem 2.2.2 ([BBKO21, corollary of Theorem 3.1]). *For every promise template* $(\mathbf{A}, \mathbf{B})$ *such that there is a minion homomorphism* $\xi\colon \operatorname{pol}(\mathbf{A}, \mathbf{B}) \to \mathscr{P}$, $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ *is NP-complete.*

This will be the hardness criterion we use to prove Theorem 2.1.2.

2.3 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 [Lov78]. We refer the reader to [Hat02] and [Mat08] for accessible general introductions to algebraic topology and topological combinatorics, respectively, and to [Koz08] 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 [Fri12] 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 2.3.1. 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 $\alpha\colon \{0,\dots,m\} \to \{0,\dots,n\}$, there is a map $\alpha^X\colon X_n \to X_m$, such that $1^X_{\{0,\dots,n\}} = 1_{X_n}$ 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 [Fri12, 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 $\alpha\colon \{0,\dots,k\} \to \{0,\dots,m\}$ 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;[6] however, degenerate simplices play an important role in specifying the glueings and the combinatorial data keeps track of them. All simplicial sets we will use 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* $f\colon X \to Y$ between simplicial sets is a collection of maps $f_n\colon X_n \to Y_n$, $n > 0$, such that $f_m \circ \alpha^X = \alpha^Y \circ f_n$ for all monotone maps $\alpha\colon \{0,\dots,m\} \to \{0,\dots,n\}$. Every simplicial map $f\colon X \to Y$ defines a continuous map $|f|\colon |X| \to |Y|$. An *isomorphism* of simplicial sets X and Y is a simplicial map $f\colon X \to Y$ with a simplicial inverse $g\colon Y \to X$ (f_n is inverse to g_n for all $n > 0$).

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 $\alpha^{X \times Y}(\sigma, \tau) = (\alpha^X(\sigma), \alpha^X(\tau))$. 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 [Fri12, Theorem 5.2] for a general statement). The nth 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 thesis (graphs, relational structures, simplicial sets, topological spaces, etc.) have a natural symmetry given by an action of a cyclic group (either $\mathbb{Z}_2$ when working with graphs or $\mathbb{Z}_3$ when dealing with the arity 3 relational structures), which is described by a structure-preserving automorphism (of order 2 or 3 depending on the

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

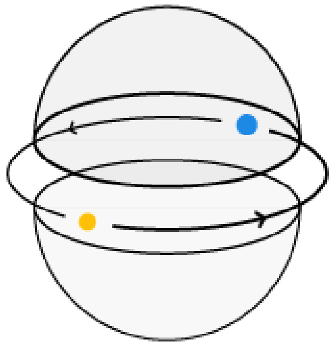

Figure 2.1: The simplicial set Σ^2

context). For instance, a $\mathbb{Z}_2$-action on a simplicial set X is given by a simplicial map $\nu\colon X \to X$ 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 actions of $\mathbb{Z}_2$ or $\mathbb{Z}_3$ simply means that ν has no fixed points. If (X, ν_X) and (Y, ν_Y) are simplicial sets with $\mathbb{Z}_k$-actions, then a simplicial map $f\colon X \to Y$ is called *equivariant* if it preserves the $\mathbb{Z}_k$-symmetry, i.e., $f \circ \nu_X = \nu_Y \circ f$.

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

Relational simplicial sets Most simplicial sets we take into consideration 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 $X_n \subset (X_0)^{n+1}$ is an $(n+1)$-ary relation, i.e., every n-simplex of X is an ordered $(n+1)$-tuple $[u_0, \ldots, u_n]$ 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 $\alpha\colon \{0, \ldots, m\} \to \{0, \ldots, n\}$, the map α^X is defined by $\alpha^X([u_0, \ldots, u_n]) = [u_{\alpha(0)}, \ldots, u_{\alpha(m)}]$. To get a simplicial set this way, the collection of relations X_n, $n > 0$, 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.

Example 2.3.1 ($\mathbb{Z}_2$-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 $\Sigma^2_0 = \{{\color{orange}\bullet}, {\color{blue}\bullet}\}$ (which we think of as a pair of antipodal points in S^2), and Σ^2_n is the set of all $(n+1)$-tuples of ${\color{orange}\bullet}$ and ${\color{blue}\bullet}$'s with at most 2 alternations. Thus, e.g., $[{\color{orange}\bullet}, {\color{blue}\bullet}, {\color{blue}\bullet}, {\color{orange}\bullet}]$ is a 3-simplex of Σ^2, but $[{\color{blue}\bullet}, {\color{orange}\bullet}, {\color{blue}\bullet}, {\color{orange}\bullet}]$ 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 $\Sigma^0 \subset \Sigma^1 \subset \Sigma^2 \subset \ldots$, such that Σ^k (whose simplices are tuples with entries in $\{{\color{orange}\bullet}, {\color{blue}\bullet}\}$ 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 ${\color{blue}\bullet}, {\color{yellow}\bullet}$; Σ^1 additionally has two non-degenerate 1-simplices $[{\color{blue}\bullet}, {\color{yellow}\bullet}]$ and $[{\color{yellow}\bullet}, {\color{blue}\bullet}]$ 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 $[{\color{yellow}\bullet}, {\color{blue}\bullet}, {\color{yellow}\bullet}]$ and $[{\color{blue}\bullet}, {\color{yellow}\bullet}, {\color{blue}\bullet}]$ which corresponds to glueing the northern and southern hemisphere, respectively (see Figure 2.1); Σ^3 adds two non-degenerate 3-simplices; etc.

Observation 1. If X is a (relational)[7] simplicial set then a simplicial map $X \to \Sigma^2$ 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 $[{\color{blue}\bullet}, {\color{yellow}\bullet}, {\color{blue}\bullet}, {\color{yellow}\bullet}]$ or $[{\color{yellow}\bullet}, {\color{blue}\bullet}, {\color{yellow}\bullet}, {\color{blue}\bullet}]$). 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.

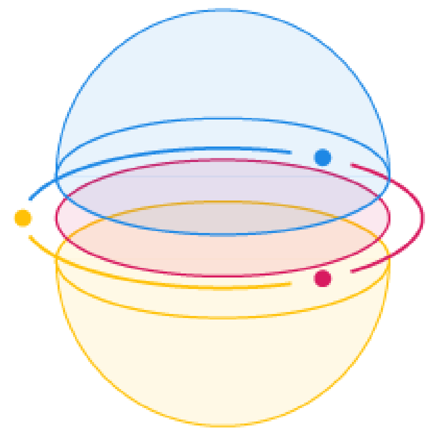

Figure 2.2: The simplicial set Θ^2.

Example 2.3.2 ($\mathbb{Z}_3$-space Θ^2). The standard sphere S^2 does not admit any free $\mathbb{Z}_3$ action on itself. Therefore we need to define a new space that will play the same role when dealing with $\mathbb{Z}_3$: Such space is the relational simplicial set Θ^2 together with a natural $\mathbb{Z}_3$-action.

The vertex set of Θ^2 is $\Theta^2_0 = \{{\color{yellow}\bullet}, {\color{blue}\bullet}, {\color{red}\bullet}\}$ (which we can think as an orbit of points that are moved around by a cyclic permutation ${\color{yellow}\bullet} \mapsto {\color{red}\bullet}$, ${\color{red}\bullet} \mapsto {\color{blue}\bullet}$), and Θ^2_n is the set of all $(n+1)$-tuples of ${\color{yellow}\bullet}$, ${\color{blue}\bullet}$'s and ${\color{red}\bullet}$'s with at most one cycle. Thus, e.g., $[{\color{yellow}\bullet}, {\color{blue}\bullet}, {\color{red}\bullet}, {\color{blue}\bullet}]$ is a 3-simplex of Θ^2, but $[{\color{blue}\bullet}, {\color{yellow}\bullet}, {\color{red}\bullet}, {\color{blue}\bullet}]$ is not. The $\mathbb{Z}_3$-action on Θ^2 is given by the cyclic permutation previously described.

The simplicial set Θ^2 has a very similar property as Σ^2 regarding (equivariant) map with itself as a target:

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

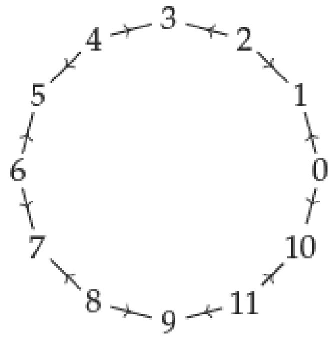

Figure 2.3: The simplicial set Γ_{12}.

Observation 2. If X is a (relational)[8] simplicial set then a simplicial map $X \to \Theta^2$ is completely described by a 3-colouring of the vertex set X_0 with colours yellow, blue or red. Conversely, a 3-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 has more than a cycle (e.g. $f(u_0) = f(u_3)$ and $f(u_1), f(u_2)$ are the other two colours). Moreover, if $\mathbb{Z}_3$-acts on X by a simplicial automorphism ν, then such a 3-colouring defines an equivariant map if and only if u, $\nu(u)$ and $\nu^2(u)$ are coloured in the correct cyclic order for every vertex u of X; e.g if u is coloured red, then $\nu(u)$ is blue and $\nu^2(u)$ is yellow.

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 $\alpha : \{0, \dots, m\} \to \{0, \dots, n\}$, $\alpha^{\Delta(P)}[u_0, \dots, u_n] = [u_{\alpha(0)}, \dots, u_{\alpha(m)}]$ 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 $f : P \to Q$ between posets naturally extends componentwise to chains and hence to a simplicial map $f : \Delta(P) \to \Delta(Q)$ between order complexes.

Example 2.3.3. 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 $a - b = \pm 1 \mod L$. 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 2.3 (as a digraph, it is a cycle of length L with edges oriented alternatingly). Moreover, the map $\mathbb{Z}_L \to \mathbb{Z}_L$, $x \mapsto x + L/2$ defines a simplicial involution $\Gamma_L \to \Gamma_L$ that corresponds to the antipodal involution on S^1.

The previous example can be modified to build a simplicial triangulation of S^1 that is equivariant with respect to the standard $\mathbb{Z}_3$-action instead:

[8]The claim is once again true for arbitrary simplicial set X; it is only required that Θ^2 is relational.

Example 2.3.4. Let L be a positive integer divisible by 6. Define a partial order $\preccurlyeq$ on $\mathbb{Z}_L = \{0, 1, \dots, L-1\}$ as in example 2.3.3 by $a \prec b$ if and only if a is even, b is odd, and $a - b = \pm 1 \mod L$. We define the simplicial set Ξ_L as the order complex of this poset,

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

The simplicial set Ξ_L is again a triangulation of S^1.

The difference is now that the map $\mathbb{Z}_L \to \mathbb{Z}_L$, $x \mapsto x + L/3$ defines a simplicial automorphism $\Xi_L \to \Xi_L$ of order 3 that corresponds to the rotation by $2\pi/3$ on S^1.

Note that the only difference in the two construction is the order of the simplicial automorphism and thus the group acting on S^1: $\mathbb{Z}_2$ in the case of Γ_L, $\mathbb{Z}_3$ in the case of Ξ_L.

Example 2.3.5. Let n be a positive integer. Define a partial order $\preccurlyeq$ on $[n] \times \mathbb{Z}_3$ by $(i, \alpha) \prec (j, \beta)$ if and only if $i < j$. Define the simplicial set L_n as the order complex of this poset,

$$L_n := \Delta([n] \times \mathbb{Z}_3, \preccurlyeq)$$

Furthermore, the map $[n] \times \mathbb{Z}_3 \to [n] \times \mathbb{Z}_3$, $(i, \alpha) \mapsto (i, \alpha + 1)$ defines a simplicial automorphism of order 3 on L_n that is free.

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, both $\Gamma_L^n = \Gamma_L \times \dots \times \Gamma_L$ and $\Xi_L^n = \Xi_L \times \dots \times \Xi_L$ are triangulations of the n-dimensional torus $T^n = S^1 \times \dots \times S^1$. Note also that in both cases the vertices are n-tuples $u \in \mathbb{Z}_L^n$, and k-simplices are $(k+1)$-tuples of vertices $[u_0, \dots, u_k]$ such that u_{i+1} is obtained from u_i by choosing a subset of coordinates of u_i all that are even and changing each of them by ± 1 modulo L.

Homomorphism complexes Given two relational structures $\mathbf{A}$ and $\mathbf{B}$, the homomorphism complex $\mathrm{Hom}(\mathbf{A}, \mathbf{B})$ is a simplicial set capturing the structure of all homomorphisms $\mathbf{A} \to \mathbf{B}$. Following [Mat08, Section 5.9], we define homomorphism complexes as order complexes of the poset of *multihomomorphisms* from $\mathbf{A}$ to $\mathbf{B}$.[9] By definition, a *multihomomorphism* is a function $f \colon A \to 2^B \setminus \{\emptyset\}$ such that, for all relational symbol R and all tuples $(a_1, \dots, a_k) \in R^A$, we have that

$$f(a_1) \times \dots \times f(a_k) \subseteq R^B.$$

We denote the set of all multihomomorphisms by $\mathrm{mhom}(\mathbf{A}, \mathbf{B})$. Multihomomorphisms are partially ordered by component-wise inclusion: $f \leq g$ if and only if $f(a) \subseteq g(a)$ for all $a \in A$.

[9]We remark that in [Mat08] 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.

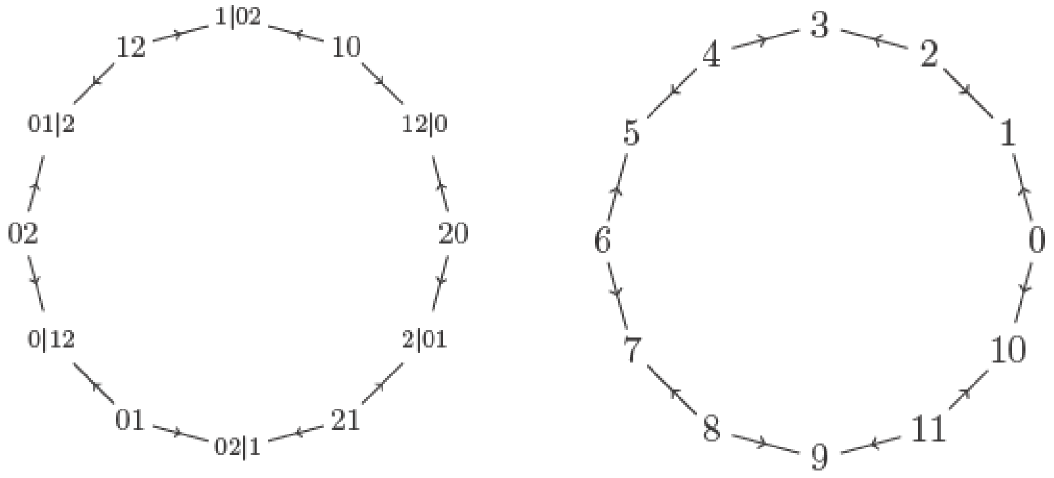

Figure 2.4: The simplicial sets $\mathrm{Hom}(K_2, C_3)$ and Γ_{12} (see Example 2.3.6).

Definition 2.3.2. Let $\mathbf{A}$ and $\mathbf{B}$ be relational structures. The *homomorphism complex* $\mathrm{Hom}(\mathbf{A}, \mathbf{B})$ is the order complex $\Delta(\mathrm{mhom}(\mathbf{A}, \mathbf{B}), \leq)$ of the poset of multihomomorphisms.

Multihomomorphisms can be composed in a natural way: if $f \in \mathrm{mhom}(\mathbf{A}, \mathbf{B})$ 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 $\mathbf{C}$. In particular, every homomorphism $f\colon \mathbf{B} \to \mathbf{C}$ induces a simplicial map $f_*\colon \mathrm{Hom}(\mathbf{A}, \mathbf{B}) \to \mathrm{Hom}(\mathbf{A}, \mathbf{C})$ defined on vertices by mapping a multihomomorphism $m \in \mathrm{mhom}(\mathbf{A}, \mathbf{B})$ to the composition $f \circ m$.

In what follows, we will focus on two special cases: graphs and symmetric relations of arity 3.

Graph homomorphism complex: When dealing with graphs, we will focus on $\mathrm{Hom}(K_2, G)$; a common tool in the study of graph colourings. Note that a multimorphism 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 $m(1)$ 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 induces an action on multihomomorphisms $m\colon K_2 \to G$, namely swapping the two sets $m(1)$ and $m(2)$, which in turn induces a $\mathbb{Z}_2$-action on the simplicial set $\mathrm{Hom}(K_2, G)$; this action is free provided G has no loops. Moreover, it is easy to check that for every graph homomorphism $f\colon G \to H$, the induced simplicial map $f_*\colon \mathrm{Hom}(K_2, G) \to \mathrm{Hom}(K_2, H)$ is equivariant.

The following two examples will play an important role in the proof of Theorem 3.1.1.

Example 2.3.6. For every odd integer $\ell \geq 3$, $\mathrm{Hom}(K_2, C_\ell)$ 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 2.3.7. The simplicial set $\mathrm{Hom}(K_2, K_4)$ is a triangulation of a sphere S^2; it is depicted in Figure 2.5 which shows two hemispheres of this sphere

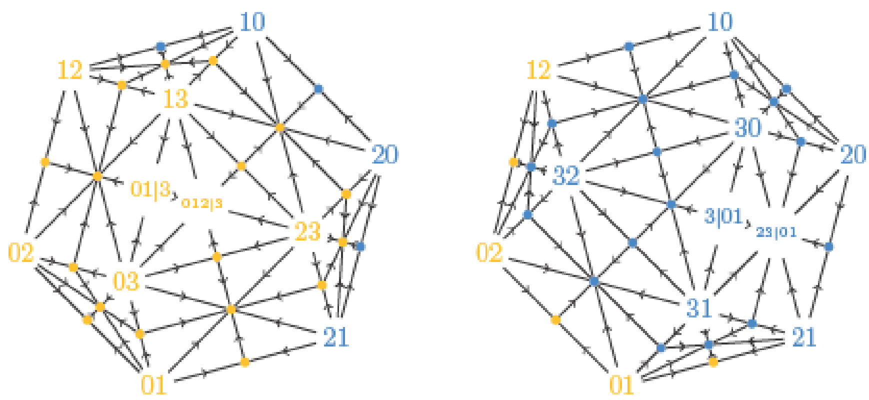

Figure 2.5: The simplicial set $\mathrm{Hom}(K_2, K_4)$; see also Example 2.3.7.

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 2.3.1. *There exists an equivariant simplicial map*

$$t\colon \mathrm{Hom}(K_2, K_4) \to \Sigma^2.$$

Proof. By Observation 1, such a map is given by a suitable 2-colouring of the vertices of $\mathrm{Hom}(K_2, K_4)$. One suitable 2-colouring is depicted in Figure 2.5. □

Homomorphism complex of relations of arity 3: In this context, we work instead with the homomorphism complex $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$ where the role of K_2 is played by $\mathbf{R}_3$, the structure with 3 elements and all *rainbow tuples*, i.e., tuples (a, b, c) such that a, b, and c are pairwise distinct.

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 $m(1) \times m(2) \times m(3) \subseteq R^{\mathbf{A}}$.

Moreover, 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 it extends naturally to an action of $\mathbb{Z}_3$ on $\mathrm{Hom}(\mathbf{R}_3, \mathbf{A})$.

It is not hard to show that the induced 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.

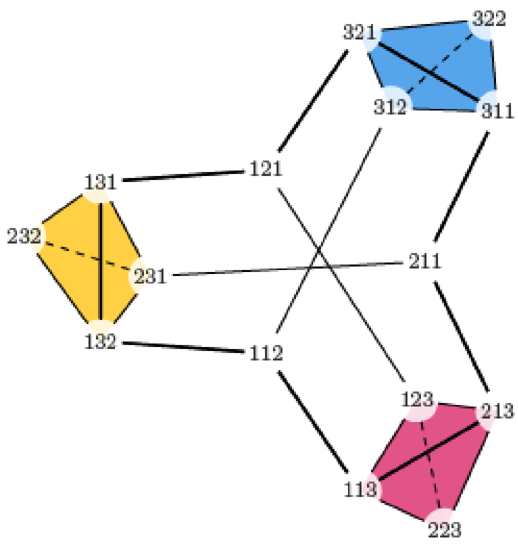

Figure 2.6: The simplicial set $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and the embedded Ξ_{18}.

For every homomorphism $f\colon \mathbf{A} \to \mathbf{B}$, the induced simplicial function $f_*\colon \mathrm{Hom}(\mathbf{R}_3, \mathbf{A}) \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{B})$ (defined on vertices by mapping multi-homomorphism 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.

Unfortunately, the simplicial sets $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ and $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ are not as easily described as it was the case for graphs; however it is still possible to show some important properties.

Lemma 2.3.2. *The simplicial set* $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ *contains an equivariant simplicial subspace isomorphic to* Ξ_{18}.

Proof. Consider the cycle in $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ given by the vertices[10]

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

This cycle is highlighted in Figure 2.6. Observe that it is invariant under the $\mathbb{Z}_3$ action. Thus we see that there is the desired $\mathbb{Z}_3$-equivariant embedding $\Xi_{18} \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$. □

In general, while the equivariant homotopy type of $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_k)$ can be hard to understand, for our purposes it is enough to show it can be (equivariantly) mapped to a sufficiently simple space.

Lemma 2.3.3. *There is a* $\mathbb{Z}_3$*-map* $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4) \to \Theta^2$.

[10] 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.

Proof. In order to provide the required equivariant map, we will first provide an equivariant simplicial map to the barycentric subdivision of L_3, then we build an equivariant map $L_3 \to \Theta^2$.

In general, the definition of barycentric subdivision for general simplicial set requires some technical care (e.g., see [May92]). However, in the case of order complexes it is possible to simplify it: The set $\mathcal{F}_X$ of non degenerate simplices in an order complex X is by itself a poset ordered by inclusion. Then, the barycentric subdivision of X is just the order complex of $\mathcal{F}_X$, (see [Mat08, Definition 1.7.2]).

Denote by $F := \mathcal{F}_{L_3}$ the poset of non-empty simplices of L_3 ordered by inclusion. To obtain the first of the claimed simplicial maps it suffices to give a monotone equivariant map from $\mathrm{mhom}(\mathbf{R}_3, \mathbf{LO}_4)$ to F.

First, we define an auxiliary function $h\colon \mathrm{hom}(\mathbf{R}_3, \mathbf{LO}_4) \to [3] \times \mathbb{Z}_3$ 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 $\phi\colon \mathrm{mhom}(\mathbf{R}_3, \mathbf{LO}_4) \to F$ 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 $m \in \mathrm{mhom}(\mathbf{R}_3, \mathbf{LO}_4)$. 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.

The second equivariant map is much easier to explicitly build: By Observation 2, it is enough to provide an equivariant 3-colouring that avoids cyclic colourings. One such colouring is shown in Figure 2.7. □

Homotopy Two continuous maps $f, g\colon X \to Y$ between topological spaces are *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 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 $f\colon X \to Y$ and $g\colon Y \to X$ such that $f \circ g \sim 1_Y$ and $g \circ f \sim 1_X$.

These notions naturally generalize to the setting of spaces with group actions: Two equivariant maps $f, g\colon X \to Y$ between spaces with $\mathbb{Z}_k$-actions are

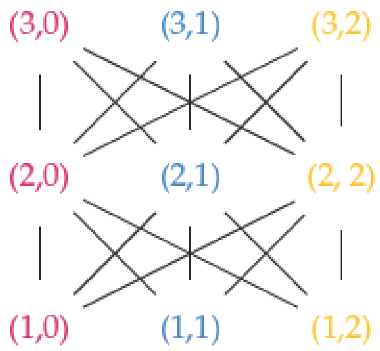

Figure 2.7: Equivariant colouring of L_3.

equivariantly homotopic, denoted by $f \sim_{\mathbb{Z}_k} g$, if there exists an *equivariant homotopy* between them, i.e., a homotopy $h\colon X \times [0,1] \to Y$ such that all maps $h(\cdot, t)\colon X \to Y$ are equivariant. We denote by $[X,Y]_{\mathbb{Z}_k}$ the set of all equivariant maps $X \to Y$ up to equivariant homotopy, i.e.,

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

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

CHAPTER 3

Outline of the arguments

Comparison with earlier work The topological approach in [KOWŽ23] for proving hardness of PCSP(C_ℓ, 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 $\mathbb{Z}^n \to \mathbb{Z}$ of the form $(x_1, \dots, x_n) \mapsto \sum_i \alpha_i x_i$, where $a_i \in \mathbb{Z}$ and $\sum_i a_i \equiv 1 \mod 2$ (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 $\sum_i |a_i|$ of absolute values of coefficients in such a map can be read of as the winding number of a simplicial map $\Gamma_{4\ell} \to \Sigma^1$, hence $O(\ell)$.

By contrast, the more careful counting argument required in the case of PCSP(C_ℓ, K_4), although elementary in hindsight, was elusive for several years.

At the same time, the case of PCSP($\mathbf{LO}_3\mathbf{LO}_4$) requires a relatively simple observation on the total structure of pol($\mathbf{LO}_3$, $\mathbf{LO}_4$) and thus it was the context where the more advanced topological ideas first bore fruit.

3.1 Outline - Hardness of PCSP(G, K_4)

We present a detailed overview of the proof of Theorem 2.1.1. Every non-bipartite, loopless graph G contains a cycle C_ℓ of odd length $\ell \geq 3$. In particular, there exists a homomorphism $C_\ell \to G$, hence every C_ℓ-colourable graph is G-colourable. This yields a trivial reduction from PCSP(C_ℓ, K_4) to PCSP(G, K_4). Thus, Theorem 2.1.1 follows from the following:

Theorem 3.1.1. *For all odd integers $\ell \geq 3$, the decision problem* PCSP(C_ℓ, K_4) *is NP-hard.*

We will prove this using Theorem 2.2.1; to this end, we need to construct a minion homomorphism from $\mathrm{pol}(C_\ell, K_4)$ to a minion $\mathscr{B}$ that contains no constant and is of bounded essential arity.

Informally speaking, as mentioned in Section 2.2.2, 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 $C_\ell^n \to K_4$, $n > 0$, 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 $T^n \to S^2$. 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 $\mathrm{pol}(C_\ell, K_4)$ to a minion $\mathscr{Z}_2$ (defined precisely below). The fact that all maps and homotopies are equivariant ensures that the minion $\mathscr{Z}_2$ 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 of all equivariant simplicial maps $X^n \to Y$, $n > 0$, is closed under taking minors, i.e., it forms a minion, which we denote by $\mathrm{spol}(X, Y)$ (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 $f\colon C_\ell^n \to K_4$ an equivariant simplicial map $\mu(f)\colon \Gamma_{4\ell}^n \to \Sigma^2$, where $\Gamma_{4\ell}$ and Σ^2 are the simplicial sets described in Examples 2.3.3 and 2.3.1, respectively. The simplicial map $\mu(f)$ is defined as a composition $t \circ f_* \circ \iota_n$:

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

where $f_*\colon\ \mathrm{Hom}(K_2, C_\ell^n) \to \mathrm{Hom}(K_2, K_4)$ is the simplicial map induced by f, $t\colon\ \mathrm{Hom}(K_2, K_4) \to \Sigma^2$ is the simplicial map from Lemma 2.3.1, the isomorphism $\Gamma_{4\ell}^n \cong \mathrm{Hom}(K_2, C_\ell)^n$ is given by the isomorphism from Example 2.3.6, and the simplicial map ι_n is given by the special case $G = C_\ell$ of the following fact:

Lemma 3.1.1. *For every graph G and $n \geq 1$, there is an equivariant simplicial map*

$$\iota_n\colon\ \mathrm{Hom}(K_2, G)^n \to \mathrm{Hom}(K_2, G^n).$$

Proof. Given an n-tuple $m = (m_1, \dots, m_n)$ of multihomomorphisms $m_i \colon K_2 \to G$, we can view m as a multihomomorphism $\iota_n(m) \colon K_2 \to G^n$ by setting $\iota_n(m)(u) = m_1(u) \times \cdots \times m_n(u)$ for each vertex u of K_2. This yields a map $\iota_n \colon \mathrm{mhom}(K_2, G)^n \to \mathrm{mhom}(K_2, G^n)$ that is monotone and equivariant and hence extends to the desired simplicial map.[1] □

The assignment $f \mapsto \mu(f)$ defines a map $\mu \colon \mathrm{pol}(C_\ell, K_4) \to \mathrm{spol}(\Gamma_{4\ell}, \Sigma^2)$ that preserves arity. The map μ does *not* strictly speaking preserve minors, i.e., for a general function $\pi \colon [n] \to [m]$, 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 $|\mu(f)| \colon T^n \cong |\Gamma^n_{4\ell}| \to |\Sigma^2| \cong S^2$, then $[\mu(f)^\pi] = [\mu(f^\pi)]$ (see Lemma A.2.3).

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 $X \to S^2$ for spaces of dimension $\dim X \geq 4$ are algorithmically undecidable [ČKM⁺13]). 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 $g \colon \Gamma^n_{4\ell} \to \Sigma^2$, we define $\eta(g) \in [T^n, Y]_{\mathbb{Z}_2}$ as the equivariant homotopy class of the composition of the geometric realization $|g| \colon T^n \to S^2$ with the inclusion map $j \colon S^2 \hookrightarrow Y$. It is easy to show that η preserves minors, and hence defines a minion homomorphism from $\mathrm{spol}(\Gamma_{4\ell}, \Sigma^2)$ to the minion $\mathrm{hpol}(S^1, Y)$ of equivariant homotopy classes of equivariant maps, i.e., the minion with $\mathrm{hpol}^{(n)}(S^1, P) = [T^n, P]_{\mathbb{Z}_2}$, 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 3.1.2 (Chapter A). *There are minion homomorphisms* $\phi \colon \mathrm{pol}(C_\ell, K_4) \to \mathrm{hpol}(S^1, Y)$ *and* $\eta \colon \mathrm{spol}(\Gamma_{4\ell}, \Sigma^2) \to \mathrm{hpol}(S^1, Y)$ *such that* $\mathrm{im}\,\phi \subseteq \mathrm{im}\,\eta$*, i.e., for each polymorphism* $f \colon C^n_\ell \to K_4$*, there is a simplicial map* $g \colon \Gamma^n_{4\ell} \to \Sigma^2$ *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 $f_\alpha \colon \mathbb{Z}_2{}^n \to \mathbb{Z}_2$ 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 $\mathscr{Z}_2^{(n)}$, and together, these sets form a function minion $\mathscr{Z}_2$. 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 $\mathscr{Z}_2$ as an abstract minion, with $\mathscr{Z}_2^{(n)} = \{\alpha \in \mathbb{Z}_2{}^n \colon \sum_i \alpha_i \equiv 1 \bmod 2\}$.

[1] It is easy to see that ι_n is injective, though generally not surjective, and it is known [Koz08, Proposition 18.17] that ι_n defines an equivariant homotopy equivalence between the spaces $|\mathrm{Hom}(K_2, G)|^n$ and $|\mathrm{Hom}(K_2, G^n)|$, but we will not need this fact in what follows.

Proposition 3.1.1. *For each $n > 0$, there is a bijection*

$$\gamma_n\colon [T^n, Y]_{\mathbb{Z}_2} \to \mathscr{Z}_2^{(n)}.$$

Moreover, these bijections preserve minors, hence they form a minion isomorphism $\gamma\colon \operatorname{hpol}(S^1, P) \to \mathscr{Z}_2$.

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

Lemma 3.1.3. *The set $[T^n, Y]_{\mathbb{Z}_2}$ has cardinality 2^{n-1}.*

One the other hand, every $\alpha \in \mathscr{Z}_2^{(n)}$ corresponds to a square-free monomial $\prod_{i\in I} z_i$ of odd degree in the variables $z_1, \dots, z_n$, where $I = \{i \in [n] \colon \alpha_i = 1\}$. If we view $S^1 = \{z \in \mathbb{C} \colon |z| = 1\}$ as the unit circle in the complex plane then each such monomial gives rise to an equivariant map $T^n = (S^1)^n \to S^1$ 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 $\Sigma^1 \subset \Sigma^2$) and then with the inclusion $j\colon S^2 \hookrightarrow Y$, we can also view each such monomial $\prod_{i\in I} z_i$ as an equivariant map $m_\alpha\colon T^n \to Y$. Using a geometrically defined set of $\mathbb{Z}_2$-valued invariants $\deg_i$, $1 \le i \le n$, we will show in Section 4.2 that these maps are pairwise non-homotopic; in fact, we will see that the map $\gamma_n\colon [T^n, Y]_{\mathbb{Z}_2} \to \mathscr{Z}_2^{(n)}$ 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 3.1.3; therefore, every equivariant map $T^n \to Y$ is equivariantly homotopic to a unique monomial map m_α with $\alpha \in \mathscr{Z}_2^{(n)}$. Moreover, we will show that the maps γ_n reserve minors, hence they form a minion isomorphism.

Step 4 Finally, we show (Theorem 5.1.1) that for every equivariant simplicial map $f\colon \Gamma_{4\ell}^n \to \Sigma^2$, 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 $\operatorname{pol}(\Gamma_{4\ell}, \Sigma^2)$ under η, and hence the image of $\operatorname{pol}(C_\ell, K_4)$ under ϕ, has bounded essential arity. This concludes the proof of Theorem 3.1.1.

3.2 Outline - Hardness of PCSP($\mathbf{LO}_3$, $\mathbf{LO}_4$)

The structure of the proof of theorem 2.1.2 is rather similar to the argument just outlined in the previous section. Here, we will present the main point where it diverges.

Here, we will use the second hardness criterion (theorem 2.2.2) provided by the algebraic theory of polymorphisms in [BBKO21]. Therefore, our overall goal is to provide a minion homomorphism from the polymorphism minion $\operatorname{pol}(\mathbf{LO}_3, \mathbf{LO}_4)$ to the minion of projections $\mathscr{P}$ (which is incidentally isomorphic to the polymorphism minion of 3-SAT).

Similarly as in the graph case, the first step replaces hypergraphs with topological spaces that are now endowed with a natural $\mathbb{Z}_3$ action using homomorphism complexes and translates the problem of understanding their polymorphisms into understanding (equivariant) homotopy classes of continuous maps $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)^n \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$. As was the case before, classifying continuous maps up to (equivariant) homotopy remains an daunting task, therefore in a second step, we further simplify the spaces involved by substituting $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)$ with a circle S^1 and $\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ with a "topologically simpler" space P.

By using the same ideas from obstruction theory, we can describe the set $[T^n, P]_{\mathbb{Z}_3}$ explicitly in terms of the corresponding equivariant cohomology group; we are thus left with a minion homomorphism χ from pol($\mathbf{LO}_3$, $\mathbf{LO}_4$) to a minion $\mathscr{Z}_3$ (defined precisely below).

The final step is significantly different from the graph case: By directly studying the polymorphisms of the pair ($\mathbf{LO}_3$, $\mathbf{LO}_4$) 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:

Step 1: As it was the case with graphs, we can use the Hom-complex construction to translate polymorphisms $f : \mathbf{LO}_3^n \to \mathbf{LO}_4$ into simplicial maps $f_* : \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3{}^n) \to \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)$ that are $\mathbb{Z}_3$-equivariant (since $\mathbf{R}_3$ has an intrinsic order 3 symmetry).

Following the same approach as before, we can define a simplicial map $\mu(f) := t \circ f_* \circ s$ where $t : \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO})_4 \to \Theta^2$ the simplicial equivariant map given in Lemma 2.3.3

Once again, the assignment $\mu : \mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathrm{spol}(\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3), \Theta^2)$ respects minors only up to equivariant homotopy; that is, if $[\mu(f)] \in [\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)^n, \mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)]_{\mathbb{Z}_3}$ denotes the equivariant homotopy class of $|\mu(f)|$, we have that $[\mu(f^\pi)] = [\mu(f)^\pi]$ for any general function $\pi : [n] \to [m]$.

Step 2: To avoid the complexity of trying to determine all possible maps up to equivariant homotopy between $|\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)|$ and $|\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_4)|$, we simplify considerably by replacing the spaces $|\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)|$ and $|L_4|$ with spaces X and P that allow equivariant maps to and from, resp., these spaces, and are better behaved from the topological perspective.

We choose X so that $X \to |\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)|$, and its powers are topologically simple but non-trivial. A natural choice is the copy of S^1 which can be seen as an equivariant subspace of $|\mathrm{Hom}(\mathbf{R}_3, \mathbf{LO}_3)|$ (Lemma 2.3.2). 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 $\mathbb{Z}_3$.

Additionally, we enlarge $|\Theta^2|$ to a "homotopically simpler" space P (once again, an Eilenberg-MacLane space) in order to simplify the computation of $[T^n, P]_{\mathbb{Z}_3}$ and obtain a minion homomorphisms ϕ : pol($\mathbf{LO}_3, \mathbf{LO}_4$) $\to$ hpol(S^1, P), η : spol(Hom($\mathbf{R}_3, \mathbf{LO}_3$), Θ^2) $\to$ hpol(S^1, P) with im $\phi \subseteq$ im η as in Lemma 3.1.2.

Step 3: Next, we give a very similar description of the sets $[T^n, P]_{\mathbb{Z}_3}$ in terms of functions $f_\alpha : \mathbb{Z}_3^n \to \mathbb{Z}_3$ of the form $f_\alpha(x_1, \dots, x_n) = \sum_i \alpha_i x_i$, where $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_3$ and $\sum_i \alpha_i = 1 \mod 3$.

The set $\mathscr{Z}_3^{(n)}$ and the function minion $\mathscr{Z}_3$ are defined analogously as the previous case.

Proposition 3.2.1. *For each $n > 0$, there is a bijection*

$$\gamma_n \colon [T^n, P]_{\mathbb{Z}_3} \to \mathscr{Z}_3^{(n)}.$$

Moreover, these bijections preserve minors, hence they form a minion isomorphism $\gamma\colon$ hpol(S^1, P) $\to \mathscr{Z}_2$.

The proof of this proposition follows the same blueprint as the corresponding Proposition 3.1.1. We first show, via equivariant obstruction theory, that:

Lemma 3.2.1. *The set $[T^n, P]_{\mathbb{Z}_3}$ has cardinality 3^{n-1}.*

Then we assign to each $\alpha \in \mathscr{Z}_3^{(n)}$ a continuous map $m_\alpha \colon T^n \to P$. We then define the analogous notion of $\deg_i$, now a set of $\mathbb{Z}_3$-valued invariants, so that the corresponding maps $\gamma_n : [T^n, P]_{\mathbb{Z}_3} \to \mathscr{Z}_3^{(n)}$ preserve minors and are such that $\gamma_n([m_\alpha]) = \alpha$, providing the desired isomorphism.

Step 4: Finally, we show that the image of the "complete" minion homomorphism χ : pol($\mathbf{LO}_3, \mathbf{LO}_4$) $\to \mathscr{Z}_3$ 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 [Wro20], and we use these ideas to connect reconfigurability with our minion homomorphism ξ.

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 $\chi(f) = \chi(g)$. Overall this means the image of $\chi_2\colon$ hpol$^{(2)}$(S^1, P) $\to \mathscr{Z}_3^{(2)}$ 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 3.2.2. *For each binary polymorphism* $f \in \text{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$, $\chi(f) \in \mathscr{P}_3^{(2)}$.

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

Corollary 3.2.1. *χ is a minion homomorphism* $\text{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathscr{P}_3$.

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 $f(x_1, \ldots, x_n) = \alpha_1 x_1 + \cdots + \alpha_n x_n$. Observe that $\alpha_1, \ldots, \alpha_n \in \{0, 1, 2\}$, 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 χ 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 2.2.2 provides the desired result, the NP-completeness of PCSP($\mathbf{LO}_3, \mathbf{LO}_4$) (Theorem 2.1.2).

CHAPTER 4

Equivariant Obstruction theory

4.1 Equivariant topology

In this section, we describe how to construct, starting from S^2 (or $|\Theta^2|$), a $\mathbb{Z}_2$-space Y ($\mathbb{Z}_3$-space P) that is homotopically simpler, together with an equivariant map $S^2 \to Y$ ($\Theta^2 \to P$).

The spaces Y and P will have the property that all of their homotopy groups $\pi_n(Y)$ and $\pi_n(P)$ for $n > 2$ are trivial (which is not the case for neither S^2 nor $|\Theta^2|$).

Moreover, the lower-dimensional homotopy groups $\pi_i(Y)$ ($\pi_i(P)$) for $i \geq 2$ are isomorphic to those of S^2 ($|\Theta^2|$); thus, both spaces are *Eilenberg–MacLane spaces*, i.e., they have only one non-trivial homotopy group, namely $\pi_2(Y) = \pi_2(S^2) \cong \mathbb{Z}$ and $\pi_2(P) = \pi_2(|\Theta^2|) \cong \mathbb{Z} \oplus \mathbb{Z}$.

The homotopy classes of maps from a complex X to an Eilenberg–MacLane space are in bijection with the elements of a suitable cohomology group of X [Hat02, Theorem 4.57]. An analogous statement is also true in the equivariant setting; this will allow us to determine both $[T^n, Y]_{\mathbb{Z}_2}$ and $[T^n, P]_{\mathbb{Z}_3}$by computing a suitable equivariant cohomology group, specifically the *Bredon cohomology group* $H^2_{\mathbb{Z}_2}(T^n; \pi_2(S^2))$ and $H^2_{\mathbb{Z}_2}(T^n; \pi_2(\Theta^2))$ respectively(see Definition 4.1.1), which will allow us to prove Lemmas 3.1.3 and 3.2.1.

Throughout this Section, we assume some familiarity with fundamental notions of algebraic topology such as homotopy, homology and cohomology. We refer to [Hat02] for background on the more basic non-equivariant setting, and to May et al. [MCC+96, Chapters I and II], [tom87], and [Bre67] for more details on equivariant homotopy and cohomology of spaces with group actions (all the definitions and constructions we use are special cases of the general theory described in these standard references).

If X is a topological space with a $\mathbb{Z}_k$-action (given by a continuous automorphism of the correct order $\nu\colon X \to X$), we will simply refer to X as a $\mathbb{Z}_k$*-space*, and we use the multiplicative notation $\nu \cdot x$ instead $\nu(x)$.

4.1.1 Construction of the equivariant Eilemberg-MacLane spaces

Construciton of Y: There are several different but ultimately equivalent ways of constructing the space Y; here (following [Hat02, Example 4.13]), we will use a simple inductive construction that starts with the sphere S^2 and achieves triviality of the higher homotopy groups $\pi_i(Y)$, $i > 2$, by successively glueing the boundaries of higher and higher-dimensional disks along non-trivial elements of the corresponding homotopy group. The formal description of this construction uses the notion of *CW complexes*.

A *CW complex* is a space X together with a increasing sequence of subspaces (called a *filtration*)

$$X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots \subseteq X,$$

with the following properties: X_0 is a discrete set of points (called *vertices* or *0-dimensional cells*) and X_{i+1} is constructed by attaching a set of $(i+1)$-dimensional discs D_α^{i+1} to X_i along their boundary via continuous maps $g_\alpha \colon \partial D_\alpha^{i+1} = S_\alpha^i \to X_i$. Thus

$$X_{i+1} = {}^{(X_i \sqcup \coprod_\alpha D_\alpha^{i+1})}\!/\!_\sim$$

where $\sim$ identifies $g_\alpha(x) \in X_i$ with $x \in \partial D_\alpha^{i+1}$. Finally, the topology on $X = \bigcup_n X_n$ is the so-called *weak topology* (i.e., a set $U \subseteq X$ is open if and only if $X \cap X_i$ is open in X_i for every i). The subspace X_i is called the i-dimensional skeleton of X.

We say that X is $\mathbb{Z}_2$-*CW complex* if, for each $i \geq 0$, $\mathbb{Z}_2$ acts on the set of i-simplices and the attaching maps respect the action. As remarked above, the geometric realization $|X|$ of simplicial set X is a CW complexes, and if X has a simplicial $\mathbb{Z}_2$-action, then $|X|$ is a $\mathbb{Z}_2$-CW complex.

Lemma 4.1.1. *There exists a $\mathbb{Z}_2$-CW complex Y such that*

1. $\pi_2(Y) = \pi_2(S^2)$;
2. $\pi_i(Y) = 0$ *for all* $i \neq 2$; *and*
3. *there is a $\mathbb{Z}_2$-map $j \colon S^2 \to Y$ that induces an isomorphism $\pi_2(j) \colon \pi_2(S^2) \to \pi_2(Y)$ of groups with a $\mathbb{Z}_2$-action.*

Proof sketch. The sphere S^2 can be viewed as a $\mathbb{Z}_2$-CW complex with the antipodal action (we can, e.g., take the geometric realization of the simplicial set Σ^2). Starting with this $\mathbb{Z}_2$-CW complex, we construct the i-skeleton Y_i of Y as follows:

1. We set $Y_2 := S^2$.
2. For $i > 2$ we create a space Y_i as follows: Start with Y_{i-1} and for every generator α of $\pi_i(Y_{i-1})$ we attach two $(i+1)$-dimensional discs $D_\alpha, D_{\nu\cdot\alpha}$ by identifying ∂D_α with α and $\partial D_{\nu\cdot\alpha}$ with $\nu \cdot \alpha$. We then extend the $\mathbb{Z}_2$ action in the natural way by "swapping" the paired discs D_α and $D_{\nu\cdot\alpha}$.

3. Finally, we take $Y = \bigcup_{i\geq 2} Y_i$.

It is not hard to check (see [Hat02, Example 4.13]) that the $\mathbb{Z}_2$-CW complex Y satisfies $\pi_i(Y) = 0$ for $i \geq 3$ and $\pi_i(Y) = \pi_i(S^2)$ for $i \leq 2$. Further, $j\colon S^2 \to Y$ is defined as the inclusion of the 2-skeleton $Y_2 = S^2$ into Y. □

Construciton of P: The process to build P is a nearly identical inductive construction starting from $|\Theta^2|$ and iteratively gluing discs along non trivial elements of the corresponding homotopy groups.

In particular, by mimicking the proof of Lemma 4.1.1 with the exception that on step 2. we glue three discs instead of two, we obtain the following result.

Lemma 4.1.2. *There exists a $\mathbb{Z}_3$-CW complex P such that*

1. $\pi_2(P) = \pi_2(\Theta^2)$;
2. $\pi_i(P) = 0$ *for all* $i \neq 2$; *and*
3. *there is a $\mathbb{Z}_3$-map $j\colon \Theta^2 \to P$ that induces an isomorphism $\pi_2(j)\colon \pi_2(\Theta^2) \to \pi_2(P)$ of groups with a $\mathbb{Z}_3$-action.*

4.1.2 Equivariant cohomology: A primer

We now introduce the Bredon cohomology that will help us classify equivariant maps. Since the theory and the necessary computations are basically identical for both cases we are interested in (i.e. $\mathbb{Z}_2$ and $\mathbb{Z}_3$ free actions), in this section we will work with spaces with a free $\mathbb{Z}_k$-action (with generator ν) and specify $k = 2$ or $k = 3$ only when needed.

Prescribing a $\mathbb{Z}_k$-action on an Abelian group M is the same as giving M the structure of a module over the *group ring* $\mathbb{Z}[\mathbb{Z}_k]$ (which is isomorphic to the the quotient $\mathbb{Z}[\nu]/(\nu^k - 1)$ of the polynomial ring by the ideal $(\nu^k - 1)$). In particular, if Y is a space with a $\mathbb{Z}_k$-action, then this action naturally induces a $\mathbb{Z}_k$-action on every homotopy group $\pi_i(Y)$ and hence turns $\pi_i(Y)$ into a $\mathbb{Z}[\mathbb{Z}_k]$-module. In what follows, we will mainly use the terminology of $\mathbb{Z}[\mathbb{Z}_k]$-modules (rather than speaking of abelian groups with $\mathbb{Z}_k$-actions). We are now ready to recall the definition of equivariant homology and cohomology groups.

Definition 4.1.1 (Equivariant homology and cohomology). Let X be a $\mathbb{Z}_k$-CW complex. Its d-dimensional chain group $C_d(X)$ has a natural structure of $\mathbb{Z}[\mathbb{Z}_k]$-module with multiplication given on a cell σ by

$$(\sum_{i=0}^{k-1} n_i \nu^i)\sigma = \sum_{i=0}^{k-1} n_i(\nu^i \cdot \sigma)$$

and extended linearly. Since the all the boundary maps commute with the action, these are $\mathbb{Z}[\mathbb{Z}_k]$-module homomorphisms, and hence $C_\bullet(X)$ can be viewed as a chain complex of $\mathbb{Z}[\mathbb{Z}_k]$-modules. We denote this chain complex by $C_\bullet^{\mathbb{Z}_k}(X)$. The homology associated to this chain complex is the *equivariant homology* of X, denoted by $H_\bullet^{\mathbb{Z}_k}(X)$.

Fix a $\mathbb{Z}[\mathbb{Z}_k]$-module N, and consider the equivariant cochain complex:

$$C^i_{\mathbb{Z}_k}(X;N) = \mathrm{Hom}_{\mathbb{Z}[\mathbb{Z}_k]}(C_i^{\mathbb{Z}_k}(X), N)$$

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

The key result we will need is the following equivariant version of the classical Brown representability theorem:

Theorem 4.1.1 ([tom87, Theorem II.3.17]; see also [MCC+96, Chapter II]). *Let Z be a $\mathbb{Z}_k$-CW complex which is an Eilenberg–MacLane space whose unique non-trivial homotopy group is $\pi_i(Z)$ (we assume that $\pi_1(Z)$ is abelian if $i = 1$). Then, for every $\mathbb{Z}_k$-CW complex X such that there is a $\mathbb{Z}_k$-map $X \to Z$, the set $[X, Z]_{\mathbb{Z}_k}$ of $\mathbb{Z}_k$-equivariant homotopy classes of equivariant maps is in bijection with $H^i_{\mathbb{Z}_k}(X;\pi_i(Z))$.*

We will apply Theorem 4.1.1 in the case where $X = T^n$ (with the diagonal action) and, for $k = 2$, $Z = Y$ is the Eilenberg–MacLane space constructed in Lemma 4.1.1 while $Z = P$ when $k = 3$ (Lemma 4.1.2). The last remaining ingredient for the proof of Lemmas 3.1.3 and 3.2.1 is the following result on the Bredon cohomology of the torus T^n, which we will prove in Section 4.1.3 below:

Proposition 4.1.1. *Let $\mathfrak{M}$ be a $\mathbb{Z}[\mathbb{Z}_k]$-module generated by a single element such that* $\mathrm{Ann}(M) = \mathfrak{I}$ *the ideal generated by the element* $(1 + \nu + \ldots \nu^{k-1}) \in \mathbb{Z}[\mathbb{Z}_k]$. *For all $n, d \geq 1$,*

$$H^d_{\mathbb{Z}_k}(T^n;\mathfrak{M}) \cong \mathbb{Z}_k^{\binom{n-1}{d-1}}.$$

This immediately implies Lemmas 3.1.3 and 3.2.1 given the following observation:

Lemma 4.1.3. *Both $\pi_2(Y)$ and $\pi_2(P)$ are generated by a single element as $\mathbb{Z}[\mathbb{Z}_2]$ and $\mathbb{Z}[\mathbb{Z}_3]$-modules with annihilators* $(1 + \nu)$ *and* $(1 + \nu + \nu^2)$ *respectively.*

Proof. The second homotopy group $\pi_2(S^2)$ is isomorphic as a $\mathbb{Z}[\mathbb{Z}_2]$-module to $\mathfrak{Z}_- = \mathbb{Z}$ where the multiplication by ν is $\nu \cdot n = -n$; in particular it is generated by 1 and its annihilator is the ideal generated by $1 + \nu$ in $\mathbb{Z}[\mathbb{Z}_2]$.

On the other hand, since Θ^2 is simply connected, by Hurewicz theorem [Hat02, Section 4.2], there is an isomorphism between $\pi_2(\Theta^2)$ and $H_2(\Theta^2)$ that is an isomorphism of $\mathbb{Z}[\mathbb{Z}_3]$-modules.

By direct computation, it is easy to show that the chains $\beta_i = d_i - d_{i+1} \in C_2(\Theta^2)$ (where $d_0 = [\bullet, \bullet, \bullet]$ and $\nu \cdot d_i = d_{i+1}$, see Figure 2.2) are cycles and any two out of $\{[\beta_i]\}_{0\leq i\leq 2}$ form a basis for $H_2(\Theta^2) \cong \mathbb{Z}\oplus\mathbb{Z}$. Fix $[\beta_0]$ and $[\beta_1]$, then the multiplication by ν is

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

Therefore, $\pi_2(P)$ is generated over $\mathbb{Z}[\mathbb{Z}_3]$ by $m = \binom{1}{0}$ and $\lambda \in \mathrm{Ann}(\pi_2(P))$ if and only if $\lambda m = 0$; hence, if $\lambda = n_0 + n_1\nu + n_2\nu^2$, then

$$\lambda m = n_0\binom{1}{0} + n_1\binom{0}{1} + n_2\binom{-1}{-1} = \binom{n_0-n_2}{n_1-n_2}.$$

Therefore, $\lambda m = 0$ if and only if $n_0 = n_1 = n_2$ if and only if $\lambda \in (1+\nu+\nu^2)$. □

Proof of Lemmas 3.1.3 and 3.2.1. In the case $k = 2$, by combining Theorem 4.1.1 and Proposition 4.1.1 and specializing to $d = 2$, we get the following bijection:

$$[T^n, Y]_{\mathbb{Z}_2} \cong H^2_{\mathbb{Z}_2}(T^n;\pi_2(S^2)) \cong \mathbb{Z}_2^{n-1}$$

Thus, $[T^n, Y]_{\mathbb{Z}_2}$ has 2^{n-1} elements, as we wanted to show.

Similarly, when $k = 3$, we get the bijection:

$$[T^n, P]_{\mathbb{Z}_3} \cong H^2_{\mathbb{Z}_3}(T^n;\pi_2(P)) \cong \mathbb{Z}_3^{n-1}$$

Thus, $[T^n, Y]_{\mathbb{Z}_2}$ contains 3^{n-1} elements. □

4.1.3 The equivariant cohomology of the torus

This section is devoted to proving Proposition 4.1.1. We begin with two technical lemmas that are useful for computing the equivariant cohomology of spaces with a free action.

Lemma 4.1.4. *If the action on X is free and cellular, then $C^{\mathbb{Z}_k}_\bullet(X)$ is a chain complex of free $\mathbb{Z}[\mathbb{Z}_k]$-modules.*

Proof. For every orbit of d-cells in X choose a representative σ, and observe that the module $C_d(X)$ is freely generated by the set of these representatives. □

The above lemma implies that the functor $\mathrm{Hom}_{\mathbb{Z}[\mathbb{Z}_k]}(C^{\mathbb{Z}_k}_d(X), -)$ is exact for all free $\mathbb{Z}_k$-CW complexes X and $d \geq 0$. Therefore, if we have a short exact sequence of $\mathbb{Z}[\mathbb{Z}_k]$-modules

$$0 \longrightarrow \mathfrak{N} \xrightarrow{g} \mathfrak{M} \xrightarrow{f} \mathfrak{Q} \longrightarrow 0$$

there is a corresponding short exact sequence of cochain complexes

$$0 \longrightarrow C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{N}) \xrightarrow{g_*} C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{M}) \xrightarrow{f_*} C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{Q}) \longrightarrow 0$$

and thus a long exact sequence in cohomology

$$\cdots \longrightarrow H^i_{\mathbb{Z}_k}(X;\mathfrak{N}) \longrightarrow H^i_{\mathbb{Z}_k}(X;\mathfrak{M}) \longrightarrow H^i_{\mathbb{Z}_k}(X;\mathfrak{Q}) \longrightarrow H^{i+1}_{\mathbb{Z}_k}(X;\mathfrak{N}) \longrightarrow H^{i+1}_{\mathbb{Z}_k}(X;\mathfrak{M}) \longrightarrow \cdots$$

Note that $\mathbb{Z}$ always admits a structure of $\mathbb{Z}[\mathbb{Z}_k]$-module: $\nu \cdot 1 = 1$; in which case the action is trivial and we denote such module as $\mathfrak{Z}_+$.

Define $\mathfrak{I} = (1+\nu+\cdots+\nu^{k-1})\mathbb{Z}[\mathbb{Z}_k]$ the ideal generated by $1+\nu+\cdots+\nu^{k-1}$ and $\mathfrak{C}$ the module $\mathbb{Z}^k$ where the action is $\nu(n_0,\ldots,n_{k-1}) = (n_{k-1}, n_0, \ldots, n_{k-2})$, i.e., it shifts the coordinates cyclically. Then, we have the following diagram of $\mathbb{Z}[\mathbb{Z}_k]$-module:

$$\begin{array}{ccc} \mathfrak{Z}_+ & \xrightarrow{\phi_1} & \mathfrak{I} \\ {\scriptstyle d}\downarrow & & \downarrow{\scriptstyle \iota} \\ \mathfrak{C} & \xrightarrow{\phi_2} & \mathbb{Z}[\mathbb{Z}_k] \end{array}$$

The maps in this diagram are defined as: $d(n) = (n,\ldots,n)$, ι is the inclusion, $\phi_1(n) = n(1+\nu+\cdots+\nu^{k-1})$, and $\phi_2(n_0,\ldots,n_{k-1}) = n_0 + n_1\nu + \cdots + n_{k-1}\nu^{k-1}$. Moreover, note that ϕ_1 and ϕ_2 are isomorphisms of $\mathbb{Z}[\mathbb{Z}_k]$-modules, hence the corresponding induced cochain maps

$$(\phi_1)_*\colon C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{Z}_+) \to C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{I}) \qquad (\phi_1)_*\colon \alpha \mapsto \phi_1 \circ \alpha$$
$$(\phi_2)_*\colon C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{C}) \to C^\bullet_{\mathbb{Z}_k}(X;\mathbb{Z}[\mathbb{Z}_k]) \qquad (\phi_2)_*\colon \alpha \mapsto \phi_2 \circ \alpha$$

are isomorphisms of cochain complexes.

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

$$h_1\colon C^\bullet(X/\mathbb{Z}_k;\mathbb{Z}) \to C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{Z}_+) \qquad h_1(\alpha)\colon \sigma \mapsto \alpha(p(\sigma))$$
$$h_2\colon C^\bullet(X;\mathbb{Z}) \to C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{C}) \qquad h_2(\alpha)\colon \sigma \mapsto \big(\alpha(\nu^i\cdot\sigma)\big)_{0\le i<k}.$$

Lemma 4.1.5. *Let X be a CW complex with a free $\mathbb{Z}_k$-action, then the following diagram commutes*

$$\begin{array}{ccccc} C^\bullet(X/\mathbb{Z}_k;\mathbb{Z}) & \xrightarrow{h_1} & C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{Z}_+) & \xrightarrow{(\phi_1)_*} & C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{I}) \\ \downarrow{\scriptstyle p^*} & & \downarrow{\scriptstyle d_*} & & \downarrow{\scriptstyle \iota_*} \\ C^\bullet(X;\mathbb{Z}) & \xrightarrow{h_2} & C^\bullet_{\mathbb{Z}_k}(X;\mathfrak{C}) & \xrightarrow{(\phi_2)_*} & C^\bullet_{\mathbb{Z}_k}(X;\mathbb{Z}[\mathbb{Z}_k]) \end{array}$$

Moreover, as already noted, all the horizontal homomorphisms (i.e. h_i and $(\phi_i)_*$) are isomorphisms of cochain complexes.

Proof. The right square commutes by functoriality of $C^\bullet_{\mathbb{Z}_k}(X;-)$, therefore it is enough to show that the left square commutes, i.e., that, for any $d \ge 0$ and any $\alpha \in C^d(X/\mathbb{Z}_k;\mathbb{Z})$, $h_2(p^*(\alpha)) = d_*(h_1(\alpha))$:

We have, for all $\sigma \in C_d^{\mathbb{Z}_k}(X)$,

$$h_2(p^*(\alpha))(\sigma) = (p^*(\alpha)(\nu^i \cdot \sigma))_{0 \le i < k} = \big(\alpha(p(\nu^i \cdot \sigma))\big)_{0 \le i < k} = \big(\alpha(p(\sigma))\big)$$
$$= (h_1(\alpha)(\sigma), h_1(\alpha)(\sigma))_{0 \le i < k} = d_*(h_1(\alpha))(\sigma)$$

as claimed. □

The last ingredient we will need is to determine what the projection map p^* does on the level of cohomology. While it is possible to compute p^* directly, it is easier to view the action on the torus from a different perspective to simplify the calculations:

Lemma 4.1.6. *Let X be the torus $T^n \subseteq \mathbb{C}^n$ with the diagonal $\mathbb{Z}_k$-action given by the multiplication with a primitive k-root of unity ω_k (i.e., $\nu \cdot (z_1, \ldots, z_n) = (\omega_k z_1, \ldots, \omega_k z_n)$). Let Y be the same torus but with $\mathbb{Z}_k$ acting only on the first coordinate (i.e., $\nu \cdot (z_1, \ldots, z_n) = (\omega_k z_1, z_2, \ldots, z_n)$). Then there is a $\mathbb{Z}_k$-equivariant homeomorphism $X \longrightarrow Y$.*

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

$$h \colon (z_1, \ldots, z_n) \mapsto (z_1, z_1^{-1} z_2, \ldots, z_1^{-1} z_n)$$
$$h' \colon (z_1, \ldots, z_n) \mapsto (z_1, z_1 z_2, \ldots, z_1 z_n)$$

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

$$h(\nu \cdot_X (z_1, \ldots, z_n)) = h(\omega_k z_1, \ldots, \omega_k z_n) = (\omega_k z_1, z_1^{-1} z_2, \ldots, z_n z_1^{-1})$$
$$= \nu \cdot_Y h(z_1, \ldots, z_n) \quad \square$$

Remark 1. If we view the torus T^n as the quotient of $\mathbb{R}^n$ by the standard lattice $\mathbb{Z}^n$, then Lemma 4.1.6 shows that factoring out the action is the same as factoring out the lattice generated by $\{\frac{1}{k} e_1, e_2, \ldots, e_n\}$. Hence, topologically, the quotient is still a torus.

Thus, for the remaining (co)homological calculations, we can assume that $\mathbb{Z}_k$ acts on T^n by changing only on the first coordinate. Using this simplified action on the torus it is much easier to compute the quotient map $p^* \colon H^\bullet(T^n/\mathbb{Z}_k) \longrightarrow H^\bullet(T^n)$. To achieve this objective, it is necessary to fix a basis for the cohomology of the torus. The ideal choice would be a basis that is "easy" to evaluate on homology classes in order to compute easily the image of p^*.

In the case of the torus, a direct application of the universal coefficient theorem [Hat02, Section 3.1] show that homology and cohomology in dimension 1 are dual to each other; hence we can choose as basis for the first cohomology group the dual of a suitable basis for the first homology group. In particular, let $\{x_i\}$ be the basis for $H_1(T^n)$ corresponding to the standard coordinate cycles in $C_1(T^n)$ (i.e., x_i corresponds to the (non-equivariant) inclusion $S^1 \hookrightarrow T^n$, $z \mapsto (0, \ldots, 0, z, 0, \ldots, 0)$ in the ith coordinate, $1 \le i \le n$), and denote by

$\{x^i\}$ the dual basis in $H^1(T^n) \cong \mathrm{Hom}\,(H_1(T^n), \mathbb{Z})$; analogously, define bases $\{q_i\}$ of $H_1(T^n/\mathbb{Z}_k)$ and $\{q^i\}$ of $H^1(T^n/\mathbb{Z}_k)$. Then

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

The ring structure on cohomology (see [Hat02, 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 $I = (i_1, \ldots, i_d)$ and $i_1 < \cdots < i_d$, form a basis for $H^d(T^n)$. Let q^I denote the analogous basis for $H^d(T^n/\mathbb{Z}_2)$. 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,

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

In particular, p_d^* is injective for all $d \geq 0$ and, in this choice of basis, p^* is the diagonal matrix with $\binom{n-1}{d-1}$ k'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 4.1.1.

Proof of Proposition 4.1.1. Fix $n \geq 2$. By hypothesis, we have a short exact sequence

$$0 \longrightarrow \mathfrak{I} \longrightarrow \mathbb{Z}[\mathbb{Z}_k] \longrightarrow \mathfrak{M} \longrightarrow 0$$

which induces short exact sequence of cochain complexes

$$0 \longrightarrow C^\bullet_{\mathbb{Z}_2}(T^n; \mathfrak{I}) \longrightarrow C^\bullet_{\mathbb{Z}_2}(T^n; \mathbb{Z}[\mathbb{Z}_k]) \longrightarrow C^\bullet_{\mathbb{Z}_2}(T^n; \mathfrak{M}) \longrightarrow 0$$

Using Lemma 4.1.5, we get that the following short sequence is also exact

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

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

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

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

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

which yields the desired result. □

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

The goal of this section is to prove Proposition 3.1.1.

To this end, we will define, for every equivariant continuous map $f: T^n \to Y$, 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 $\sum_{i=1}^n \deg_i(f) \equiv 1 \bmod 2$. Thus, by assigning to every equivariant homotopy class $[f] \in [T^n, Y]_{\mathbb{Z}_2}$ the sequence $\gamma_n([f]) = (\deg_1(f), \ldots, \deg_n(f)) \in \mathbb{Z}_2{}^n$, we get a well-defined map $\gamma_n : [T^n, Y]_{\mathbb{Z}_2} \to \mathcal{Z}_2^{(n)}$.

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 [Hat02] for general background, and to [MCC+96, Chapters I and II],[tom87], and [Bre67] for more details on the equivariant setting.

We will use the fact that the space Y is is a CW complex constructed from Σ^2 by attaching higher-dimensional cells (see the proof of Lemma 4.1.1 in Section 4.1.1); 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 and equivariant map $f: T^2 \to Y$. Consider a loop in T^2 that wraps around the first coordinate direction, say the circle $x_1 = \{(z_1, 1): z_1 \in S^1\} \cong S^1 \subset T^2$. 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 $(z_1, z_2) \mapsto\mapsto j(z_1)$, 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 $g(w) = \bullet$. 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 $\nu \cdot x_1$ 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 4.2.1. 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

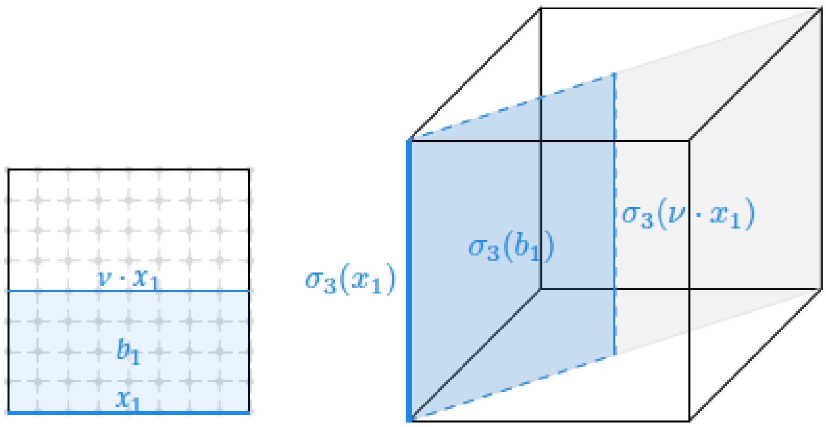

Figure 4.1: Coordinate cycle and band in T^2 and T^3.

the dual of $e_0 = [\bullet, \bullet] \in C_1(\Sigma^2) = C_1(Y)$ and $d_0 = [\bullet, \bullet, \bullet] \in C_2(\Sigma^2) = C_2(Y)$ respectively (i.e., $e^0([\bullet, \bullet]) = 0$ and $e^0([\bullet, \bullet]) = 1$, 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 $\nu \cdot x_1$ (i.e., $\partial b_1 = x_1 + \nu \cdot x_1$; see Figure 4.1). Let $f\colon |\Gamma_L \times \Gamma_{L'}| \to Y$ be an equivariant map. By the (equivariant) cellular approximation theorem, f induces an equivariant cochain map $f^*\colon C^\bullet(Y) \to C^\bullet(\Gamma_L \times \Gamma_{L'})$ (i.e., equivariant homomorphisms $f^*\colon C^i(Y) \to C^i(X)$ that commute with the coboundary map). We define

$$\deg_1(f) = \left(f^*(e^0)(x_1) + f^*(d^0)(b_1)\right) \bmod 2$$

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

Lemma 4.2.1. *Fix positive integers L, L_0 and L_1 divisible by 4. Let $f_0\colon |\Gamma_L \times \Gamma_{L_0}| \to Y$ and $f_1\colon |\Gamma_L \times \Gamma_{L_1}| \to Y$ 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 $f_0^*, f_1^*\colon C^\bullet(Y) \to C^\bullet(X)$, i.e., there exist equivariant homomorphisms $h\colon C^i(Y) \to C^{i-1}(X)$ satisfying

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

Therefore, on the cochains of dimension 1 we have:

$$f_0^*(e^0)(x_1) + f_1^*(e^0)(x_1) = (\delta h(e^0))(x_1) + (h\delta(e^0))(x_1) = 0 + h(d^0 + d^1)(x_1)$$

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:

$$f_0^*(d^0)(b_1) + f_1^*(d^0)(b_1) = (\delta h(d^0))(b_1) + (h\delta(d^0))(b_1) = h(d^0)(x_1 + \nu \cdot x_1) + 0$$

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

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

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

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 $\iota\colon C_\bullet(\Gamma_L \times \Gamma_{L_0}) \to C_\bullet(\Gamma_L \times \Gamma_{L_1})$ 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, $\deg_1(f_1 \circ \iota) = \deg_1(f_1)$. □

We can now define $\deg_i(f)$ of an equivariant map $f\colon T^n \to Y$ as $\deg_1(f^{\sigma_i})$ for a suitable 2-minor $f^{\sigma_i}\colon T^2 \to Y$ (see Figure 4.1):

Definition 4.2.2. Let L a positive integer divisible by 4, and let $f\colon |\Gamma_L^n| \to Y$ be a $\mathbb{Z}_2$-equivariant map. For $i \in [n]$, we define $\sigma_i\colon [n] \to [2]$ by $\sigma_i(i) = 1$ and $\sigma_i(j) = 2$ for $j \neq i$. Then the i-degree of f is defined as

$$\deg_i(f) = \deg_1(f^{\sigma_i}) = \big((f \circ \sigma_i)^*(e^0)(x_1) + (f \circ \sigma_i)^*(d^0)(b_1)\big) \bmod 2$$

An immediate consequence of Lemma 4.2.1 is the invariance of the i-degree under equivariant homotopies:

Corollary 4.2.1. *Let $f_0, f_1\colon |\Gamma_L^n| \to Y$ 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 $f_1^{\sigma_i}$. □

It will be convenient to extend the notation for monomial maps to general integer coefficients. As before, let us view $S^1 = \{z \in \mathbb{C}\colon |z| = 1\}$ as the unit circle in the complex plane. Given an n-tuple of integers $\alpha = (\alpha_1, \dots, \alpha_n)$ with $\sum_i \alpha_i \equiv 1 \bmod 2$, we get an equivariant map from T^n to S^1 defined by $(z_1, \dots, z_n) \mapsto z_1^{\alpha_1} \cdots z_n^{\alpha_n}$.

By composing this map first with a fixed equivariant inclusion $S^1 \hookrightarrow S^2$ and then with the inclusion $j\colon S^2 \to Y$, we get an equivariant *monomial map* $m_\alpha\colon T^n \to Y$ given by

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

Remark 2. 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, $\mathscr{Z}_2$ is *not* a subminion of the minion of tuples $\alpha \in \mathbb{Z}^n$ with $\sum \alpha_i \equiv 1 \pmod 2$ 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 4.2.2. *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 $\deg_i(m_\alpha) = \deg(m_\beta) = \alpha_i + 0 \bmod 2$. □

Corollary 4.2.2. *Let $\alpha, \beta \in \mathscr{Z}_2^{(n)}$. 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 4.2.1. Thus, by Lemma 4.2.2, $\alpha_i = \beta_i$, for all $i \in [n]$. □

We are now ready to prove Proposition 3.1.1:

Proof of Proposition 3.1.1. For every $n \geq 1$, consider the map $\gamma_n : [T^n, Y]_{\mathbb{Z}_2} \to \mathbb{Z}_2{}^n$ given by

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

By Corollary 4.2.1, this mapping is well-defined. Moreover, by Lemma 4.3.3, if $\alpha \in \mathscr{Z}_2^{(n)}$, 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 $[m_\alpha]$, $\alpha \in \mathscr{Z}_2^{(n)}$, are pairwise distinct, and by Lemma 3.1.3, they account for all elements of $[T^n, Y]_{\mathbb{Z}_2}$, i.e., every equivariant map $f : T^n \to Y$ is equivariantly homotopic to m_α for a unique $\alpha \in \mathscr{Z}_2^{(n)}$. It follows that $\gamma_n([f]) \in \mathscr{Z}_2^{(n)}$ and that γ_n is a bijection.

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

$$\gamma_n([m_\alpha]) = (\alpha_1 \bmod 2, \ldots, \alpha_n \bmod 2)$$

by Lemma 4.3.3, 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 $\gamma\colon \operatorname{hpol}(S^1, Y) \to \mathscr{Z}_2$. □

Finally, we show non zero degree guarantees a colour swapping edge, result that we will use in the final combinatorial step of our proof of Theorem 2.1.1.

Lemma 4.2.3. *Let $f : \Gamma_L \times \Gamma_{L'} \to \Sigma^2$ 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 $f(v_1, v_2) = {\color{blue}\bullet}$ and $f(v_1 + 1, v_2) = {\color{orange}\bullet}$.*

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., $f([u, v, w]) = [{\color{blue}\bullet}, {\color{orange}\bullet}, {\color{blue}\bullet}]$ or $[{\color{orange}\bullet}, {\color{blue}\bullet}, {\color{orange}\bullet}]$); 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 $\deg_1(f) = f^*(e_0)(x_1) + f^*(d^0)(b_1) = 0$ □

4.3 Monomial Maps, Degrees, and $[T^n, P]_{\mathbb{Z}_3}$

The goal of this section is to show the minion isomorphisms claimed in Proposition 3.2.1.

We start with describing the corresponding $\mathbb{Z}_3$-equivariant version of the $\mathbb{Z}_2$-monomial maps defined in the previous section. In particular, each n-tuple $(\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_3^n$ with $\sum \alpha_i = 1 \mod 3$ induces a $\mathbb{Z}_3$-equivariant map $T^n \to S^1$. 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 4.3.1. Fix an embedding $\iota : S^1 \hookrightarrow P$ (e.g., the inclusion of the 1-skeleton), and let $n \geq 1$. We assigns to each tuple $\alpha \in \mathbb{Z}^n$ with $\sum_i \alpha_i = 1$ (mod 3), a *monomial map* $m_\alpha : T^n \to P$, 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

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

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

Remark 3. As it was previously observed in the case with $\mathbb{Z}_2$-monomial maps, it is not hard to observe that the assignment $\alpha \mapsto m_\alpha$ preserves minors when α is interpreted as a function $f : \mathbb{Z}^n \to \mathbb{Z}$. As it was the case previously we implicitly use this minion homomorphism it is not the one we are looking for since $\mathscr{Z}_3$ is *not* a subminion of the minion of tuples $\alpha \in \mathbb{Z}^n$ with $\sum \alpha_i = 1$ (mod 3).

The overall goal is now to build the algebraic machinery needed to show that the $\mathbb{Z}_3$-monomial maps are all non-homotopic; that is the following result.

Lemma 4.3.1. *Let $\alpha, \beta \in \mathscr{Z}_3^{(n)}$. If m_α is equivariantly homotopic to m_β, then $\alpha = \beta$.*

The proof of this lemma is build on a similar informal geometric intuition as before, but adapted for the $\mathbb{Z}_3$-equivariant case.

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 formalization of this idea follows a similar blueprint as the previous case.

Definition 4.3.2. 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 i-th coordinate cycle in T^n and $b_i, B_i \in C_2(T^n)$ fillings for $x_i - \nu \cdot x_i$ and $x_i - \nu^2 \cdot x_i$ respectively.

Then, the i-degree of an equivariant map $f : T^n \to P$ 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 4.3.2. *Let $f_0, f_1 \colon T^n \to P$ 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:

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

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

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

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

□

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 4.3.3. *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 4.3.1 which follows immediately from the above.

Proof of Lemma 4.3.1. 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 4.3.3, and therefore m_α and m_β are not equivariantly homotopic by Lemma 4.3.2. □

Before, we progress further, let us discuss one more consequence of Lemma 4.3.1, namely, the following.

Lemma 4.3.4. *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, for each $\alpha \in \mathbb{Z}^n$ *with* $\sum_i \alpha_i = 1$, *and* $\pi\colon [n] \to [m]$, $\gamma([m_\alpha]^\pi) = \gamma([m_\alpha])^\pi$, *i.e., it preserves minors when restricted to classes of monomial maps.*

Proof. Note that γ is well-defined since the degrees do note depend on the choice of representative (Lemma 4.3.2). Further, using Lemma 4.3.3, 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

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

Putting everything together, we can now provide the required isomorphism of minions $\mathscr{Z}_3$ and $\operatorname{hpol}(S^1, P)$.

Proof of Lemma 3.2.1. We show that μ and γ are mutually inverse minion homomorphisms. Recall that μ is injective by Lemma 4.3.1 and it is surjective by Lemma 3.2.1.

The above, in particular, means that $\mu_n\colon \mathscr{Z}_3^{(n)} \to \operatorname{hpol}^{(n)}(S^1, P)$ is onto, and hence a bijection. Consequently, γ is a minion homomorphism by Lemma 4.3.4 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 4.3.3, 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. □

CHAPTER 5

The combinatorics of polymorphisms

5.1 Bounding Essential Arity of Maps induced by Graph Polymorphisms

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

Theorem 5.1.1. *Let $L \geq 4$ be an integer divisible by 4, let $f\colon \Gamma_L^n \to \Sigma^2$ be an equivariant simplicial map such that the composition with the map $\Sigma^2 \to Y$ 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 1) that equivariant simplicial maps $f\colon \Gamma_L^n \to \Sigma^2$ 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 u and v are different and u and v differ only in the ith 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 $u = (u_1, \ldots, u_n) \in \mathbb{Z}_L^n$; edges (1-simplices) are pairs $[u, v]$ of vertices such that v is obtained from u by choosing a non-empty subset of coordinates of u that are all even, and changing each of them by ± 1 modulo L; and the k-simplices are $(k+1)$-tuples $[u_0, u_1, \ldots, u_k]$ such that $[u_{j-1}, u_j]$ is an edge for $1 \leq j \leq k$. We define the *height* $\mathrm{ht}(u)$ of a vertex $u = (u_1, \ldots, u_n)$ as the number of coordinates $i \in [n]$ such that u_i is odd; moreover, we define the height of an edge $[u, v]$ as the height of u. Note that *every* edge $[u, v]$, we have $\mathrm{ht}(u) < \mathrm{ht}(v)$. A special role will be played by edges $[u, v]$ such that $\mathrm{ht}(v) = \mathrm{ht}(u)+1$, or equivalently, such that u and v differ in exactly one coordinate; we call such edges *coordinate edges*. More precisely, we say that an edge $[u, v]$ is *in coordinate direction i* if u and v differ exactly in the ith coordinate. For $i \in [n]$, we denote the set of all edges in coordinate direction i by E_i, and denote by $E := E_1 \sqcup \cdots \sqcup E_n$ 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 $E(h) = E_1(h) \sqcup \cdots \sqcup E_n(h)$ 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 $[u, v]$ of Γ_L^n is *colour-swapping* if u and v have different colours. We now state a key lemma used in the proof of Theorem 5.1.1. 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 5.1.1. *Let $f\colon \Gamma_L^n \to \Sigma^2$ 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 5.1.1.

Proof of Theorem 5.1.1 assuming Lemma 5.1.1. 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 $\pi\colon [n] \to [m]$ that is injective on I. Then the minor f^π is an equivariant simplicial map $f^\pi\colon \Gamma_L^m \to \Sigma^2$ that is equivariantly homotopic to the monomial map $\prod_{i\in[m]} z_i$, by Lemma A.2.5.

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 $\sigma = [u_0, \ldots, u_n]$ 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 $[u_h, u_{h+1}]$ 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 $f(u_i)$ changes as we traverse the vertices of σ in their given order. Observe that, for every $0 \leq h \leq n-1$, the edge $[u_h, u_{h+1}]$ 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 $\mathbf{E}[X_h(\sigma)]$ is the probability that a uniformly random edge in $E(h)$ is colour-swapping. Moreover, by Lemma 5.1.1 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, $\mathbf{E}[X_h(\sigma)] + \mathbf{E}[X_{n-1-h}(\sigma)] \geq \frac{1}{CL^2}$ for $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$. Consequently,

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

Thus, there exists some n-simplex $\sigma = [u_0, \dots, u_n]$ of Γ_L^n such that the colour of $f(u_i)$ 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 $[u_{i_0}, u_{i_1}, \dots, u_{i_k}]$ 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 $|I| = n = O(L^2)$. □

The rest of this section is dedicated to proving Lemma 5.1.1.

In this proof, we will use Lemma 4.2.3 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* $\mathrm{diag}\colon \Gamma_L \hookrightarrow \Gamma_L^{n-1}$ given by $\mathrm{diag}(y) = (y, \dots, y)$. This is an equivariant simplicial map, which induces an equivariant simplicial embedding $s_{\mathrm{diag}}\colon \Gamma_L \times \Gamma_L \hookrightarrow \Gamma_L^n$ given by $s_{\mathrm{diag}} := 1_{\Gamma_L} \times \mathrm{diag}$, i.e., $s_{\mathrm{diag}}(x, y) = (x, y, \dots, y)$.

More generally, let L' be an integer divisible by 4, and let $\zeta\colon \Gamma_{L'} \to \Gamma_L^{n-1}$ be an equivariant simplicial map; we call ζ a *generalized diagonal* if its geometric realization $|\zeta|$, seen as an equivariant embedding $S^1 \to T^{n-1}$, is equivariantly homotopic to the diagonal embedding $S^1 \to T^{n-1}$ (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 $s_\zeta\colon \Gamma_L \times \Gamma_{L'} \to \Gamma_L^n$ 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 5.1.2. *Let $f\colon \Gamma_L^n \to \Sigma^2$ be an equivariant simplicial map such that $\deg_1(f) = 1$, and let $s_\zeta\colon \Gamma_L \times \Gamma_{L'} \to \Gamma_L^n$ 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 $f \circ s_{\mathrm{diag}}$ is the same as the 2-minor f^π of f given by the map $\pi\colon [n] \to [2]$, $\pi(1) = 1$ and $\pi(j) = 2$ for $2 \leq j \leq n$. Thus, $\deg_1(f \circ s_{\mathrm{diag}}) = \deg_1(f^\pi) = \deg_1(f) = 1$ by Definition 4.3.2. Moreover, by definition of generalized diagonals, it follows that $|f \circ s_\zeta|$ and $|f \circ s_{\mathrm{diag}}|$ are equivariantly homotopic as maps $T^2 = S^1 \times S^1 \to S^2$, hence $\deg_1(f \circ s_\zeta) = \deg_1(f \circ s_{\mathrm{diag}}) = 1$ (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 4.2.3. □

The last puzzle piece we need to prove Lemma 5.1.1 (and thus to complete the proof of Theorem 5.1.1) is the following lemma which constructs a generalised diagonal of a special shape.

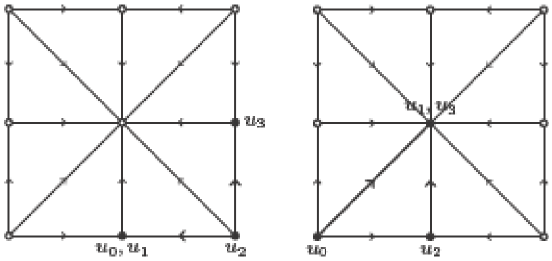

Figure 5.1: A path starting with the point $u_0 = (1, 0, 0, 0)$ shown as projection on the first two (left) and last two coordinates (right).

Lemma 5.1.3. *Let* $0 \leq h < \lfloor \frac{n-1}{3} \rfloor$*. Then there exists a generalised diagonal* $\zeta_0 \colon \Gamma_{3L} \to \Gamma_L^{n-1}$ *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 $\zeta_0 \colon \Gamma_{3L} \to \Gamma_L^{n-1}$, 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 $n-1-3h$).

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

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 bloc and decrease the value in the last bloc, 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. 5.1 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.,

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

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 $\zeta_0 \colon \Gamma_{3L} \to \Gamma_L^{n-1}$ 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). □

We may now finish the proof of Lemma 5.1.1 and, consequently, of Theorem 5.1.1.

Proof of Lemma 5.1.1. Let us fix a coordinate direction, without loss of generality $i = 1$, and let $f : \Gamma_L^n \to \Sigma^2$ 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 $n - h - 1$ 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 5.1.3 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 $\pi \colon [n-1] \to [n-1]$ 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 $n - h - 1$, 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 $s_\zeta \colon \Gamma_L \times \Gamma_{3L} \to \Gamma_L^n$ 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 5.1.2. 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 $\frac{1}{3L^2}$. □

5.2 Combinatorics of reconfigurations for $\mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4)$

The goal of this section is a careful combinatorial analysis of the binary polymorphisms. In particular, we will describe how the minion homomorphism $\xi\colon \mathrm{pol}(\mathbf{LO}_3, \mathbf{LO}_4) \to \mathscr{P}$ 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 *reconfigurable* one to the other 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 5.2.1. *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 $f = f_0, \ldots, f_k = g$ 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* for polymorphisms from K_d to K_{2d-2}.

Lemma 5.2.2. *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 $((1, 1), (x, x'), (y, y')) = ((1, 1), (1, y), (x', 1)) \in R^{\mathbf{LO}_3^2}$ 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. First, the value 1 is not possible since $((1, 2), (x, y), (x', y')) = ((1, 2), (1, 3), (x', 1)) \in R^{\mathbf{LO}_3^2}$, $f(x, y) = 1$, and $f(x', y') = 1$. $f(1, 2) = 2$ is not possible since $((1, 2), (1, 1), (x', y')) = ((1, 1), (1, 2), (x', 1)) \in R^{\mathbf{LO}_3^2}$, and $f(x', y') = 1, f(1, 1) = 2$. $f(1, 2) = 3$ is not possible since $((1, 2), (2, 2), (x, y)) = ((1, 2), (2, 2), (1, 3)) \in R^{\mathbf{LO}_3^2}$, and $f(x, y) = 1, f(2, 2) = 3$. Finally, $f(1, 2) < f(3, 3) = 4$, so $f(1, 2) \neq 4$.

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

$x' = y = 2$. We consider the pair of values $f(1, 3)$ and $f(3, 1)$. First, we have $f(1, 3) > f(1, 2) = f(x, y) = 1$ and $f(3, 1) > f(2, 1) = f(x', y') = 1$. As $((1, 3), (1, 1), (x', y')) = ((1, 3), (1, 1), (2, 1)) \in R^{\mathbf{LO}_3^2}$ and $f(1, 1) = 2, f(x', y') = 2$ we have that $f(1, 3) \neq 2$; symmetrically $f(3, 1) \neq 2$. We also have $f(1, 3) \neq 3$ since $((1, 3), (x, y), (2, 2)) = ((1, 3), (1, 2), (2, 2)) \in R^{\mathbf{LO}_3^2}$ and $f(1, 2) = 1, f(2, 2) = 3$; symmetrically $f(3, 1) \neq 2$. Thus $f(1, 3) = f(3, 1) = 4$. However, then $(f(1, 2), f(1, 3), f(3, 1)) = (1, 4, 4) \notin R^{\mathbf{LO}_4}$, which is not possible, as $((1, 2), (1, 3), (3, 1)) \in R^{\mathbf{LO}_3^2}$, which yields our contradiction.

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

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

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

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

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

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

This concludes the proof of the lemma. □

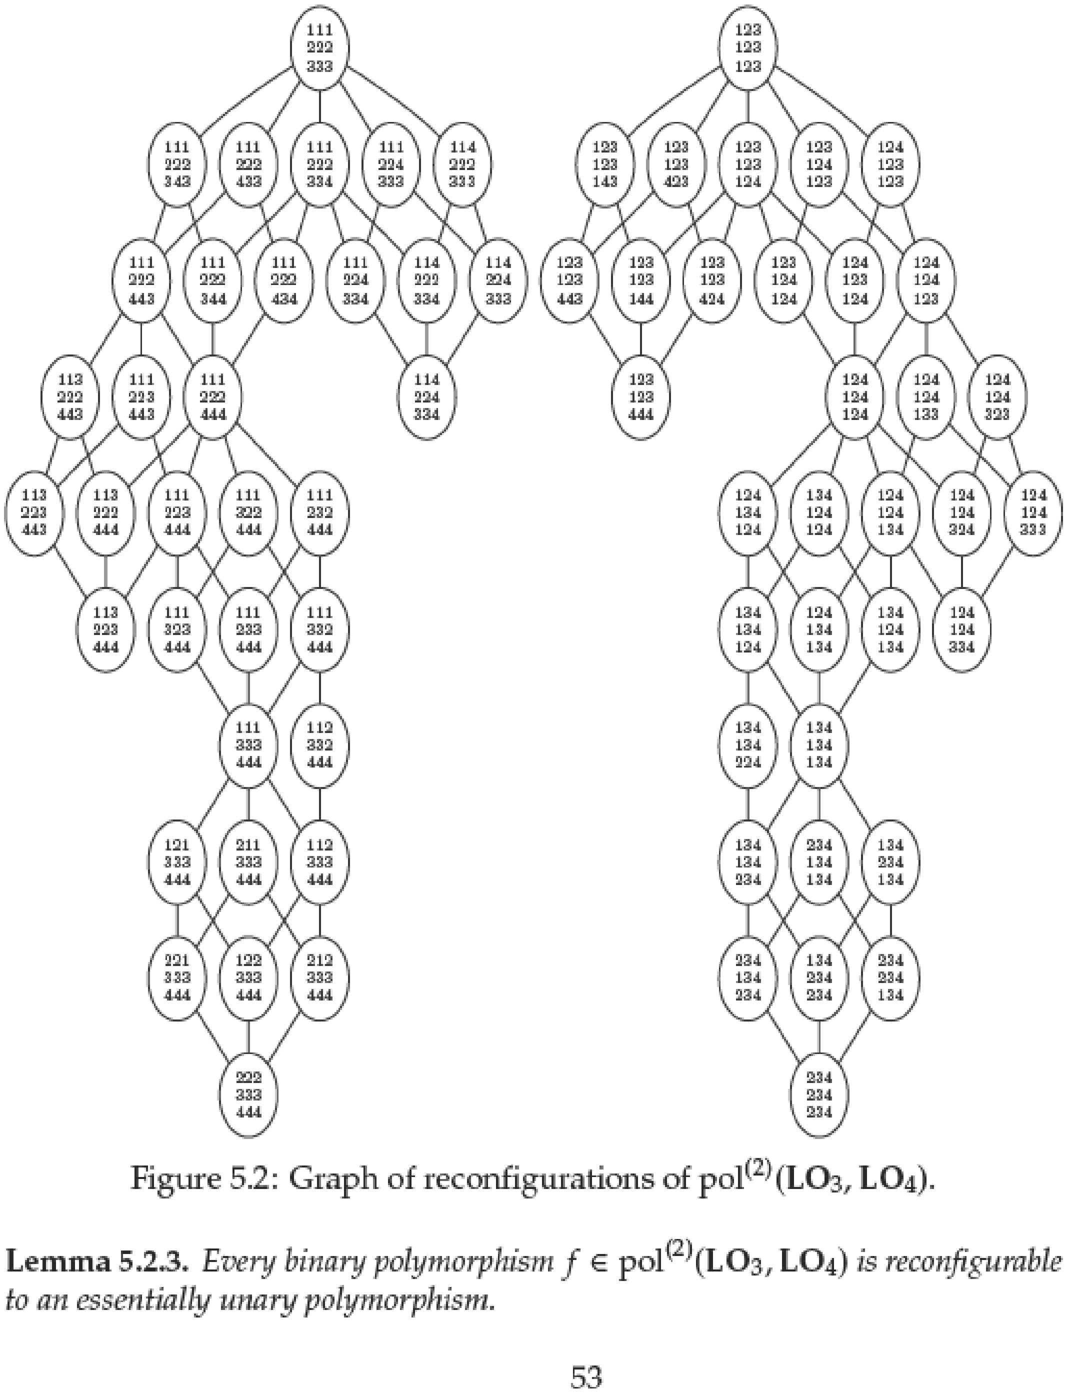

Figure 5.2: Graph of reconfigurations of $\mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$.

Lemma 5.2.3. *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 5.2.2. 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 5.2, we can see the reconfiguration graph of $\mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$. This shows how one can reconfigure all polymorphisms to essentially unary ones. In the diagram, we show a polymorphism in its matrix representation.

It can be also observed that unary polymorphisms that depend on the same coordinate are reconfigurable to each other. Moreover, since every connected component of $\mathrm{Hom}(\mathbf{LO}_3^2, \mathbf{LO}_4)$ 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. We have shown in Section 4.3 that is has at least two connected components (the two projections are not homotopic).

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

Lemma 5.2.4. *Let* $\mathbf{A}$*,* $\mathbf{B}$*, and* $\mathbf{C}$ *be three structures,* G *a group acting on* $\mathbf{A}$*, and assume that* $f, g \in \mathrm{hom}(\mathbf{B}, \mathbf{C})$ *are 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 $\mathrm{mhom}(\mathbf{A}, \mathbf{B}) \to \mathrm{mhom}(\mathbf{B}, \mathbf{C}) \to \mathrm{mhom}(\mathbf{A}, \mathbf{C})$ is monotone. This means that the composition extends linearly to a continuous map

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

(see also [Koz08, 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. □

The following corollary then follows directly from the above and Lemma 5.2.3.

Corollary 5.2.1. *For every binary polymorphism $f \in \mathrm{pol}^{(2)}(\mathbf{LO}_3, \mathbf{LO}_4)$, the induced map*

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

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 3.2.2 and the main theorem.

CHAPTER 6

Grünbaum mass partitioning problem

Geometric methods for partitioning space, point sets, or other geometric objects are a central topic in discrete and computational geometry. Partitioning results are often proved using topological methods and also play an important role in topological combinatorics [DLGMM19, Mat08, RS21]. As anticipated in Chapter 1, a classical example is the famous *Ham-Sandwich Theorem*, which goes back to the work of Steinhaus, Banach, Stone, and Tukey (see [RS21, Sec. 1] for more background and references). A "discrete" version of this theorem asserts that, given any d finite point sets $P_1, \dots, P_d$ in $\mathbb{R}^d$, there is an (affine) hyperplane H that *simultaneously bisects* all P_i, i.e., each of the two open half-spaces determined by H contains at most $|P_i|/2$ points, $1 \leq i \leq d$. This follows (by a standard limit argument, see [Mat08, Sec. 3.1]) from the following "continuous" version: Let $\mu_1, \dots, \mu_d$ be *mass distributions* in $\mathbb{R}^d$, i.e., finite measures such that every open set is measurable and every hyperplane has measure zero. Then there exists a hyperplane H such that $\mu_i(H^+) = \mu_i(H^-) = \frac{1}{2}\mu_i(\mathbb{R}^d)$ for $1 \leq i \leq d$, where H^+ and H^- are the two open half-spaces bounded by H.

In this paper, we are interested in another classical equipartitioning problem, first posed by Grünbaum [Grü60] in 1960: Given a mass distribution (respectively, a finite point set) in $\mathbb{R}^d$, can one find d hyperplanes that subdivide $\mathbb{R}^d$ into 2^d open orthants, each of which contains exactly $1/2^d$ of the mass (respectively, at most $1/2^d$ of the points)? We call such a d-tuple of hyperplanes a 2^d*-partition* of the mass distribution (respectively, of the point set).

For $d = 2$, it is an easy consequence of the planar Ham-Sandwich theorem that any mass distribution (or point set) in $\mathbb{R}^2$ admits a four-partition; moreover, the four-partition can be chosen such that one of the lines has a prescribed direction (indeed, start by choosing a first line in the prescribed direction that bisects the given mass distribution; by the Ham-Sandwich Theorem, there exists a second line that simultaneously bisects the two parts of the mass on either side of the first line). Alternatively, one can also show that there is always a four-partition such that the two lines are orthogonal. Intuitively, the reason that we can impose such additional conditions is that the four-partitioning problem in the plane is *underconstrained*: A line in the plane can be described by two independent parameters, so a pair of lines have four degrees of freedom, while the condition that the four quadrants have the same mass can be expressed by three equations, leaving one degree of freedom; either one of the additional constraints uses this extra degree of freedom.

In 1966, Hadwiger [Had66] gave an affirmative answer to Grünbaum's question for $d = 3$ and showed that any mass distribution in $\mathbb{R}^3$ admits an eight-partition; moreover, the normal vector of one of the planes can be prescribed arbitrarily. This result was later re-discovered by Yao, Dobkin, Edelsbrunner, and Paterson [YDEP89].

Theorem 6.0.1 ([Had66, YDEP89]). *Let μ be a mass distribution on $\mathbb{R}^3$, and let $v \in S^2$. Then there exists a triple of planes (H_1, H_2, H_3) that form an eight-partition for μ and such that the normal vector of H_1 is v.*

More recently, Blagojević and Karasev [BK16] gave a different proof for the existence of eight-partitions and showed the following variant:

Theorem 6.0.2 ([BK16]). *Let μ be a mass distribution on $\mathbb{R}^3$. Then there exists an eight-partition (H_1, H_2, H_3) of μ such that the plane H_1 is perpendicular to both H_2 and H_3.*

Our first result is the following alternative version of eight-partitioning, which to the best of our knowledge is new:

Theorem 6.0.3. *Given a mass distribution μ in $\mathbb{R}^3$ and a vector $v \in S^2$, there exists an eight-partition (H_1, H_2, H_3) of μ such that the intersection of the two planes H_1 and H_2 is a line in the direction of v.*

As in the case of the Ham-Sandwich Theorem, each of the three theorems above also implies the existence of the corresponding type of eight-partition for finite point sets, again by a standard limit argument (see Lemma B.1.1).

We remark that, in general, d hyperplanes in $\mathbb{R}^d$ are described by d^2 independent parameters, while the condition that 2^d orthants have equal mass can be expressed by $2^d - 1$ equations. For $d = 3$, this leaves $9 - 7 = 2$ degrees of freedom, which allows for any one of the additional conditions imposed in Theorems 6.0.1, 6.0.2, and 6.0.3, respectively. On the other hand, for $d \geq 5$, we have $d^2 < 2^d - 1$, so intuitively Grünbaum's problem is overconstrained. Avis [Avi84] made this precise and constructed explicit counterexamples using the well-known *moment curve* $\gamma = \{(t, t^2, \ldots, t^d) : t \in \mathbb{R}\}$ in $\mathbb{R}^d$. The crucial fact is that any hyperplane intersects the moment curve γ in at most d points ([Mat08, Lemma 1.6.4]). Thus, for $d \geq 5$, a mass distribution supported on γ admits no 2^d-partition because any d hyperplanes intersect γ in at most d^2 points, which subdivide γ into at most $d^2 + 1$ intervals, hence there are always at least $2^d - d^2 - 1 > 0$ orthants that do not intersect γ and hence contain no mass. The last remaining case $d = 4$ of Grünbaum's problem, i.e., the question whether any mass distribution in $\mathbb{R}^4$ admits a 16-partition by four hyperplanes, remains stubbornly open (see [BFH18], [DLGMM19, Conjecture 7.2], [Mat08, pp. 50–51], and [RS21, Problem 2.1.4] for more background and related open problems).

We now turn to the algorithmic question of computing eight-partitions in $\mathbb{R}^3$.

Problem 1. Given a set P of n points in $\mathbb{R}^3$, in sufficiently general position, compute three planes H_1, H_2, H_3 that form an eight-partition of the points.

As noted above, the corresponding problem of computing a four-partition of a planar point set can be reduced to finding a Ham-Sandwich cut of two planar point sets that are separated by a line. Megiddo [Meg85] showed that this can be done in linear time.

To characterize the complexity of Problem 1, we introduce the following concept. A *halving line* (resp., *halving plane*) for an n-point set in $\mathbb{R}^2$ (resp., $\mathbb{R}^3$) in general position is a line (resp., plane) that passes through two (resp., three) of the points and divides the remaining ones as equally as possible. Let $h_2(n)$ (resp., $h_3(n)$) denote the maximum number of halving lines (resp., planes) for an n-point set in $\mathbb{R}^2$ (resp., $\mathbb{R}^3$). The best known upper and lower bounds for $h_2(n)$ are $O(n^{4/3})$, due to Dey [Dey98], and $\Omega(ne^{\sqrt{\log n}})$, due to Tóth [Tót01], respectively. For $h_3(n)$, the best-known bounds are are $O(n^{5/2})$, due to Sharir, Smorodinsky, and Tardos [SST01], and $\Omega(n^2 e^{\sqrt{\log n}})$, due to Tóth [Tót01].

By a result of Lo, Matoušek, and Steiger [LMS94, Proposition 2], a Ham-Sandwich cut of n points in $\mathbb{R}^3$ can be computed in time $O^*(h_2(n)) = O^*(n^{4/3})$ (see also [HS17]), where the $O^*(\cdot)$-notation suppresses polylogarithmic factors. However, computing eight-partitions in $\mathbb{R}^3$ appears to be significantly more difficult. Unlike the planar four-partition problem, there is no known way of reducing it to the computation of a Ham-Sandwich cut. In particular, given two planes H_1 and H_2 that four-partition a finite point set P in $\mathbb{R}^3$ (in the sense that every one of the four open quadrants determined by H_1 and H_2 contains at most $|P|/4$ points), there generally need not exist a third plane H_3 such that H_1, H_2, H_3 form an eight-partition.

We note that, for fixed dimension $d \geq 3$, the best known algorithm for computing Ham-Sandwich cuts in $\mathbb{R}^d$ runs in time $O(n^{d-1-\alpha_d})$ where $\alpha_d > 0$ is a constant depending only on d [LMS94]. When the dimension is part of the input, a decision variant of the problem becomes computationally hard, see, e.g., [KTW11].

A brute-force algorithm that checks every triple of halving planes solves Problem 1 in time comparable to $O(h_3(n)^3) = O(n^{15/2})$. Yao et al. [YDEP89] and Edelsbrunner [Ede86] gave a $O(n^6)$-time algorithm that computes an eight-partition (with a prescribed normal direction for one of the planes, as in Theorem 6.0.1) by an exhaustive search, using the fact that only two planes need to be identified. Fixing one plane and performing a brute-force search for the remaining two would yield an algorithm with a running time comparable to $O(h_3(n)^2) = O(n^5)$.

Here, we present, to our knowledge, the fastest known algorithm for Problem 1. Roughly speaking, our algorithm runs in time near-linear in $h_3(n)$ rather than quadratic in it. Slightly more precisely, our algorithm runs in time near-linear in $nh_2(n)$, which is not known to be $o(h_3(n))$, but for which the best known upper bound is strictly stronger; see Theorem 8.2.1 and Fact 8.1.2:

Theorem 6.0.4 (Algorithm). *An eight-partition of n points in general position in $\mathbb{R}^3$, with a prescribed normal vector for one of the planes, can be computed in time $O^*(nh_2(n))$, hence $O^*(n^{7/3})$; here, the $O^*(\cdot)$-notation suppresses polylogarithmic factors.*

Our algorithm can be seen as a constructive version of Hadwiger's proof [Had66]. We start by bisecting the point set by a plane with a fixed normal direction, which partitions the initial point set into two subsets of "red" and "blue" points, respectively, of equal size. After that, our algorithm finds two more planes that simultaneously four-partition both the red and the blue points.

It remains an open question whether Theorem 6.0.2 or our own Theorem 6.0.3 can also be used to obtain an efficient algorithm for Problem 1. It would also be interesting to decide whether there is an algorithm for Problem 1 with running time $o(nh_2(n))$.

CHAPTER 7

Eight-partitioning lines with a fixed intersection direction

7.1 Notation and preliminaries

In what follows, it will often be convenient to assume that the mass distributions we work with have *connected support*, where the support of a mass distribution μ is $\mathrm{Supp}(\mu) := \{x \in \mathbb{R}^3 : \mu(B_r(x)) > 0 \text{ for every } r > 0\}$ and $B_r(x)$ denotes the ball of radius $r > 0$ centred at x.

By a standard limit argument (see Lemma B.1.2), the existence of eight-partitions for mass distributions with connected support implies the existence of eight-partitions for the general case. Hereafter, unless stated otherwise, we assume, without loss of generality, that every mass distribution has connected support.

We denote the *scalar product* of two vectors $x, y \in \mathbb{R}^3$ by $x \cdot y := \sum_{i=1}^{3} x_i y_i$. A vector $v \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ and a scalar $a \in \mathbb{R}$ determine an (affine) plane

$$H = H_v(a) := \{x \in \mathbb{R}^3 : x \cdot v = a\},$$

together with an orientation of H (given by the direction of the normal vector v). We denote by $-H := H_{-v}(-a)$ the affine plane with the same equation as H but with opposite orientation. The oriented plane H determines two open half-spaces, denoted by

$$H^+ := \{x \in \mathbb{R}^3 : x \cdot v > a\} \quad \text{and} \quad H^- := \{x \in \mathbb{R}^3 : x \cdot v < a\}.$$

More generally, let $\mathcal{H} = (H_1, \dots, H_k)$ be an ordered k-tuple of (oriented) planes in $\mathbb{R}^3$, $k \leq 3$. In what follows, it will be convenient to identify the set $\{+, -\}$ with the group $\mathbb{Z}_2$ (where the group operation is multiplication of signs). Elements of $\{+, -\}^k = \mathbb{Z}_2^k$ are strings of signs of length k, and we will denote by $\mathbf{+} = + \cdots +$ the identity element of $\mathbb{Z}_2^k$.

For $\alpha = (\alpha_1, \dots, \alpha_k) \in \mathbb{Z}_2^k = \{+, -\}^k$, we define the *open orthant* determined by $\mathcal{H}$ and α as $O_\alpha^{\mathcal{H}} := H_1^{\alpha_1} \cap \cdots \cap H_k^{\alpha_k}$. Given a mass distribution μ in $\mathbb{R}^3$, we say that an ordered k-tuple $\mathcal{H} = (H_1, \ldots, H_k)$ of planes ($k \leq 3$) forms a 2^k-*partition* of μ if every orthant contains $1/2^k$ of the mass, i.e., $\mu(O_\alpha^{\mathcal{H}}) = \mu(\mathbb{R}^3)/2^k$ for every $\alpha \in \{+,-\}^k$. For $k = 1, 2, 3$, this corresponds to the notions of bisecting, four-partitioning, and eight-partitioning μ as mentioned in the introduction. Analogously, we say that $\mathcal{H}$ forms a 2^k-*partition* of a *finite* point set P in $\mathbb{R}^3$ if $|P \cap O_\alpha^{\mathcal{H}}| \leq \frac{|P|}{2^k}$ for all α.

We will parameterize oriented planes in $\mathbb{R}^3$ by S^3, where the north pole e_4 and the south pole $-e_4$ map to the plane at infinity with opposite orientations. For this we embed $\mathbb{R}^3$ into $\mathbb{R}^4$ via the map $(x_1, x_2, x_3) \mapsto (x_1, x_2, x_3, 1)$. An oriented plane in $\mathbb{R}^3$ is mapped to an oriented affine 2-dimensional subspace of $\mathbb{R}^4$ and is extended (uniquely) to an oriented linear hyperplane. The unit normal vector on the positive side of the linear hyperplane defines a point on the sphere S^3. Hence, there is a one-to-one correspondence between points v in $S^3 \setminus \{e_4, -e_4\}$ and oriented affine planes H_v in $\mathbb{R}^3$. The positive side of the plane at infinity is $\mathbb{R}^3$ for $v = e_4$ and $\emptyset$ for $v = -e_4$. Hence $H_{-v}^+ = H_v^-$ for every v. Note that planes at infinity cannot arise as solutions to the measure partitioning problem, since they produce empty orthants. Therefore we do not need to worry about the fact that the sphere includes these.

We parameterize triples of planes (called *plane configurations*) in $\mathbb{R}^3$ by $(S^3)^3$, and denote by $\mathcal{H}_v$ the triple corresponding to $v \in (S^3)^3$. Given a mass distribution μ on $\mathbb{R}^3$, for each $v \in (S^3)^3$ and $\alpha \in {\mathbb{Z}_2}^3 \setminus \{+\}$, we set

$$F_\alpha(v, \mu) = \sum_{\beta \in \mathbb{Z}_2^3} (-1)^{p(\alpha,\beta)} \mu(O_\beta^{\mathcal{H}_v}).$$

where $p(\alpha, \beta)$ is the number of coordinates where both α and β are $-$. The functions F_α were also utilized in the proof of Theorem 6.0.1 in [YDEP89].

As an example, with $\mathcal{H} := \mathcal{H}_v = (H_1, H_2, H_3)$ and $\alpha = --+ \in \mathbb{Z}_2^3 \setminus \{+\}$, we obtain

$$\begin{aligned} F_{--+}(\mathcal{H}, \mu) &= \sum_{\beta \in \mathbb{Z}_2^3} (-1)^{p(\alpha,\beta)} \mu(O_\beta^{\mathcal{H}}) = \sum_{\beta \in \mathbb{Z}_2^3 :\, p(\alpha,\beta)=0} \mu(O_\beta^{\mathcal{H}}) - \sum_{\beta \in \mathbb{Z}_2^3 :\, p(\alpha,\beta)=1} \mu(O_\beta^{\mathcal{H}}) \\ &= \left(\mu(O_{+++}^{\mathcal{H}}) + \mu(O_{++-}^{\mathcal{H}}) + \mu(O_{--+}^{\mathcal{H}}) + \mu(O_{---}^{\mathcal{H}}) \right) \\ &\quad - \left(\mu(O_{-++}^{\mathcal{H}}) + \mu(O_{-+-}^{\mathcal{H}}) + \mu(O_{+-+}^{\mathcal{H}}) + \mu(O_{+--}^{\mathcal{H}}) \right) \\ &= \mu(H_1^+ \cap H_2^+) + \mu(H_1^- \cap H_2^-) - \mu(H_1^- \cap H_2^+) - \mu(H_1^+ \cap H_2^-). \end{aligned}$$

When μ is clear from context, we write $F_\alpha(\mathcal{H})$ instead of $F_\alpha(\mathcal{H}, \mu)$. The definitions of alternating sums for a pair of planes or a single plane are analogous.

The alternating sums have the following properties which will play an important role in the proof below.

Observation 3. Let μ be a mass distribution and fix $k = 2, 3$.

(i) Let $\alpha \in {\mathbb{Z}_2}^{k-1} \setminus \{+\}$ and let $\mathcal{H} = (H_1, \ldots, H_k)$ be a k-tuple of planes. Then $F_{+\alpha}(\mathcal{H}) = F_\alpha((H_2, \ldots, H_k))$ (the equivalent statement holds for any other entry of a k-tuple $(\alpha_1, \cdots, \alpha_k)$ instead of just for α_1).

(ii) A k-tuple $\mathcal{H}$ of planes 2^k-partitions if and only if $F_\alpha(\mathcal{H}) = 0$ for every $\alpha \in {\mathbb{Z}_2}^k \setminus \{+\}$.

Proof. (i) Since every hyperplane has null measure it follows that, for any $\beta \in {\mathbb{Z}_2}^{k-1}$

$$\mu(O_\beta^{(H_2,\ldots,H_k)}) = \mu(O_{+\beta}^{\mathcal{H}}) + \mu(O_{-\beta}^{\mathcal{H}}).$$

If $\tilde{\alpha} = +\alpha$ then, from the definition of $F_{\tilde{\alpha}}$, for any $\beta \in {\mathbb{Z}_2}^{k-1}$ the two orthants $O_{+\beta}^{\mathcal{H}}$ and $O_{-\beta}^{\mathcal{H}}$ are counted with the same sign in the sum, therefore

$$F_{\tilde{\alpha}}(\mathcal{H}) = \sum_{\substack{\beta\in{\mathbb{Z}_2}^{k-1}\\ p(\alpha,\beta)=0}} \left(\mu(O_{+\beta}^{\mathcal{H}}) + \mu(O_{-\beta}^{\mathcal{H}})\right) - \sum_{\substack{\beta\in{\mathbb{Z}_2}^{k-1}\\ p(\alpha,\beta)=1}} \left(\mu(O_{+\beta}^{\mathcal{H}}) + \mu(O_{-\beta}^{\mathcal{H}})\right) = F_\alpha((H_2,\ldots,H_k)).$$

((ii)) It is clear that, if $\mathcal{H}$ is a 2^k-partition, then all the alternating sums are 0. We will prove the other implication.

Suppose first that $k = 1$ and that $\mathcal{H} = H$. The only alternating sum is $F_-(H) = \mu(H^+) - \mu(H^-)$ and $F_-(H) = 0$ implies that H bisects μ.

Suppose now that $k = 2$ and that $\mathcal{H} = (H_1, H_2)$. By ((i)) and the statement for a single plane, $F_{+-}(\mathcal{H}) = 0$ and $F_{-+}(\mathcal{H}) = 0$ imply that both H_1 and H_2 bisect. Therefore, if $\lambda := \mu(O_{++}^{\mathcal{H}})$, we have that

$$0 = F_{--}(\mathcal{H}) = \mu(O_{++}^{\mathcal{H}}) + \mu(O_{--}^{\mathcal{H}}) - \mu(O_{-+}^{\mathcal{H}}) - \mu(O_{+-}^{\mathcal{H}}) = \lambda + \lambda - (\frac{1}{2} - \lambda) - (\frac{1}{2} - \lambda) = 4\lambda - 1;$$

hence $\lambda = \frac{1}{4}$ as desired.

Finally, suppose that $k = 3$ and that $\mathcal{H} = (H_1, H_2, H_3)$. By ((i)) and the statement for single planes and for pairs of planes, we have that all planes H_i bisect and all pairs (H_i, H_j) four-partition. Therefore, if $\lambda := \mu(O_{+++}^{\mathcal{H}})$, we have that

$$\begin{aligned} 0 = F_{---}(\mathcal{H}) &= \mu(O_{+++}^{\mathcal{H}}) + \mu(O_{+--}^{\mathcal{H}}) + \mu(O_{-+-}^{\mathcal{H}}) + \mu(O_{--+}^{\mathcal{H}}) \\ &\quad - \mu(O_{-++}^{\mathcal{H}}) - \mu(O_{+-+}^{\mathcal{H}}) - \mu(O_{++-}^{\mathcal{H}}) - \mu(O_{---}^{\mathcal{H}}) \\ &= \lambda + \lambda + \lambda + \lambda - (\frac{1}{4} - \lambda) - (\frac{1}{4} - \lambda) - (\frac{1}{4} - \lambda) - (\frac{1}{4} - \lambda) = 8\lambda - 1; \end{aligned}$$

hence $\lambda = \frac{1}{8}$ as claimed. □

7.1.1 The main topological result

Our goal is to prove the following result, which is a more precise statement of Theorem 6.0.3:

Theorem 7.1.1. *Given a mass distribution μ and a direction $u \in S^2$, there exists a triple $\mathcal{H} = (H_1, H_2, H_3)$ of oriented planes that eight-partition μ so that the oriented direction of the intersection $H_1 \cap H_2$ is u.*

By Lemma B.1.2, it is sufficient to prove Theorem 7.1.1 for mass distributions with connected support. We require the following technical lemma about partitioning a mass distribution on $\mathbb{R}^2$, due to Blagojević and Karasev [BK16]. For completeness, the proof is given in Appendix B.1.2.

Lemma 7.1.1 (Four-partitioning a mass distribution in $\mathbb{R}^2$ [BK16]). *Let $\mu^\#$ be a mass distribution (with connected support) on $\mathbb{R}^2$ and $v \in S^1$. Then there exists a pair (ℓ_1, ℓ_2) of lines in $\mathbb{R}^2$ that four-partitions $\mu^\#$ and such that v bisects the angle between ℓ_1 and ℓ_2.*

Moreover, if we orient ℓ_1 and ℓ_2 so that ℓ_1 is in the first direction clockwise from v, and ℓ_2 is in the first direction counterclockwise, the oriented pair is unique and the lines depend continuously on v.

Proof of Theorem 7.1.1. Without loss of generality, let $u = (0, 0, 1)$. Our proof proceeds in two steps. In the first step, we construct a map $\Phi \colon S^1 \times S^3 \to \mathbb{R}^4$ whose zeros codify eight-partitions of μ; then we prove that Φ is equivariant with respect to a suitable choice of actions of $G := \mathbb{Z}_4 \times \mathbb{Z}_2$ on the two spaces. In the second step we show that any continuous G-equivariant map $\Psi \colon S^1 \times S^3 \to \mathbb{R}^4$ has to have a zero.

Step 1: The key step in constructing the map Φ is to show that we can parameterize pairs of planes that have intersection direction u and four-partition μ, by a vector in S^1.

We project μ onto $u^\perp$, the plane orthogonal to u, to obtain a mass distribution $\mu^\#$ on $\mathbb{R}^2$. Specifically, identifying the plane with $\mathbb{R}^2$, let $A \subseteq \mathbb{R}^2$ and set $A \times \mathbb{R}$ to be the cylinder over A in the u-direction. Then $\mu^\#(A) = \mu(A \times \mathbb{R})$.

Let $v \in S^1 \subseteq \mathbb{R}^2$. By Lemma 7.1.1, there are two oriented lines $\ell_1(v)$ and $\ell_2(v)$ (that we can interpret as points $(\ell_1^i, \ell_2^i, a^i) \in S^2$) in the plane $u^\perp$ such that v bisects the angle between the two and ℓ_1, ℓ_2 four-partition the projected measure $\mu^\#$. Define $H_i(v) := (\ell_1^i(v), \ell_2^i(v), 0, a^i(v)) \in S^3$ to be $\ell_i(v) \times \mathbb{R}$ the (oriented) span of $\ell_i(v)$ and u; the two planes now four-partition μ and have the desired intersection direction.

Now let g_1 be a generator of $\mathbb{Z}_4 \times \{+\} \subseteq G$ and define its action on S^1 by a counterclockwise rotation by $\frac{\pi}{2}$. We use $g_1 \cdot v$ to denote the action of g_1 on v. Then, by the uniqueness in Lemma 7.1.1, we have that (see Figure 7.1):

$$\vec{\ell}_1(g_1 \cdot v) = \vec{\ell}_2(v) \quad \text{and} \quad \vec{\ell}_2(g_1 \cdot v) = -\vec{\ell}_1(v). \tag{7.1}$$

Using this construction, we can define a function $S^1 \to S^3 \times S^3$ by $v \mapsto (H_1(v), H_2(v))$. It follows from eq. (7.1) that $g_1 \cdot v$ is mapped to $(H_2(v), -H_1(v))$. Therefore, if we fix the corresponding action[1] of $\mathbb{Z}_4$ on $S^3 \times S^3$, the map is $\mathbb{Z}_4$-equivariant.

The group $\{e\} \times \mathbb{Z}_2$ acts by antipodality on S^3; therefore, if G acts on $(S^3 \times S^3) \times S^3$ component-wise, the map $\Phi \colon S^1 \times S^3 \to (S^3 \times S^3) \times S^3$ defined as $\Phi(v, w) := (H_1(v), H_2(v), H_w)$ is G-equivariant.

[1]Formally, for any $(x, y) \in S^3 \times S^3$ the generator g_1 of $\mathbb{Z}_4 \times \{+\} \subseteq G$ acts by $g_1 \cdot (x, y) = (y, -x)$.

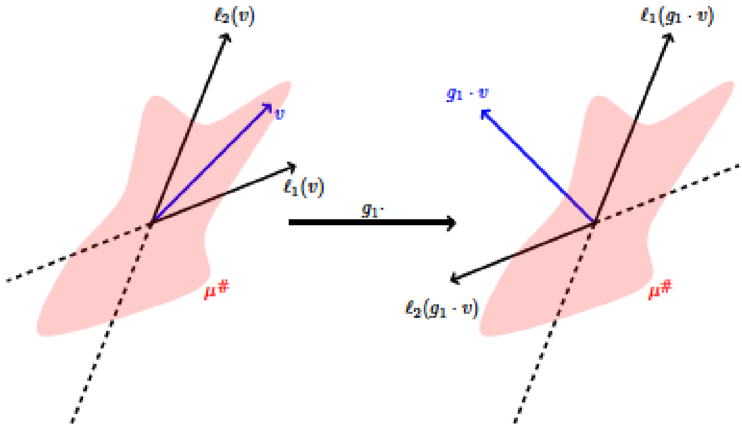

Figure 7.1: Example of the action of g_1.

By construction, the first two planes are always a four-partition of the mass distribution, therefore by Observation 3, a configuration $\Phi(v,w)$ is an eight-partition if and only if the four alternating sums with $\alpha_3 = -$ (i.e., $\alpha = ++-$, $-+-$, $+--$ and $---$) are 0.

To compute the action of G on the alternating sums, it is enough to specify what happens on g_1 (a generator of $\mathbb{Z}_4 \times \{e\}$) and g_2 (the generator of $\{e\} \times \mathbb{Z}_2$). Recalling that $g_1 \cdot \Phi(v,w) = (H_2(v), -H_1(v), H_w)$ and applying Observation 3((i)), we obtain

$$
\begin{aligned}
&F_{+--}(g_1 \cdot \Phi(v,w)) = F_{+--}((H_2(v), -H_1(v), H_w)) = F_{--}((-H_1(v), H_w)) \\
&= \mu(-H_1(v)^+ \cap H_w^+) + \mu(-H_1(v)^- \cap H_w^-) - \mu(-H_1(v)^- \cap H_w^+) - \mu(-H_1(v)^+ \cap H_w^-) \\
&= \mu(H_1(v)^- \cap H_w^+) + \mu(H_1(v)^+ \cap H_w^-) - \mu(H_1(v)^+ \cap H_w^+) - \mu(H_1(v)^- \cap H_w^-) \\
&= -F_{-+-}(\Phi(v,w))
\end{aligned}
$$

Similar computations imply that, if we act with g_1, we obtain

$$
\begin{aligned}
&F_{++-}(g_1 \cdot \Phi(v,w)) = F_{++-}(\Phi(v,w)), \\
&F_{-+-}(g_1 \cdot \Phi(v,w)) = F_{+--}(\Phi(v,w)), \text{ and} \\
&F_{---}(g_1 \cdot \Phi(v,w)) = -F_{---}(\Phi(v,w)),
\end{aligned}
$$

while acting with g_2 produces

$$
\begin{aligned}
&F_{++-}(g_2 \cdot \Phi(v,w)) = -F_{++-}(\Phi(v,w)), \\
&F_{+--}(g_2 \cdot \Phi(v,w)) = -F_{-+-}(\Phi(v,w)), \\
&F_{-+-}(g_2 \cdot \Phi(v,w)) = -F_{+--}(\Phi(v,w)), \text{ and} \\
&F_{---}(g_2 \cdot \Phi(v,w)) = -F_{---}(\Phi(v,w)),
\end{aligned}
$$

for every $(v,w) \in S^1 \times S^3$.

Finally, we can choose a linear G-action on $\mathbb{R}^4$ that is consistent with the previous equations. In particular, if we define

$$
g_1 \cdot (x,y,z,u) = (x,-z,y,-u) \quad \text{and} \quad g_2 \cdot (x,y,z,u) = (-x,-y,-z,-u),
$$

then the map $\Psi\colon S^1 \times S^3 \to \mathbb{R}^4$, given by

$$(v, w) \mapsto (F_{++-}(v, w), F_{+--}(v, w), F_{-+-}(v, w), F_{---}(v, w))$$

is G-equivariant, i.e. $\Psi(g_1 \cdot v, w) = g_1 \cdot \Psi(v, w)$ and $\Psi(v, -w) = -\Psi(v, w)$. By Observation 3, the zeros of Ψ are exactly the configurations of planes that eight-partition the measure and have the desired intersection property.

Step 2: Suppose now, for a contradiction, that Ψ does not have a zero. This means that it is possible to define a G-equivariant map $\overline{\Psi}\colon S^1 \times S^3 \to S^3$ by $\overline{\Psi}(v, w) := \frac{\Psi(v,w)}{\|\Psi(v,w)\|}$.

Denote by Ψ_a, for $a \in S^1$, the map $\Psi_a\colon S^3 \to S^3$, $\Psi_a(w) = \overline{\Psi}(a, w)$; this function has two key properties:

(i) for any $a \in S^1$, Ψ_a is antipodal;

(ii) for any $a, b \in S^1$, let $\gamma : [0, 1] \to S^1$ be a parametrization of the arc between $\gamma(0) = a$ and $\gamma(1) = b$. Then Ψ_a and Ψ_b are homotopic via $H : S^3 \times [0, 1] \to S^3$ with $H(t, x) = \Psi_{\gamma(t)}(x)$.

For any $n \geq 1$, the group of orthogonal matrices $O(n)$ contains exactly two connected components, distinguished by the sign of the determinant. Since the map $g_1 : S^3 \to S^3$ is induced by a matrix with determinant -1, it is homotopic to any other orthogonal linear map with determinant -1. In particular, it is homotopic to the reflection r along the last coordinate and, thus, $\deg(g_1) = \deg(r) = -1$. Combining everything together we have:

$$\deg(\Psi_a) = \deg(\Psi_{g_1 \cdot a}) = \deg(g_1 \cdot \Psi_a) = \deg(g_1)\deg(\Psi_a) = -\deg(\Psi_a).$$

Hence $\deg(\Psi_a) = 0$, contradicting the Borsuk-Ulam theorem (see [Mat08, Theorem 2.1.1]). □

Theorem 7.1.1, along with Lemma B.1.1, immediately implies the following.

Theorem 7.1.2. *Let $P \subseteq \mathbb{R}^3$ be a finite set of points and $p \in S^2$ a fixed direction. Then there exists a triple $\mathcal{H} = (H_1, H_2, H_3)$ of oriented planes that eight-partitions P, so that the oriented direction of the intersection $H_1 \cap H_2$ is p.*

CHAPTER 8

An algorithm for eight-partitioning points in 3D

8.1 Levels in arrangements of planes

Let $P \subseteq \mathbb{R}^3$ be a set of n points in general position. Specifically, we assume that the points in P satisfy the following: no four in a plane, no three in a vertical plane, and no two in a horizontal plane. Recall that a *halving plane* for a point set in $\mathbb{R}^3$ in general position is a plane that passes through three of the points and divides the remaining points as equally as possible; $h_3(n)$ is the maximum number of halving planes for an n-point set in $\mathbb{R}^3$.

The duality transform maps points in $\mathbb{R}^3$ to planes in $\mathbb{R}^3$ and vice versa. Specifically, the point $p = (p_1, p_2, p_3) \in \mathbb{R}^3$ is mapped to the non-vertical plane $p^\star\colon z = p_1x + p_2y - p_3$ in $\mathbb{R}^3$, and vice versa. See [Har11, Chapter 25.2] for standard properties of the duality transform.

Let H be a set of planes in $\mathbb{R}^3$ in general position. Specifically, we assume that the planes are non-vertical, every triple of planes in H meets in a unique point, and no point in $\mathbb{R}^3$ is incident with more than three of the planes. The planes in H partition $\mathbb{R}^3$ into a complex of convex cells, called the *arrangement* of H and denoted by $\mathcal{A}(H)$. The *k-level* in $\mathcal{A}(H)$ is defined as the closure of the set of all points which lie on a unique plane of the arrangement and have exactly $k-1$ planes below it. Note that the k-level is a piecewise linear surface in $\mathbb{R}^3$ whose faces are contained in planes of H. The *complexity* of a level is the total number of vertices, edges and faces contained in the level. When $k = \lfloor(|H| + 1)/2\rfloor$, the k-level is called the *median level* of the arrangement.

Duality, k-levels, and complexity of levels are defined analogously in $\mathbb{R}^2$. Let $g_2(n)$ (resp., $g_3(n)$) be the maximum complexity of any k-level in an arrangement of n planes in $\mathbb{R}^2$ (resp., $\mathbb{R}^3$). It is well-known that $h_2(n) = \Theta(g_2(n))$ and that $h_3(n) = \Theta(g_3(n))$ (see [AAH+98, Theorem 3], [Ede87, Theorem 3.3]).

The main object of interest in bounding the complexity of our algorithm is the intersection of median levels of two arrangements of disjoint sets of planes. We first show that the complexity of this is proportional to $h_3(n)$, in the worst case.

Fact 8.1.1. *Let $\ell(n)$ be the maximum complexity of the intersection of a level in the arrangement $\mathcal{A}(R)$ and a level in the arrangement $\mathcal{A}(B)$ of disjoint sets R and B of planes in $\mathbb{R}^3$, with $H := A \cup B$ in general position and $n := |H|$. Then $\ell(n) = \Theta(h_3(n))$.*

Proof. Let L be the intersection of level k of $\mathcal{A}(R)$ and level k' of $\mathcal{A}(B)$. As the planes of H are in general position, L is (a disjoint union of) a collection of (open) edges and vertices in $\mathcal{A}(R \cup B)$. In fact, L is a collection of cycles and bi-infinite curves so its complexity is asymptotically determined by the number of its edges.

A point on an edge of L has the property that it lies on one plane of R and one plane of B, and has exactly $k + k' - 2$ planes of $R \cup B$ below it. Hence L is contained in the $(k + k' - 1)$-level of $\mathcal{A}(R \cup B)$. It follows that the complexity of L is bounded above by the complexity of a level in an arrangement of n planes, implying $\ell(n) \leq g_3(n) = O(h_3(n))$. This proves the upper bound.

For the lower bound, suppose first that n is of the form $4k + 6$ and let $P \subseteq \mathbb{R}^3$ be a set $n/2 = 2k + 3$ points in general position achieving the maximum number $h_3(2k + 3)$ of halving planes. Let P' be a copy of P translated by a sufficiently small distance $\epsilon > 0$ in a generic direction. We then slightly perturb the points of $P \cup P'$ to ensure general position. For a point $p \in P$, we denote by p' its copy (or *twin*) in P'. Consider the sets R^* (red) and B^* (blue) of k points each obtained as follows: for each pair (p, p'), we assign one point to R^* and the other to B^* uniformly and independently at random.

Recall that a halving plane is defined by a triple of points and divides the remaining points as evenly as possible. Let π_1 be a halving plane of P defined by a triple (a, b, c), so that π_1 has exactly k points of P on each side. Consider the set $S := \{a, b, c, a', b', c'\}$. We claim that, with constant probability, there exists a plane π_2 that passes through a red point and a blue point of S that are not twins, and has precisely one red and one blue point of S on each side.

As our perturbation ϵ is arbitrarily small, π_2 partitions the remaining points of $R^* \cup B^*$ in the same manner as π_1 partitions P. Let R and B be the sets of planes dual to points in R^* and B^*, respectively. In the dual, the plane π_2 corresponds to a point π_2^* on an edge of the arrangement $\mathcal{A}(R \cup B)$ that lies on level k of $\mathcal{A}(R)$ and of $\mathcal{A}(B)$. Therefore it lies on the curve L of intersection of two k-levels and this curve contains all the points π_2^* (as we range over halving planes π_1 of P).

Since different halving planes correspond to partitioning P in different ways, the planes π_2 partition $P \cup P'$ in different ways and, hence, points π_2^* arising from different halving planes of P lie on different edges of $\mathcal{A}(R \cup B)$. By construction, the number of edges of L is $\Omega(h_3(n))$ in expectation, completing the lower bound proof for $n = 4k + 6$.

Finally, for a general n, we write $n = 4k + 6 + c$, for $0 \leq c \leq 5$, we apply the above construction to $n - c$ and, at the end, add c planes in general position to the set R lying above all vertices of $\mathcal{A}(R \cup B)$. It is easily checked that this addition does not reduce the size of L. □

Note, however, that in our application, the sets of points R^* and B^* are strictly separated, which is not the case in the above lower-bound construction. We show that, with this additional constraint, the complexity of L is $O(nh_2(n))$ and that there exist pairs of separated point sets R^* and B^* that achieve this bound.

Fact 8.1.2. *Let $\ell^*(n)$ be the maximum complexity of the intersection of two levels in $\mathcal{A}(R)$ and $\mathcal{A}(B)$ as in Fact 8.1.1, with the additional constraint that the dual sets R^* and B^* are strictly separated by a plane. Then $\ell^*(n) = \Theta(nh_2(n))$.*

Proof. Upper bound: Let L be the intersection of level k in $\mathcal{A}(R)$ and level k' in $\mathcal{A}(B)$. Without loss of generality, we consider vertices of L that are the intersection of one plane in R and two planes in B. Such a vertex has $k-1$ planes of R and either $k-1$ or $k-2$ planes of B below it.

We work with the dual point sets R^* (red) and B^* (blue), strictly separated by the plane π. Here, a vertex as above corresponds to a plane passing through one red point and two blue points, with $k-1$ red points and $k-1$ or $k-2$ blue points above it. We count all such planes σ_a passing through a fixed red point $a \in R^*$.

Let B' be a set of points obtained by radially projecting points of B^* onto π with center a. Let ℓ_a be the line $\pi \cap \sigma_a$. Since R^* and B^* are separated by π, ℓ_a has $k-1$ or $k-2$ points of B' on one side of it. Hence, the dual of ℓ_a is a vertex on level k (or its complement) in the planar arrangement of the dual of B' in π. Therefore, there are at most $2g_2(|B'|)$ choices for the plane σ_a, and at most $|R| \cdot 2g_2(|B|) = O(nh_2(n))$ such planes overall.

Planes passing through exactly one blue point and two red points are handled symmetrically, completing the proof of the upper bound.

Lower bound: Once again, it is sufficient to make the argument for n of the form $4k+3$ for a positive integer k; the general case is handled as in the lower bound proof of Fact 8.1.1. Consider a set R^* of $2k+2$ points realizing $h_2(2k+2)$ lying in the xy-plane, scaled to fit in the rectangle $(0,1) \times (0,\epsilon)$, for a sufficiently small $\epsilon > 0$.

Let B^* be a set of $2k+1$ points equally spaced on a unit circle in the plane $x = 0$, centred at the origin, so that no point lies in the xy-plane. Note that R^* and B^* are separated by the plane $x = \delta$, for a sufficiently small $\delta > 0$.

By making ϵ small enough, we can ensure that any halving line of R^* passes arbitrarily close to the origin. In particular, any pair of a halving line of R^* and a point of B^* define a halving plane that passes through two points of R^* and one of B^*, and has exactly $2k$ points of $R^* \cup B^*$ on each side. The number of such halving planes is $(2k+1)h_2(2k+2) = \Omega(nh_2(n))$. Finally, the points of $R^* \cup B^*$ can be perturbed to satisfy the general position assumption without reducing this number.

So we have constructed two separated sets of points R^* and B^* with the property that there are $\Omega(nh_2(n))$ halving planes spanned by three of the points that simultaneously bisect both sets. The lower bound follows by considering the dual sets R and B. □

8.2 The algorithm

We can deduce the existence of eight-partitions of a finite point set $P \subset \mathbb{R}^3$ of a certain advantageous form from Theorem 6.0.1.

Observation 4. Let $k > 0$ be an integer and $P \subseteq \mathbb{R}^3$ be a set of $n = 8k+7$ points in general position. Then, there exists a triple of planes (H_1, H_2, H_3) that eight-partitions P with the following properties:

(i) H_1 is horizontal (i.e., parallel to the xy-plane) and passes through the z-median point of P. From here on, we refer to the $4k+3$ points that lie below (resp., above) H_1 as *red* (resp., *blue*) points and denote the sets R (resp. B).

(ii) H_2 and H_3 each contain exactly three points, and each open octant contains exactly k points.

(iii) H_2, H_3 each bisect R and B, and the pair (H_2, H_3) four-partitions both R and B. Furthermore, H_2 and H_3 contain at least one point of each color.

Proof. Since the set $X := \{(H_1, H_2, H_3) : H_1 \text{ is horizontal}\} \subset (S^3)^3$ is compact, by Theorem 6.0.1 and Lemma B.1.1, there exists a configuration $\mathcal{H}_\infty = (H_1, H_2, H_3)$ that eight-partitions the point set with H_1 horizontal. Along with the general position assumption, this implies that H_1 contains only the z-median point. This proves ((i)).

To see ((ii)), note that any eight-partition has at most k points of P in each of the eight open octants, one point in H_1, and at most three points in each of H_2 and H_3, by general position, for a total of at most $8k+1+2\cdot 3 = 8k+7 = n$ points. So, in fact, all the inequalities are equalities: there *must* be *exactly* k points in each open quadrant and exactly three points of $R \cup B$ in each of H_2 and H_3.

It remains to show ((iii)). By the preceding paragraph, it is straightforward to see that (H_2, H_3) four-partitions both R and B. By Corollary B.1.1, we have that any pair (H_i, H_j) four-partitions P. Since (H_1, H_2) four-partitions P, each quadrant formed by (H_1, H_2) has at most $\lceil (8k+7)/4 \rceil = 2k+1$ points. H_1 has $4k+3$ points on each side. Hence, we obtain that H_2 bisects R and B, and, in particular, contains at least one point of each color. A symmetric argument shows that H_3 bisects both R and B, and contains at least one point of each color. This completes the proof. □

Theorem 8.2.1 (Computation of an eight-partition). *Let $P \subseteq \mathbb{R}^3$ be a set of $n > 0$ points in general position and $v \in S^2$. An eight-partition (H_1, H_2, H_3) of P, with v being the normal vector of H_1, can be computed in time $O^*(n + \ell^*(n))$.*

Remark 4. Since $\ell^*(n) = \Theta(nh_2(n)) = O(n^{7/3})$ by Fact 8.1.2, we can compute an eight-partition in time $O^*(n^{7/3})$.

The rest of this section is devoted to the proof of Theorem 8.2.1. We assume, without loss of generality, that $v = e_3 = (0, 0, 1)$ is the vertical vector, so H_1 is required to be horizontal. If $n \leq 7$, the statement holds trivially — set H_1 to be the horizontal plane containing any point of P, and H_2, H_3 to contain at most three distinct points each, so that the octants do not contain any points. From here on, we will assume that $n = 8k + 7$, for an integer $k > 0$. If n is not of this form, we may add dummy points to P (in general position) until the number of points is of the required form and run the algorithm. Once the algorithm terminates, we discard the dummy points, resulting in an eight-partition with at most k points in each octant.

We now describe the algorithm to construct an eight-partition of P satisfying the properties in Observation 4. Let H_1 be the horizontal plane containing the z-median point of P, and, without loss of generality, identify H_1 with the xy-plane. Now consider the sets R and B of $4k + 3$ points each lying below and above, respectively, H_1. We further assume, without loss of generality, that B is contained in the half-space $x < 0$ and R is contained in the half-space $x > 0$. Otherwise, since no point in $R \cup B$ has $z = 0$ by the general position assumption, there exists a plane H containing the y-axis and with sufficiently small negative slope in the x direction such that all red points are below H and all blue points are above H. Applying a generic sheer transformation (so as not to violate the general position assumption) that fixes the xy-plane and maps H to the plane $x = 0$, we obtain point sets with the required properties.

Let $R^* = \{p^* : p \in R\}$ be the set of red planes dual to points in R and set $\mathcal{A}(R) := \mathcal{A}(R^*)$ to be the arrangement formed by the set R^*. The set of blue planes B^* and the blue arrangement $\mathcal{A}(B)$ are defined analogously. We will write $\mathcal{A} := \mathcal{A}(R \cup B)$ for the arrangement formed by the planes in $R^* \cup B^*$. For a (dual) point $p \in \mathbb{R}^3$, we set $R_p^+, R_p^- \subseteq R^*$ to be the set of red planes lying strictly above and below p, respectively. For a pair p, q of (dual) points, put

$$X(p,q) := |R_p^+ \cap R_q^+| - |R_p^+ \cap R_q^-| - |R_p^- \cap R_q^+| + |R_p^- \cap R_q^-|.$$

The sets $B_p^+, B_p^- \subseteq B^*$ and the function $Y(p, q)$ are defined analogously for B^*.

Let L be the intersection of the median levels of $\mathcal{A}(B)$ and $\mathcal{A}(R)$. Let m be the complexity, i.e., the number of vertices and edges, of L. By the following lemma, we have that L is a connected y-monotone polygonal curve and is an alternating sequence of edges and vertices of $\mathcal{A}$ terminated by half-lines.

Lemma 8.2.1. *L is a connected y-monotone curve.*

Proof. Recall that B lies in the quadrant $x < 0, z < 0$ and R lies in the quadrant $x > 0, z > 0$. Hence, the dual planes in B^* and R^* have equations of the form $z = ax + by + c$ with $a < 0$ and $a > 0$, respectively.

Consider the intersection of $R^* \cup B^*$ with the plane $\Pi_d : y = d$. The intersection of the plane $z = ax + by + c$ is the line $z = ax + (bd + c)$, so planes in B^* correspond to lines with negative slope and planes in R^* correspond to lines with positive slope. In particular, the median levels of lines corresponding to B^* and R^* are graphs of piecewise-linear total functions that are decreasing and increasing, respectively. It follows that the two curves intersect exactly once. This intersection point corresponds to the intersection of L with the plane Π_d.

By general position, L is a union of vertex-disjoint cycles and bi-infinite paths composed of edges and vertices of $\mathcal{A}$, since incident to every vertex of A contained in L are precisely two edges belonging to L. By monotonicity in y, L must be a single bi-infinite chain. □

In fact, L can be computed efficiently using standard tools [AM95, Cha10], which we outline now.

Lemma 8.2.2 (Computing the intersection of two levels [AM95, Cha10]). *The intersection curve L of the median level of $\mathcal{A}(B)$ and the median level of $\mathcal{A}(R)$ can computed in time $O^*(n+m)$, where m is the complexity of the curve and $O^*(\cdot)$ notation hides polylogarithmic factors.*

Proof. We use the standard dynamic data structure for ray-shooting queries in the intersection of half-spaces; the currently fastest algorithm is due to Chan [Cha10], see also earlier work of Agarwal and Matoušek [AM95].

A starting ray of L can be computed by computing the intersection of the median levels in the vertical plane $\Pi_d : y = d$ for a small enough d, defined as in the proof of Lemma 8.2.1, in linear time, using an algorithm of Megiddo [Meg85].

Consider a point p on the initial edge of L (infinite in the $-y$-direction). It lies on one plane of $\pi \in B$ and one plane $\pi' \in R$. Let ℓ be the intersection line of π and π', and consider the half line ρ of ℓ starting at p and infinite in the $+y$-direction. The planes of $B \cup R$ (besides π and π') are classified into those lying above p and those lying below it. Call the first set U and the second D. We preprocess the intersection of the set of lower half-spaces defined by the planes of U and the intersection of the set of upper half-spaces defined by the planes of D for dynamic ray shooting and shoot with ρ. The earlier of the two intersections identifies the first plane π_2 of $(B \cup R) \setminus \{\pi, \pi'\}$ that ρ meets. This is the next vertex of $\mathcal{A}(B \cup R)$ on L; L turns here. If π_2 belongs to B, L now follows the intersection line of π_2 and π'. Otherwise it follows the intersection line of π and π_2. Past the intersection, the sets U and D need to be updated and we continue, following the next ray, until we trace all of L.

The only cost besides the initial computation of p are identifying U and D in $O(n)$ time, initializing the dynamic structure, in $O^*(n)$ time, and performing two ray shots and $O(1)$ updates on U and D per vertex of L, each at a cost of $O^*(1)$. □

Note. As we construct L, we can store it as a sequence of vertices and edges. Each edge is associated with the red-blue pair of planes containing it. An endpoint of an edge is contained in an additional plane. For each consecutive edge/vertex pair (e, v), in either direction, we record which new plane contains v together with its color.

We now return to the computation of the eight-partition (H_1, H_2, H_3). By the general position assumption, H_2 and H_3 cannot be vertical, so H_2 and H_3 correspond to vertices in $\mathcal{A}$, by Observation 4. With the above setup, we can reformulate the problem of computing H_2 and H_3 as follows.

Claim 1 (The dual alternating sign functions). Computing H_2 and H_3 is equivalent to identifying a pair of vertices $p, q \in L$ such that $Y(p, q) = X(p, q) = 0$.

Proof. By Observation 4((ii)), the eight-partition (H_1, H_2, H_3) has exactly k points in each of the eight open octants. Setting $p := H_2^*$ and $q := H_3^*$, we obtain that $|R_p^\pm \cap R_q^\pm| = |B_p^\pm \cap B_q^\pm| = k$ for all combinations of signs. Therefore $Y(p, q) = X(p, q) = 0$, as claimed.

We now argue the other direction. Let $p, q \in L$ be vertices such that $X(p, q) = Y(p, q) = 0$. Since p and q lie on L, $H_2 := p^*$ and $H_3 := q^*$ bisect both R and B and contain exactly three points each, at least one of each color. Hence, it suffices to show that (H_2, H_3) is a four-partition of both R and B, i.e., $|R_p^\pm \cap R_q^\pm|, |B_p^\pm \cap B_q^\pm| \le k$ for all combinations of signs. Indeed, this implies that each octant formed by (H_1, H_2, H_3) contains exactly k points, completing the proof.

Let $a_r := |R_p^+ \cap R_q^+|$, $b_r := |R_p^+ \cap R_q^-|$, $c_r := |R_p^- \cap R_q^+|$, and $d_r := |R_p^- \cap R_q^-|$. Define a_b, b_b, c_b, d_b analogously for the blue planes. Without loss of generality, for a contradiction, suppose $a_r > k$.

We first consider the case $a_r \ge k + 2$. Since p lies on the median level of $\mathcal{A}(R)$, we have $a_r + b_r \le |R_p^+| = 2k + 1$, implying $b_r \le k - 1$. Similarly, since q lies on the median level of $\mathcal{A}(R)$, we have $c_r \le k - 1$. Recall that, by assumption, $X(p, q) = a_r + d_r - b_r - c_r = 0$, implying $d_r = b_r + c_r - a_r \le k - 4$. Hence, $a_r + b_r + c_r + d_r \le 4k - 4$, so p and q together are contained in $4k + 3 - (a_r + b_r + c_r + d_r) \ge 7$ red planes, contradicting the general position assumption.

We may now assume $a_r = k + 1$. Following the same reasoning we obtain $b_r \le k$, $c_r \le k$, and $d_r = b_r + c_r - a_r \le k - 1$. This implies $a_r + b_r + c_r + d_r \le 4k$, and, in particular, that p and q together are contained in at least 3 red planes. Now consider the blue planes and note that $a_b + b_b + c_b + d_b \le 4k$ — this is clear if each of sets $B_p^\pm \cap B_q^\pm$ contains at most k blue planes, otherwise it follows by the same argument as above. Hence, p and q together are contained in $4k + 3 - (a_b + b_b + c_b + d_b) \ge 3$ blue planes.

By Observation 4((ii)), p and q are contained in at most 6 planes of $R^* \cup B^*$. Combined with the argument above, this implies p and q together are contained in exactly 3 planes of each color. It follows that $a_r + b_r + c_r + d_r = a_b + b_b + c_b + d_b = 4k$, which, by the assumption $a_r = k + 1$, implies $b_r = c_r = k$ and $d_r = k - 1$. Since $|R_p^-| = 2k + 1$ and $b_r + d_r = 2k - 1$, there are exactly 2 red planes containing q below p. Similarly, since $|R_q^-| = 2k + 1$ and $b_r + d_r = 2k - 1$, there are exactly 2 red planes containing p below q. But then p and q are contained in a total of 4 red planes, a contradiction.

This exhausts all possibilities and, hence, $|R_p^\pm \cap R_q^\pm|, |B_p^\pm \cap B_q^\pm| \le k$ for all combinations of signs, completing the proof. □

To summarize, once we construct L in time $O^*(n+m)$, to compute an eight-partition, it is sufficient, by Claim 1, to find two vertices $p, q \in L$ satisfying $X(p,q) = Y(p,q) = 0$. In particular, it is possible to construct an eight-partition by enumerating all the $\Theta(m^2)$ pairs of vertices in L; the exact running time depends on how efficiently one can check candidate pairs. Below, we describe how to reduce the amount of remaining work to $O((m+n)\log m)$.

Speed up For simplicity of later calculations, we orient L in the y-direction (which is possible by Lemma 8.2.1) and view it as an alternating sequence of edges and vertices, starting and ending with a half-line. We denote these elements by $x_1, x_2, \dots, x_m$. Recall that the goal is to identify $i, j \in [m]$ so that x_i, x_j are vertices and $X(x_i, x_j) = Y(x_i, x_j) = 0$.

We extend the definition of X, Y as follows. If x_i, x_j are both edges, we pick arbitrary points p and q in the open edges x_i and x_j, respectively, and set $X(x_i, x_j) := X(p,q)$ and $Y(x_i, x_j) := Y(p,q)$; the cases where x_i is an edge or x_j is an edge, but not both, are handled analogously. Note that specifying the (open) edges containing p and q is sufficient to determine X and Y, hence the definition is unambiguous. Define $\pi\colon [m]^2 \to \mathbb{Z}^2$ by

$$\pi(i,j) := (X(x_i, x_j), Y(x_i, x_j)).$$

With this setup, our goal is to identify a point $(i,j) \in [m]^2$ (corresponding to a pair of vertices on L) such that $\pi(i,j) = \mathbf{0}$.

We define a *grid curve* C to be a sequence of points $(i_1, j_1), \dots, (i_t, j_t)$ in $\mathbb{Z}^2$ such that $(i_{\ell+1}, j_{\ell+1}) \in \{(i_\ell, j_\ell), (i_{\ell\pm1}, j_\ell), (i_\ell, j_{\ell\pm1})\}$ for each $\ell \in [t-1]$. In words, a grid curve is a walk in $\mathbb{Z}^2$ which, at each step, does not move at all or moves by exactly one unit up/down/left/right. A curve is *closed* if $(i_1, j_1) = (i_t, j_t)$. A grid curve is *simple* if non-consecutive points are distinct (we think of the start and end points as consecutive) — so the curve does not revisit a point after it moves away from the point.

To each grid curve C, we associate a piecewise linear curve $\overline{C}$ in $\mathbb{R}^2$, consisting of line segments connecting consecutive points $(i_\ell, j_\ell), (i_{\ell+1}, j_{\ell+1})$ of C for each $\ell \in [t-1]$. For a curve $\overline{C}$ not passing through the origin $\mathbf{0}$, the winding number $w(\overline{C})$ about $\mathbf{0}$ is defined in the standard way. Slightly abusing notation, we set $w(C) := w(\overline{C})$. In particular, provided $\overline{C}$ misses the origin,

$$w(C) = w(\overline{C}) = \begin{cases} 0 & \text{if } \overline{C} \text{ does not wind around } \mathbf{0}, \\ n > 0 & \text{if } \overline{C} \text{ winds around } \mathbf{0}\ n \text{ times counterclockwise}, \\ n < 0 & \text{if } \overline{C} \text{ winds around } \mathbf{0}\ {-n} \text{ times clockwise}. \end{cases}$$

We omit the rigorous definition of $w(C)$ as a contour integral in the complex plane since it does not add to the discussion and, instead, refer the reader to [Kra99, Chapter 4.4.4].

Our algorithm proceeds as follows:

Step 1 Set $C := T$ (see Definition 8.2.1). If $\pi(C)$ meets $\mathbf{0}$, then stop — we have found a point that maps to $\mathbf{0}$ (see Lemma 8.2.4). Otherwise $\pi(\overline{C})$ has odd winding number, by Lemma 8.2.5.

Step 2 Construct two simple closed curves C_1, C_2 so that (a) $\overline{C} = \overline{C}_1 + \overline{C}_2$, (b) at least one of $\pi(C_1), \pi(C_2)$ has odd winding number (unless they meet $\mathbf{0}$), (c) the regions enclosed by $\overline{C}_1$ and $\overline{C}_2$ partition the region enclosed by $\overline{C}$, and (d) the area enclosed by each of $\overline{C}_1, \overline{C}_2$ is a fraction of that enclosed by $\overline{C}$ (see Lemma 8.2.7).

Step 3 If $\pi(C_1)$ or $\pi(C_2)$ meets $\mathbf{0}$, then stop — we found a point that maps to $\mathbf{0}$, by Lemma 8.2.4.

Step 4 Compute $w(\pi(C_1))$ and $w(\pi(C_2))$, and replace C with the one with the odd winding number. Goto Step Step 2.

We now proceed to fill in the details, starting with the definition of the initial curve T.

Definition 8.2.1 (The triangular grid curve T)**.** The simple closed grid curve T traverses a *triangular* path defined as follows:

- Starting with the bottom horizontal side of the grid $[m]^2$, T traverses the points
$$(x_1, x_1), (x_2, x_1), \ldots, (x_m, x_1),$$
- continuing along the right vertical side of the grid $[m]^2$ along the points
$$(x_m, x_1), (x_m, x_2), \ldots, (x_m, x_m),$$
- finally, traversing back *diagonally* along
$$(x_m, x_m), (x_{m-1}, x_m), (x_{m-1}, x_{m-1}), (x_{m-2}, x_{m-1}), \ldots, (x_1, x_2), (x_1, x_1).$$

Along the diagonal side of T, we are really only interested in points of the form (x_ℓ, x_ℓ) with $\ell \in [m]$. However, since this doesn't give a grid curve, we "patch" it up by introducing intermediate points. Fortunately, this does not change the desired properties of T.

Lemma 8.2.3. *If C is a grid curve, then $\pi(C)$ is a grid curve. Moreover, if L has already been computed, $\pi(C)$ can be computed in time $O(n + |C|)$.*

Proof. Consider a step in C from (x_i, x_j) to (x_{i+1}, x_j), where x_i is an edge of L and x_{i+1} is a vertex. Then x_{i+1} is contained in the planes that contain x_i and one additional plane H. Suppose, without loss of generality, that H is red. This means that the cardinality of one of the sets $R_p^\pm$ changes by one. Hence, the cardinality of $R_p^\pm \cap R_q^\pm$, for each combination of signs, changes by at most one — if H contains q, nothing changes. It follows that the function X changes by at most one, and the function Y remains unchanged.

Note that, up to symmetry, only one such transition or its reverse occurs in a single step of C. We've shown that each step causes either X or Y (but not both) to change by at most one, and, hence, $\pi(C)$ is a grid curve.

The computation can be carried out in constant time per incident edge-vertex pair of C, since L has been already computed, after a $O(n)$-time initialization that computes X, Y at an arbitrary starting point of C by brute force. □

Lemma 8.2.3 immediately implies the following.

Lemma 8.2.4. *If $\overline{\pi(C)}$ meets $\mathbf{0}$, then some point of C is mapped to $\mathbf{0}$.*

A key property of the triangular grid curve T is the following.

Lemma 8.2.5. *If $\mathbf{0} \notin \pi(T)$, then $w(T)$ is odd.*

Proof. Let $N := 4k + 2$, and let H, V, D be the images (under π) of the horizontal, vertical, diagonal sides of T, respectively. Note that $\pi(T)$ is the concatenation of H, V, and D in that order.

Observe that if $p = q = x_i$ with $i \in [m]$, then $|R_p^+ \cap R_q^-| = |R_p^- \cap R_q^+| = 0$. Hence, $X(x_i, x_i) \in \{4k + 1, 4k + 2\}$ depending on whether x_i is contained in one or two red planes. Similarly, $Y(x_i, x_i) \in \{4k + 1, 4k + 2\}$. Hence $\pi(x_i, x_i) \in \{(N, N), (N - 1, N - 1)\}$ and, in particular, $\pi(x_i, x_i) = (N, N)$ if x_i is an edge. Along with Lemma 8.2.3, this implies that the grid curve D is a closed walk on the points in $\{N - 2, N - 1, N, N + 1\}^2 \setminus \{\mathbf{0}\}$ starting and ending at the point (N, N).

Noting that x_1 and x_m are half-lines contained in the same red plane, and that every red plane that lies above x_1 lies below x_m and vice versa, we obtain $\pi(x_m, x_1) = (-N, -N)$. Hence, H is a grid curve from the point (N, N) to $(-N, -N)$ and V is a grid curve from the point $(-N, -N)$ to (N, N).

The discussion above implies that $w(T)$ is equal to the winding number of the concatenation of V and H. We argue below that V is the image of H under a rotation by 180° around the origin, i.e., the map $(x, y) \mapsto (-x, -y)$. Since, by assumption, neither H nor V contain $\mathbf{0}$, the concatenation of H and V has odd winding number as claimed.

Specifically, we need to show that $\pi(x_i, x_1) = -\pi(x_m, x_i)$. Since π is symmetric in the two arguments, it suffices to show that $\pi(x_1, x_i) = -\pi(x_m, x_i)$. As mentioned before, every plane that lies above x_1 lies below x_m and vice versa. That is, $R_{x_1}^+ = R_{x_m}^-$ and $R_{x_1}^- = R_{x_m}^+$, and similarly $B_{x_1}^+ = B_{x_m}^-$ and $B_{x_1}^- = B_{x_m}^+$. The claim is now obvious from the definition of the functions X and Y. □

Lemma 8.2.6. *If $w(\pi(C))$ is odd, then there is a point $(i, j) \in \mathbb{Z}^2$ enclosed by $\overline{C}$ with $\pi(i, j) = \mathbf{0}$.*

Proof. A *grid square* S is a simple closed grid curve of the form

$$(i, j), (i + 1, j), (i + 1, j + 1), (i, j + 1), (i, j)$$

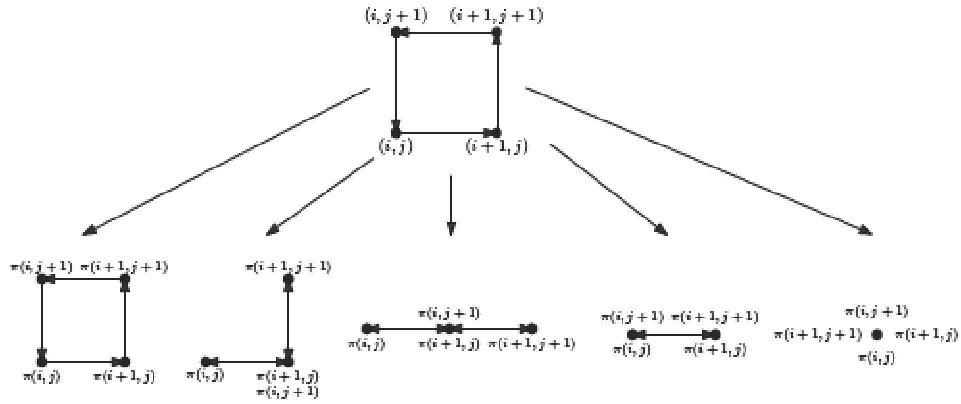

Figure 8.1: Up to symmetries, the different possibilities for the image under π of a grid square S, which is always always a grid curve in $\mathbb{Z}^2$, by Lemma 8.2.3. Note that the image cannot have odd winding number.

with $(i, j) \in \mathbb{Z}^2$. A *square* is $\overline{S}$ for some grid square S. If there is a grid square S enclosed by $\overline{C}$ such that $\pi(S)$ meets $\mathbf{0}$, then we are done by Lemma 8.2.4. Otherwise, note that $\overline{\pi(C)}$ is the sum of the images of the corresponding squares. Hence, there is a grid square S with $w(\pi(S))$ odd. By Lemma 8.2.3, $\pi(S)$ is a grid curve. By enumerating all possibilities (see Fig. 8.1), we conclude that $w(\pi(S))$ cannot be odd. □

Next, we show how to decompose a curve C. We restrict our attention to "trapezoidal" curves: Such a curve is the boundary of the intersection of the region bounded by the initial triangle T with a grid-aligned rectangle. This property is maintained inductively.

Lemma 8.2.7. *Given a trapezoidal curve C whose image misses $\mathbf{0}$, with $w(\pi(C))$ odd, we can construct two trapezoidal curves C_1 and C_2 so that*

(i) the region R surrounded by $\overline{C}$ is partitioned into region R_1 surrounded by $\overline{C}_1$ and region R_2 surrounded by $\overline{C}_2$.

(ii) $\text{area}(R_1), \text{area}(R_2) \leq c \cdot \text{area}(R)$, *for an absolute constant $c < 1$.*

(iii) either $\mathbf{0}$ is in the image of C_1 and C_2 or $w(\pi(C)) = w(\pi(C_1)) + w(\pi(C_2))$.

Proof. We already noted that the image of a grid square cannot have odd winding number, therefore R is not a grid square. As long as R is at least two units high, divide it by a horizontal grid chord into two near-equal-height pieces (that is, the two parts have equal height, or the lower one is one smaller) producing two regions R_1 and R_2. The curves C_1 and C_2 are the boundaries of the regions (refer to Fig. 8.2). If the height of R is one, perform a similar partition by a vertical chord into to near-equal-width pieces.

Property ((i)) is satisfied by construction. If the image of the new chord misses $\mathbf{0}$, then both C_1 and C_2 avoid $\mathbf{0}$ and property ((iii)) follows from the properties of the winding number on the plane. Finally, an easy calculation shows that, if the split height/width is even, then each part contains at most

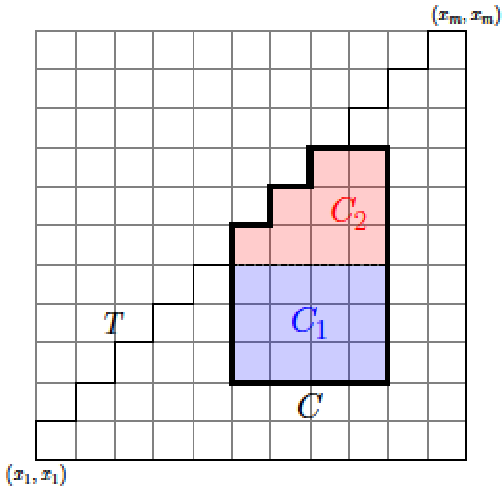

Figure 8.2: Curve C; the blue region is bounded by C_1, and the red by C_2, with the horizontal dividing chord drawn dashed.

3/4 of the original area; this fraction can rise to 5/6 if R has odd width or length (the extreme case is achieved at width/length of three), which proves property ((ii)). □

Lemma 8.2.8. *Given a simple closed grid curve C in $[m]^2$ we can determine whether $\pi(C)$ contains a zero. If not, we can compute $w(\pi(C))$, all in time $O(|C| + n)$.*

Proof. By Lemma 8.2.3, we can trace $\pi(C)$ step by step and, in particular, detect whether $\mathbf{0} \in \pi(C)$. So suppose this is not the case.

Consider the (open) ray ρ from the origin directed to the right in $\mathbb{Z}^2$. To determine the winding number of the curve $\overline{\pi(C)}$ not passing through the origin, it's sufficient to count how many times the curve crosses the ray ρ. We can compute the number of times $\overline{\pi(C)}$ crosses ρ by computing π for every vertex of C in order and counting the number of times $(X, 0)$ occurs along it, with $X > 0$.

As $\overline{\pi(C)}$ may partially overlap ρ, we need to check whether $\overline{\pi(C)}$ arrives at $(X, 0)$ with $X > 0$ from below the X-axis and (possibly after staying on the axis for a while) departs into the region above X-axis, or vice versa. That would count as a signed crossing. Arriving from below and returning below, or arriving from above and returning above, does not count as a crossing.

All of the above calculations can be done in time O(1) per step of $\pi(C)$, after proper initialization, by Lemma 8.2.3. □

Running time We now analyze the running time of the algorithm we described. We can traverse a length-$O(m)$ closed grid curve C, compute its image $\pi(C)$, and check whether it passes through the origin in time $O(m + n)$ by Lemma 8.2.3. One can check whether $\pi(C)$ winds around the origin an odd number of times by Lemma 8.2.8, also in time $O(m + n)$.

The number of rounds of the main loop of the algorithm is $O(\log m)$, as the starting curve cannot enclose an area larger than $O(m^2)$ and areas shrink by a constant factor in every iteration, by Lemma 8.2.7. Combining everything together, we conclude that L can be computed in $O^*(n + m)$ time, and the algorithm can then identify the pair of vertices of L corresponding to an eight-partition in at most $O(\log m)$ rounds, each costing at most $O(m + n)$. This concludes the proof of Theorem 8.2.1.

CHAPTER 9

Conclusions

In this dissertation we presented our work on two areas of theoretical computer science and combinatorics. In the first half we established new hardness results for approximate 4-colourings of graphs and 4-linearly ordered 3-uniform hypergraphs by combining the general algebraic theory of promise constraint satisfaction problems, combinatorial restrictions on the simplicial structure of the Hom complex construction for graphs or relational structures and the topological tools from equivariant obstruction theory.

However, we strongly believe that this is just the beginning of a much deeper and fascinating connection between topology and the theory of promise constraint satisfaction problems. One immediate approach that could lead to new results is simply to adapt different hardness criterions (different from the bounded essential arity criterion we use) from the algebraic theory to the topological context. At the same time, another promising avenue is to use more versatile construction rather than the Hom-complex to obtain a simpler functorial translation of the algebraic notions in the topological language; yet another possible direction is to use more advanced ideas from algebraic topology (e.g. spectral sequence or higher levels of Postinkov towers) to obtain the constraints on the polymorphisms required by the theory of PCSPs.

More concretely, a natural open question to explore next would be to try and establish hardness for $\mathrm{PCSP}(G, K_5)$ whenever G is 5-colourable and non-bipartite. However, trying to mimic the approached we used for the K_4-case can not work in this context as it is a simple homotopy argument to show that any two $\mathbb{Z}_2$-equivariant map from $T^2 \to S^3$ are equivariantly homotopic and thus simply considering maps up to homotopy as the target minion will cause the appearance of constant maps, invalidating the hardness criterion.

In general, it is worth investigating what kind of PCSP templates are amenable to be studied via topology. Would it be possible to give a sufficient criterion on the pair of structure so that the topological section of the argument would go through directly? One example of such criterion for the graph case is [KOWŽ23, Theorem 1.4]: if (G, H) is such that $\mathrm{Hom}(K_2, H) \to S^1$, then $\mathrm{PCSP}(G, H)$ is NP-hard. However, one interesting question would be to try rephrase this criterion in terms of graph properties of H without explicitly invoking the Hom construction.

Note that, unlike the previous mention result of [KOWŽ23], none of our hardness proofs can be easily used to give any more general criterion since Theorem 2.1.1 relies on the fact that $\mathrm{Hom}(K_2, C_\ell)$ has a very specific simplicial structure while Theorem 2.1.2 exploit the structure of the reconfigurability graph of binary polymorphisms $\mathbf{LO}_3^2 \to \mathbf{LO}_4$.

Furthermore, throughout our argument we constantly exploit the inherit $\mathbb{Z}_2$ or $\mathbb{Z}_3$ action that is given on the simplicial sets; however this is possible because the structures we are interested in are inherently symmetric. Is it possible to adapt the topological tools to work in the case of non-symmetric structure, e.g. digraphs?

As we have briefly mentioned, most of the objects and constructions we work with can be rephrased fully in general categorical terms [HJO25], however it is not clear at this point in time if such a point of view can lead to fruitful new insights in the study of the complexity of PCSPs. A natural line of research is thus to use fresh categorical insights to establish new hardness results.

The second half of this work was dedicated to the Grünbaum mass partitioning problem in 3D. In particular, we show a new existence result in three dimension by showing that any "nice" measure can be eight-partitioned by three planes, two of which intersect in a prescribed oriented direction. We also give a new algorithm for computing an eight partition of a point set in 3D where one of the plane has a prescribed normal direction.

One very tantalizing question in this area is clearly the existence of 16-partition for any nice measure in $\mathbb{R}^4$. Such question has a very rich history and it has been extensively studied; however it remains stubbornly open to this day. While from the euristic dimension counting, one would expect to have a one dimensional space of possible equipartitions, it turns out that configuration space/test map set up fails: In fact, it is possible to find equivariant maps $(S^4)^4 \to S^{15}$.

More generally, the Grünbaum-Hadwiger-Ramos problem asks for which triples (n, d, k) it is possible to simultaneously 2^k-partition n nice measure in $\mathbb{R}^d$ with k hyperplanes. By extending Avis' argument [Avi84] for the Grünbaum problem in dimension $d \geq 5$, Ramos [Ram96] showed that a necessary condition for a triple (n, d, k) to be a solution is $d \geq (\frac{2^k-1}{k})n$. He also conjectured that such condition is also sufficient:

Conjecture 9.0.1. Let n, d, k positive integers. The triple (n, d, k) is a solution for the Grünbaum-Hadwiger-Ramos problem if and only if

$$d \geq \lceil (\frac{2^k - 1}{k})n \rceil$$

At the moment the best known bound is due to Mani-Levitska, Vrećica, and Živaljević [MVZ06]:

Theorem 9.0.1. *Let n, d, k positive integers and let $a = \lfloor \log_2(n) \rfloor$. Then the triple (n, d, k) is a solution if*

$$d \geq n + (2^{k-1} - 1)2^a$$

Furthermore, even when the a solution is known to exists, we have very little information on the global topology of the solution space. Therefore a natural problem to investigate is the following:

Problem 2. Let (n, d, k) a triple that is a solution for the Grünbaum-Hadwiger-Ramos problem and fix $\mu_1, \ldots, \mu_n$ nice measures in $\mathbb{R}^d$. Define

$$S_{(\{\mu_i\},d,k)} := \{\mathcal{H} \in (S^d)^k \mid \mathcal{H}\ 2^k\text{-partitions all } \mu_i\}.$$

Determine $S_{(\{\mu_i\},d,k)}$ up to homeomorphism/homotopy/compute all the homology groups...

Finally, many different algorithmic questions are still wide open, even just in the setting of classical Grünbaum problem in 3D: Is it possible to give an algorithm to compute an 8-partition with a different requirement, e.g. two planes are orthogonal to the third or intersect in a prescribed direction?

More generally, is it possible to algorithmically compute solutions for the Grünbaum-Hadwiger-Ramos problem when such solutions are known to exists?

Another natural question is to try and speed up our algorithm: the bottleneck for the approach described in Chapter 8 is the fact that we need to explicitly compute the intersection of the two median levels of the separated points. However, a priori it might be possible to give an implicit description of the intersection curve that allows to update the alternating sums X and Y without explicitly walking along the full curve.

Bibliography

[AAH+98] Artur Andrzejak, Boris Aronov, Sariel Har-Peled, Raimund Seidel, and Emo Welzl. Results on k-sets and j-facets via continuous motion. In *Proceedings of the 14th Annual Symposium on Computational Geometry*, pages 192–199, 1998.

[ABR+24] Boris Aronov, Abdul Basit, Indu Ramesh, Gianluca Tasinato, and Uli Wagner. Eight-partitioning points in 3D, and efficiently too. In *40th International Symposium on Computational Geometry*, volume 293 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 8, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024.

[AFO+25] Sergey Avvakumov, Marek Filakovský, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. Hardness of 4-colouring *G*-colourable graphs. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, Stoc '25, pages 72–83, Prague, Czechia and New York, NY, USA, 2025. Association for Computing Machinery.

[AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad. (2+ϵ)-Sat is NP-hard. *SIAM Jounal on Computing*, 46(5):1554–1573, 2017.

[AM95] Pankaj K. Agarwal and Jirí Matousek. Dynamic half-space range reporting and its applications. *Algorithmica. An International Journal in Computer Science*, 13(4):325–345, 1995.

[Avi84] David Avis. Non-partitionable point sets. *Information Processing Letters*, 19(3):125–129, 1984.

[BBB21] Libor Barto, Diego Battistelli, and Kevin M. Berg. Symmetric promise constraint satisfaction problems: Beyond the Boolean case. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16–19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

[BBKO21] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *Journal of The Acm*, 68(4):28:1–66, 2021.

[BFH18] Pavle V. M. Blagojević, Florian Frick, and Günter M. Haase, Albert and Ziegler. Topology of the Grünbaum-Hadwiger-Ramos hyperplane mass partition problem. *Transactions of the American Mathematical Society*, 370(10):6795–6824, 2018.

[BG16] Joshua Brakensiek and Venkatesan Guruswami. New hardness results for graph and hypergraph colorings. In Ran Raz, editor, *Conference on Computational Complexity (CCC 2016)*, volume 50 of *LIPIcs*, pages 14:1–14:27, Dagstuhl, Germany, 2016. Schloss Dagstuhl–LZI.

[BG21] Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric Boolean dichotomy. *SIAM Jounal on Computing*, 50(6):1663–1700, 2021.

[Bjö95] A. Björner. Topological methods. In *Handbook of Combinatorics, Vol.1,2*, pages 1819–1872. Elsevier Sci. B. V., Amsterdam, 1995.

[BK07] Eric Babson and Dmitry N Kozlov. Proof of the Lovász conjecture. *Annals of Mathematics*, 165(3):965–1007, 2007.

[BK16] Pavle V. M. Blagojević and Roman Karasev. Partitioning a measure in $\mathbb{R}^3$. 2016.

[BKLM22] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. An invariance principle for the multi-slice, with applications. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021*, pages 228–236. IEEE Computer Soc., Los Alamitos, CA, 2022.

[BKO19] Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. In *Proc. of the 51st Annual ACM-SIGACT Symposium on the Theory of Computing, STOC 2019*, pages 602–613, New York, USA, 2019. ACM.

[Blu94] Avrim Blum. New approximation algorithms for graph coloring. *Journal of The Acm*, 41(3):470–516, 1994.

[Bre67] Glen E. Bredon. *Equivariant Cohomology Theories*, volume 34 of *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1967.

[Bul17] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, Berkeley, CA, USA, 2017. IEEE.

[Cha10] Timothy M. Chan. A dynamic data structure for 3-D convex hulls and 2-D nearest neighbor queries. *Journal of The Acm*, 57(3):16:1–16:15, 2010.

[CKK+23] Lorenzo Ciardo, Marcin Kozik, Andrei A. Krokhin, Tamio-Vesa Nakajima, and Stanislav Živný. On the complexity of the approximate hypergraph homomorphism problem, 2023.

[ČKM+13] Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner. Extending continuous maps: Polynomiality and undecidability. In *Proceedings of the 2013 ACM Symposium on Theory of Computing, STOC 2013*, pages 595–604. ACM, New York, 2013.

[CKP13] Panagiotis Cheilaris, Balázs Keszegh, and Dömötör Pálvölgyi. Unique-maximum and conflict-free coloring for hypergraphs and tree graphs. *SIAM J. Discret. Math.*, 27(4):1775–1787, 2013.

[Dey98] Tamal K Dey. Improved bounds for planar k-sets and related problems. *Discrete & Computational Geometry*, 19:373–382, 1998.

[DLGMM19] Jesús A. De Loera, Xavier Goaoc, Frédéric Meunier, and Nabil H. Mustafa. The discrete yet ubiquitous theorems of Carathéodory, Helly, Sperner, Tucker, and Tverberg. *Bulletin of the American Mathematical Society*, 56(3):415–511, 2019.

[DMR09] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *Siam Journal On Computing*, 39(3):843–873, 2009.

[DRS05] Irit Dinur, Oded Regev, and Clifford Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica. An International Journal on Combinatorics and the Theory of Computing*, 25(5):519–535, 2005.

[DS23] Anton Dochtermann and Anurag Singh. Homomorphism complexes, reconfiguration, and homotopy for directed graphs. *European Journal of Combinatorics*, 110:103704, May 2023.

[Ede86] Herbert Edelsbrunner. Edge-skeletons in arrangements with applications. *Algorithmica. An International Journal in Computer Science*, 1:93–109, 1986.

[Ede87] Herbert Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1987.

[FNO+24] Marek Filakovský, Tamio-Vesa Nakajima, Jakub Opršal, Gianluca Tasinato, and Uli Wagner. Hardness of Linearly Ordered 4-Colouring of 3-Colourable 3-Uniform Hypergraphs. *LIPIcs, Volume 289, STACS 2024*, 289:34:1–34:19, 2024.

[Fri12] Greg Friedman. Survey article: An elementary illustrated introduction to simplicial sets. *The Rocky Mountain Journal of Mathematics*, 42(2):353–423, 2012.

[GK04] Venkatesan Guruswami and Sanjeev Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discret. Math.*, 18(1):30–40, 2004.

[Grü60] Branko Grünbaum. Partitions of mass-distributions and of convex bodies by hyperplanes. *Pacific Journal of Mathematics*, 10(4):1257–1261, 1960.

[GS20] Venkatesan Guruswami and Sai Sandeep. d-to-1 hardness of coloring 3-colorable graphs with O(1) colors. In *47th International Colloquium on Automata, Languages, and Programming*, volume 168 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 62, 12. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020.

[Had66] H. Hadwiger. Simultane Vierteilung zweier Körper. *Archiv der Mathematik (Basel)*, 17:274–278, 1966.

[Har11] Sariel Har-Peled. *Geometric Approximation Algorithms*. American Mathematical Society, USA, 2011.

[Hat02] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.

[HJO25] Maximilian Hadek, Tomáš Jakl, and Jakub Opršal. A categorical perspective on constraint satisfaction: The wonderland of adjunctions, March 2025.

[HKR14] Distributed Computing Through Combinatorial Topology. In Maurice Herlihy, Dmitry Kozlov, and Sergio Rajsbaum, editors, *Distributed Computing Through Combinatorial Topology*, page iii. Morgan Kaufmann, Boston, January 2014.

[HMNŽ24] Johan Håstad, Björn Martinsson, Tamio-Vesa Nakajima, and Stanislav Živný. A logarithmic approximation of linearly-ordered colourings. In Amit Kumar and Noga Ron-Zewi, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2024*, volume 317 of *LIPIcs*, pages 7:1–7:6, August 28-30, 2024, London, UK, 2024. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[HN90] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory Series B*, 48(1):92–110, 1990.

[HS17] Dan Halperin and Micha Sharir. Arrangements. In *Handbook of Discrete and Computational Geometry*, pages 723–762. Chapman and Hall/CRC, 2017.

[Kar72] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger, editors, *Complexity of Computer Computations: Proceedings of a Symposium on the Complexity of Computer Computations, Mar 20–22, 1972, Yorktown Heights, New York*, pages 85–103, Boston, MA, 1972. Springer US.

[KLS00] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the chromatic number. *Combinatorica. An International Journal on Combinatorics and the Theory of Computing*, 20(3):393–415, 2000.

[KO19] Andrei Krokhin and Jakub Opršal. The complexity of 3-colouring H-colourable graphs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS'19)*, pages 1227–1239, Baltimore, MD, USA, 2019. IEEE.

[KO22] Andrei Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, 2022.

[KOWŽ23] Andrei A. Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM Journal on Computing*, 52(1):38–79, 2023.

[Koz08] Dmitry Kozlov. *Combinatorial Algebraic Topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, 2008.

[Kra99] Steven George Krantz. *Handbook of Complex Variables*. Springer, 1999.

[KSS84] Jeff Kahn, Michael Saks, and Dean Sturtevant. A topological approach to evasiveness. *Combinatorica*, 4(4):297–306, December 1984.

[KT17] Ken-ichi Kawarabayashi and Mikkel Thorup. Coloring 3-colorable graphs with less than $n^{1/5}$ colors. *Journal of The Acm*, 64(1):4:1–4:23, 2017.

[KTW11] Christian Knauer, Hans Raj Tiwary, and Daniel Werner. On the computational complexity of Ham-Sandwich cuts, Helly sets, and related problems. In *STACS'11*, pages 649–660, 2011.

[KTY24] Ken-ichi Kawarabayashi, Mikkel Thorup, and Hirotaka Yoneda. Better coloring of 3-colorable graphs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024*, pages 331–339, Vancouver, BC, Canada, June 24–28, 2024, 2024. ACM.

[LMS94] Chi-Yuan Lo, J. Matoušek, and W. Steiger. Algorithms for ham-sandwich cuts. *Discrete & Computational Geometry*, 11(4):433–452, April 1994.

[Lov78] László Lovász. Kneser's conjecture, chromatic number, and homotopy. *Journal of Combinatorial Theory*, 25(3):319–324, 1978.

[Mat08] Jiří Matoušek. *Using the Borsuk-Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry*. Universitext. Springer, Berlin Heidelberg, 2., corrected printing edition, 2008.

[Mat17] Takahiro Matsushita. Box complexes and homotopy theory of graphs, July 2017.

[May92] J. Peter May. *Simplicial Objects in Algebraic Topology*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992.

[MCC+96] J. Peter May, M. Cole, G. Comezana, S. Costenoble, A. Elmendorf, J. Greenlees, R. Lewis, L. and Piacenza, G. Triantafillou, and S. Waner. *Equivariant Homotopy and Cohomology Theory*. American Mathematical Society, October 1996.

[Meg85] Nimrod Megiddo. Partitioning with two lines in the plane. *Journal of Algorithms*, 6(3):430–433, 1985.

[MO25] Sebastian Meyer and Jakub Opršal. A topological proof of the Hell–Nešetřil dichotomy. In Yossi Azar and Debmalya Panigrahi, editor, *Proceedings of the 2025 Annual ACM–SIAM Symposium on Discrete Algorithms, SODA 2025*, pages 4507–4519, New Orleans, LA, USA, January 12–15, 2025, 2025. SIAM.

[MVZ06] Peter Mani-Levitska, Siniša Vrećica, and Rade Zivaljević. Topology and combinatorics of partitions of masses by hyperplanes. *Advances in Mathematics*, 207(1):266–296, 2006.

[MZ03] Jiri Matousek and Günter M. Ziegler. Topological lower bounds for the chromatic number: A hierarchy, 2003.

[NŽ23a] Tamio-Vesa Nakajima and Stanislav Živný. Boolean symmetric vs. functional PCSP dichotomy. In *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12, 2023.

[NŽ23b] Tamio-Vesa Nakajima and Stanislav Živný. Linearly ordered colourings of hypergraphs. *ACM Transactions on Computation Theory*, 14(3–4), February 2023.

[Ram96] E. A. Ramos. Equipartition of mass distributions by hyperplanes. *Discrete & Computational Geometry*, 15(2):147–167, 1996.

[RS21] Edgardo Roldán-Pensado and Pablo Soberón. A survey of mass partitions. *Bulletin of the American Mathematical Society*, 2021.

[SST01] Micha Sharir, Shakhar Smorodinsky, and Gábor Tardos. An improved bound for k-sets in three dimensions. *Discrete & Computational Geometry*, 26(2):195–204, 2001.

[tom87] Tammo tom Dieck. *Transformation Groups*. De Gruyter, Berlin, New York, 1987.

[Tót01] Géza Tóth. Point sets with many k-sets. *Discrete & Computational Geometry*, 26(2):187–194, 2001.

[Wro20] Marcin Wrochna. Homomorphism reconfiguration via homotopy. *SIAM J. Discret. Math.*, 34(1):328–350, 2020.

[WŽ20] Marcin Wrochna and Stanislav Živný. Improved hardness for H-colourings of G-colourable graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*, pages 1426–1435, Salt Lake City, UT, USA, 2020. SIAM.

[Yao88] Andrew Chi-Chih Yao. Monotone bipartite graph properties are evasive. *SIAM Journal on Computing*, 17(3):517–520, 1988.

[YDEP89] F. Frances Yao, David P. Dobkin, Herbert Edelsbrunner, and Mike Paterson. Partitioning space for range queries. *SIAM Journal on Computing*, 18(2):371–384, 1989.

[Zhu20] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *Journal of The Acm*, 67(5):30:1–30:78, 2020.

[Zuc07] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing. An Open Access Journal*, 3:103–128, 2007.

APPENDIX A

Minion homomorphisms

The goal of this appendix is to construct construct all the minion homomorphisms we use in the proof of Theorems 2.1.2 and 3.1.1 and we prove Lemma 3.1.2 (in bigger generality).

We will give here a very short presentation (based on [FNO+24]) of the results we need; however it is worth noting that all the constructions here are inherently categorical in nature and fit in a much more general framework when seen in these lenses.

We will not discuss here the broader categorical approach in details, but we refer to [HJO25] for a more thorough investigation.

A.1 Categories, functors, and natural transformations

A *category* is a collection of objects and morphism, e.g., the objects of the category of graphs are graphs, and morphisms are graph homomorphisms. Analogously, we have a category of structures in a given signature together with homomorphisms. We will also work with the category of (G-)simplicial sets with (equivariant) simplicial maps, and with a category of topological spaces. Formally, a category is defined as follows.

Definition A.1.1. A *category* Cat is a class (or a set) of *objects*, and a mapping hom which assigns to each pair of objects $A, B \in \mathrm{Cat}$ a set of *morphisms* $\hom(A, B)$. We further require that these morphisms compose in the intuitive way, and that the composition has units and is associative, i.e., we require that:

- For each $f \in \hom(A, B)$ and $g \in \hom(B, C)$, there is a morphism $g \circ f \in \hom(A, C)$.
- For all morphisms f, g, h where the composition makes sense,

$$(h \circ g) \circ f = h \circ (g \circ f).$$

- For each object $A \in \mathsf{Cat}$, there is a morphism $1_A \in \hom(A, A)$ that satisfies $f \circ 1_A = f$ for each $f \in \hom(A, B)$ and $1_A \circ f = f$ for each $f \in \hom(B, A)$.

Instead of identity between objects in a category, it is preferable to use a notion of *isomorphism* — we say that a morphism $f \in \hom(A, B)$ is an isomorphism if there exists $g \in \hom(B, A)$ such that $g \circ f = 1_A$ and $f \circ g = 1_B$. Let us fix the few categories we will work with.

Example A.1.1 (The categories of sets and finite sets)**.** The objects of the category of sets are sets, and morphisms are mappings, i.e., $\hom(A, B) = \{f \mid f : A \to B\}$. An isomorphism is a bijective mapping, i.e., two sets are considered to be isomorphic if they have the same cardinality. Sometimes, we will restrict the objects to include only finite sets, in which case, we will talk about the category of finite sets. Note that each object in the category of finite sets is isomorphic to $[n]$ for some non-negative integer n.

Example A.1.2 (The category of graphs). Graphs together with homomorphisms form a category which we denote by Graph. The categorical notion of isomorphism coincides with the usual notion of graph isomorphism.

Example A.1.3 (The category of relational structures with signature Σ). Relational structures of a fixed signature Σ together with homomorphisms form a category which we denote by Str_Σ. In particular we will be interested when the structures have one single relation of arity 3 that we will denote by Str. The categorical notion of isomorphism coincides with structural isomorphism.

Example A.1.4 (The category of equivariant simplicial sets)**.** The objects of this category are simplicial sets X endowed with a simplicial action of a fixed finite group G. The morphisms from X to Y are then simplicial maps $f : \mathsf{X} \to \mathsf{Y}$ that preserve the G-action, i.e., maps such that $f(\alpha(\sigma)) = \alpha(f(\sigma))$ for each $\alpha \in G$ and any σ simplex in X. The isomorphisms here are again such equivariant simplicial maps that are bijective on simplices in all dimensions.

Example A.1.5 (The category of topological spaces with homotopy classes of maps)**.** This is one of categories whose objects are topological spaces. The key idea is that we do not want to distinguish maps that are homotopic, and consequently, we will not distinguish homotopically equivalent spaces. As in the previous category, we will work with spaces that have an action of a fixed group G.

The objects of the category hTop_G are topological spaces X together with an action of G, and the morphisms from X to Y are then defined as the equivariant homotopy classes of equivariant maps from X to Y, i.e.,

$$\hom(X, Y) = [X, Y]_G = \{[f]_G \mid f : X \to Y, f\alpha = \alpha f \text{ for any } \alpha \in G\}$$

where $[f]_G$ denotes the class of all equivariant maps that are equivariantly homotopic to f. If it is clear from the context that we are considering equivariant classes, we will drop the index in $[f]_G$.

The composition in this category is defined as $[f] \circ [g] = [f \circ g]$. It is well-known, and not hard to prove, that the result does not depend on the choice of representatives. The identity is simply the homotopy class of the identity map. Finally, isomorphisms in this category are *equivariant homotopy equivalences*, i.e., $[f] \in [X, Y]_G$ is an isomorphism if there is an equivariant map $g \colon Y \to X$ such that $f \circ g$ is equivariantly homotopic to 1_Y and $g \circ f$ is equivariantly homotopic to 1_X.

We will work with polymorphisms in several categories — a *polymorphism* from X to Y of arity n is a morphism from the n-th power of X to Y. Therefore, in order to give a formal definition we need to define powers, or more generally products. The intuition behind the abstract definition of a products is the following observation about Cartesian products of sets: There is 1-to-1 correspondence between mappings $f \colon C \to A \times B$ and pairs of mappings $a \colon C \to A$ and $b \colon C \to B$. (Given f, we may define $a = p_A \circ f$, and $b = p_B \circ f$ where p_A and p_B denote the projection on the first and second coordinate, respectively. In the other directions, given a and b, we define f by $f(c) = (a(c), b(c))$.)

Definition A.1.2. Fix a category Cat and a positive integer n. Let $X_1, \dots, X_n \in \mathsf{Cat}$, we say that an object P together with morphisms $p_i \in \hom(P, X_i)$ for $i \in [n]$ is a *product* of $X_1, \dots, X_n$ if, for each object $C \in \mathsf{Cat}$ and morphisms $c_i \in \hom(C, X_i)$ for each i, there is a unique map $c \in \hom(C, P)$ such that $c_i = p_i \circ c$ for all $i \in [n]$. We denote the product P by $X_1 \times \dots \times X_n$, and if $X_i = X$ for all $i \in [n]$, we call P the n-th power of X_1 and write $P = X^n$.

It is straightforward to check that a product is unique up to isomorphisms, or more precisely, that if both P and C are products then the homomorphism c given by the definition above is an isomorphism.

Definition A.1.3. Assume that Cat has all finite products (i.e., the product exists for all n-tuples of objects as above), and let $A, B \in \mathsf{Cat}$. A *polymorphism* of arity n from A to B is a homomorphism from A^n to B.

A *functor* is a natural notion of a morphism between categories. A functor $\mathscr{F} \colon \mathsf{Cat}_1 \to \mathsf{Cat}_2$ is a mapping between objects and morphisms that preserves composition and identity morphisms. More specifically, we require that $\mathscr{F}(A) \in \mathsf{Cat}_2$ for each $A \in \mathsf{Cat}_1$, and $\mathscr{F}(f) \in \hom(\mathscr{F}(A), \mathscr{F}(B))$ for each $f \in \hom(A, B)$ where $A, B \in \mathsf{Cat}_1$.

Example A.1.6. 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 A.1.7. An abstract minion is a functor $\mathscr{M} \colon \mathsf{Fin} \to \mathsf{Set}$ such that $\mathscr{M}(n) \neq \emptyset$ for all $n \neq 0$.

Finally, *natural transformations* are morphisms between functors.

Formally, a natural transformation $\eta\colon \mathscr{F} \to \mathscr{G}$ where $\mathscr{F}, \mathscr{G}\colon \mathsf{Cat}_1 \to \mathsf{Cat}_2$ is a collection $\eta_A \in \hom(\mathscr{F}(A), \mathscr{G}(A))$ indexed by objects $A \in \mathsf{Cat}_1$. Such that for each $f \in \hom(A, B)$, the following square commutes:

$$\begin{array}{ccc} \mathscr{F}(A) & \xrightarrow{\mathscr{F}(f)} & \mathscr{F}(B) \\ {\scriptstyle \eta_A}\downarrow & & \downarrow{\scriptstyle \eta_B} \\ \mathscr{G}(A) & \xrightarrow[\mathscr{G}(f)]{} & \mathscr{G}(B) \end{array}$$

Example A.1.8. 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.

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

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

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

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

Definition A.1.4. Let Cat be a category with finite products, $A, B \in \mathsf{Cat}$ be two objects, and $n \geq 0$. We define a *polymorphism* from A to B of arity n to be any element $f \in \hom(A^n, B)$ 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 Cat, 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 $\mathsf{Fin} \to \mathsf{Cat}$ and $\mathsf{Cat} \to \mathsf{Set}$: 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 $A^{-}\colon \mathsf{Fin} \to \mathsf{Cat}$ for each $A \in \mathsf{Cat}$. 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 $\mathscr{A}(\pi) = p_\pi^A$ preserves compositions and identities.

The second of these assignments is known as *contravariant hom-functor* $\hom(-, B)$. It maps an object $A \in \mathsf{Cat}$ 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 5. If $\mathsf{Cat} = \mathsf{Set}$, 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 A.1.5. Let Cat be a category with finite products, and $A, B \in \mathsf{Cat}$ be such that $\hom(A, B)$ is non-empty. We define the *polymorphism minion* $\operatorname{pol}_{\mathsf{Cat}}(A, B)$ as the composition of functors A^- and $\hom(-, B)$. We will omit the index Cat whenever the category is clear from the context.

Example A.1.9. 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 $\operatorname{pol}_{\mathsf{hTop}}(X, Y) = \operatorname{hpol}(X, Y)$.

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 A.1.6. We say that a functor $\mathscr{F}\colon \mathsf{Cat}_1 \to \mathsf{Cat}_2$ *preserves finite products* if, for each $A_1, \dots, A_n \in \mathsf{Cat}_1$ 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 A.1.1. *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. ◻

A.1.1 Two general lemmata

Given the definitions above, the first of the two lemmata (Lemma A.1.2) 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 [WŽ20], although in different flavours it can be tracked down to Lawvere theories.

Lemma A.1.2. *If a functor* $\mathcal{F}\colon \mathsf{Cat}_1 \to \mathsf{Cat}_2$ *preserves products, then there is a minion homomorphism*

$$\xi\colon \operatorname{pol}(A, B) \to \operatorname{pol}(\mathcal{F}(A), \mathcal{F}(B))$$

for all $A, B \in \mathsf{Cat}_1$ *such that* $\hom(A, B) \neq \emptyset$.

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

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

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

The second lemma is a direct generalisation of [BBKO21, Lemma 4.8(1)], and the proof is analogous.

Lemma A.1.3. *Let* $A, A', B, B' \in \mathsf{Cat}$ *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*

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

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.

$$\begin{array}{ccccccc}
(A')^n & \xrightarrow{a^n} & A^n & & & & \\
\downarrow{\scriptstyle p^{A'}_\pi} & & \downarrow{\scriptstyle p^A_\pi} & \searrow^{f^\pi} & & & \\
(A')^m & \xrightarrow{a^m} & A^m & \xrightarrow[f]{} & B & \xrightarrow{b} & B'
\end{array}$$

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

A.2 Constructing the Minion Homomorphisms

The goal of this section is to construct the assignment μ and show it is indeed a minion homomorphism. As before, we will focus explicitly on the case of graphs for simplicity, but the argument is identical (up to substituting K_2 with $\mathbf{R}_3$ and $\mathbb{Z}_2$ with $\mathbb{Z}_3$) in the case of ternary relations.

A.2.1 From Graphs to Simplicial Sets

We prove that the functor $\mathrm{Hom}(K_2, -)$ (seen as a functor $\mathsf{Graph} \Rightarrow \mathsf{hTop}_{\mathbb{Z}_2}$) preserves products in the categorical sense. Here, we focus of graphs for ease of notation, but all the discussion is identical in the case of structures with one relation of arity 3 by just changing K_2 to $\mathbf{R}_3$ and $\mathbb{Z}_2$ to $\mathbb{Z}_3$.

Essentially, it follows by the same argument as the well-known fact that the product of homomorphisms complexes is homotopically equivalent to the homomorphism complex of the product (see, e.g., [Koz08, 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.

Let P, Q be posets. By definition, an n-ary *poset polymorphism* from P to Q is a map $f : P^n \to Q$ that is monotone (where the partial order on P is defined componentwise). We use the usual notation $\mathrm{pol}^{(n)}(P, Q)$ and $\mathrm{pol}(P, Q)$ and the sets of polymorphisms.

Monotone maps between posets are naturally partially ordered: $f \leq g$ if $f(x) \leq g(x)$ for all x. This allows us to relax the notion of minion homomorphism: Let $\mathscr{M}$ be a minion, and P, Q posets. A *lax minion homomorphism* $\mathscr{M} \to \mathrm{pol}(P, Q)$ is a collection of mappings $\lambda_n \colon \mathscr{M}^{(n)} \to \mathrm{pol}^{(n)}(P, Q)$ such that $\lambda_m(f^\pi) \leq \lambda_n(f)^\pi$. The following is a straightforward generalisation of [MO25, Lemma 4.1].

Lemma A.2.1. *Let G, H be graphs. There is a lax minion homomorphism*

$$\mu' \colon \mathrm{pol}(G, H) \to \mathrm{pol}(\mathrm{mhom}(K_2, G), \mathrm{mhom}(K_2, H)).$$

Proof. Let $f : G^n \to H$ be a homomorphism. We define

$$\mu'(f) \colon \mathrm{mhom}(K_2, G)^n \to \mathrm{mhom}(K_2, H)$$

by setting $\mu'(f)(m_1, \dots, m_n)$ to be the multihomomorphism

$$u \mapsto \{f(v_1, \dots, v_n) \mid v_i \in m_i(u) \text{ for } i \in [n]\}$$

where $u \in V(C)$. It is easy to check that $\mu'(f)(m_1, \dots, m_n)$ is a multihomomorphism using that all m_i's are multihomomorphisms and f is a polymorphism.

Now consider a map $\pi \colon [n] \to [k]$ and multihomomorphisms $m_j \in \mathrm{mhom}(C, G)$ for $j \in [k]$. For every vertex u of K_2, we have

$$\begin{aligned}
\mu'(f)^\pi&(m_1, \dots, m_k)(u) \\
&= \mu'(f)(m_{\pi(1)}, \dots, m_{\pi(n)})(u) \\
&= \{f(v_1, \dots, v_n) \mid v_i \in m_{\pi(i)}(u) \text{ for all } i \in [n]\} \\
&\supseteq \{f(v'_{\pi(1)}, \dots, v'_{\pi(n)}) \mid v'_j \in m_j(u) \text{ for all } j \in [k]\} \\
&= \{f^\pi(v'_1, \dots, v'_k) \mid v'_j \in m_j(u) \text{ for all } j \in [k]\} \\
&= \mu'(f^\pi)(m_1, \dots, m_k)(u)
\end{aligned}$$

Thus, $\mu'(f)^\pi \geq \mu'(f^\pi)$ as we wanted to show. Checking that $\mu'(f)$ preserves the $\mathbb{Z}_2$-symmetry is straightforward. □

By applying the monotone map $\mu'(f)$ elementwise to chains, it naturally extends to a simplicial map

$$\mu'(f)\colon \operatorname{Hom}(K_2, G)^n \to \operatorname{Hom}(K_2, H).$$

In this way, μ' can be treated as a map

$$\mu'\colon \operatorname{pol}(G, H) \to \operatorname{spol}(\operatorname{Hom}(C, G), \operatorname{Hom}(C, H)).$$

In order to show that μ' preserves minors up to homotopy, we use the following well-known result about order complexes (see, e.g., [Bjö95, Theorem 10.11] or [MO25, Lemma 2.3]):

Lemma A.2.2. *If $f, g\colon P \to Q$ monotone are monotone maps between posets such that $f \geq g$, then the induced continuous maps $|f|, |g|\colon |\Delta(P)| \to |\Delta(Q)|$ are homotopic. Moreover, if $\mathbb{Z}_2$ acts on both P and Q and f and g are equivariant, then $|f|$ and $|g|$ are equivariantly homotopic.*

Proof. Consider the poset $P \times \{0, 1\}$ with the componentwise partial order, where $\mathbb{Z}_2$ acts trivially on the first coordinate. Since $f \geq g$, the map $H\colon P \times \{0, 1\}$ defined by $H(p, 0) = f(p)$ and $H(p, 1) = g(p)$ is monotone and equivariant. Further observe that $|\Delta(P \times \{0, 1\})|$ is $\mathbb{Z}_2$-homeomorphic to $|\Delta(P)| \times [0, 1]$, hence $|H|$ induces an equivariant homotopy $|\Delta(P)| \times [0, 1] \to |\Delta(Q)|$ with $|H|(-, 0) = |f|$ and $|H|(-, 1) = |g|$. □

Using Lemma A.2.1 (applied with $G = C_\ell$ and $H = K_4$), Lemma A.2.2, and the equivariant simplicial map $s\colon \operatorname{Hom}(K_2, K_4) \to \Sigma^2$ described in Lemma 2.3.1, we get the required homomorphism μ:

Lemma A.2.3. *There is a mapping $\mu\colon \operatorname{pol}(C_\ell, K_4) \to \operatorname{spol}(\Gamma_{4\ell}, \Sigma^2)$ such that $|\mu(f^\pi)|$ and $|\mu(f)^\pi|$ are $\mathbb{Z}_2$-homotopic for all polymorphisms $f \in \operatorname{pol}^{(n)}(C_\ell, K_4)$ and $\pi\colon [n] \to [m]$.*

Proof. Let $s\colon \operatorname{Hom}(K_2, K_4) \to \Sigma^2$ be the equivariant simplicial map described in Lemma 2.3.1. Then μ is defined by $\mu(f) = s \circ \mu'(f)$. We have $\mu'(f^\pi) \leq \mu'(f)^\pi$, hence the geometric realisations of these two maps are equivariantly homotopic by Lemma A.2.2. Composing with s preserves both minors and equivariant homotopies. □

A.2.2 From simplicial sets to topological spaces

Here, we use the fact that geometric realisation preserves finite products [Fri12, Theorem 5.12],[1] and hence, for any simplicial map $f\colon \mathsf{X}^n \to \mathsf{Y}$, we can treat $|f|$ as a function $|\mathsf{X}^n| \to |\mathsf{Y}|$. The following is then an instance of the more general Lemma A.1.2.

[1] In our case, the simplicial sets are locally finite, hence the statement is true for the usual product of topological spaces.

Lemma A.2.4. *Let* X *and* Y *be two simplicial sets, then the mapping*

$$\eta'\colon \operatorname{spol}(\mathsf{X},\mathsf{Y}) \to \operatorname{hpol}(|\mathsf{X}|,|\mathsf{Y}|)$$

defined by $\eta'(f) = [|f|]$ *is a minion homomorphism.*

Proof sketch. After we have identified $|\mathsf{X}^n|$ with $|\mathsf{X}|^n$, there is not much happening here. The functions $|f|^\pi$ and $|f^\pi|$ agree on vertices since on those they are both defined as f^π. Similarly, they map each of the faces to the same face. Finally, on internal points of the faces, they are both defined as a linear extension of f^π, and hence they are equal, and consequently, their homotopy classes coincide. □

Combining the above with the relaxation lemma using the $\mathbb{Z}_2$-equivariant continuous map $S^2 \to Y$ constructed in Lemma 3.1.2 below, we obtain the desired minion homomorphism.

Lemma A.2.5. *There is a minion homomorphism*

$$\eta\colon \operatorname{spol}(\Gamma_{4\ell},\Sigma^2) \to \operatorname{hpol}(S^1, Y).$$

A.2.3 From graphs to topological spaces

Finally, let us discuss the composition $\phi = \eta \circ \mu$ defined in Lemmas A.2.3 and A.2.5. The claim is the following.

Lemma A.2.6. *The composition* $\phi = \eta \circ \mu$ *is a minion homomorphism*

$$\phi\colon \operatorname{pol}(C_\ell, K_4) \to \operatorname{hpol}(S^1, Y).$$

Proof. It is enough to show that the composition preserves minors. For that, let $f \in \operatorname{pol}^{(n)}(C_\ell, K_4)$ and $\pi\colon [n] \to [m]$. We have that $\mu(f)^\pi$ and $\mu(f^\pi)$ are $\mathbb{Z}_2$-homotopic by Lemma A.2.3. Furthermore,

$$\phi(f)^\pi = (\eta\mu(f))^\pi = \eta(\mu(f)^\pi) = \eta\mu(f^\pi) = \phi(f^\pi)$$

where the third equality uses the fact that η is constant on homotopy classes (which is true since η is a composition of η' and postcomposition with $S^2 \to Y$, and η' is constant on homotopy classes by definition). □

This concludes the proof of Lemma 3.1.2.

APPENDIX B

Auxiliary Lemmas

B.1 Mass Partitioning

B.1.1 Limit arguments

We prove some standard facts using limit arguments.

Lemma B.1.1 (Limit argument for finite point sets). *Let $X \subseteq \left(S^3\right)^3$ be a compact subset such that, for all μ mass distributions (with connected support) on $\mathbb{R}^3$ there is a plane configuration $\mathcal{H} \in X$ that eight-partitions μ; then for any set P of points in $\mathbb{R}^3$, there is a configuration $\mathcal{H}_\infty \in X$ that eight-partitions the point set.*

Proof. Let n be the number of points in P. Let μ_i be the measure defined as follows. At every point in P, place a ball of radius $\epsilon_i = \frac{1}{2^i}$ with uniform density and total mass $(1 - \epsilon_i)/n$; finally, add a normal Gaussian distribution "on the background," with total mass ϵ_i/n. Note that the total measure μ_i of the complement of the union of balls is less than ϵ_i.

By choosing i large enough, we can assume that a plane can intersect a collection of balls only if their centres are coplanar; hence, without loss of generality, we can assume that this happens for $i = 1$.

By assumption, there is a plane configuration $\mathcal{H}_i$ that eight-partitions the mass μ_i for each i; by compactness of X, there is a subsequence $\mathcal{H}_{i_j}$ that converges to some limit $\mathcal{H}_\infty$; up to reindexing we can assume that the original sequence $\mathcal{H}_i$ does. The obtained limit point eight-partitions the original point set P: in fact, there is a i_0 big enough such that for every orthant $\alpha \in \mathbb{Z}_2^3$ and any $m \geq i_0$, every point $p \in P \cap O_\alpha^{\mathcal{H}_\infty}$ is "far away" (e.g., at least $1/2^{i_0}$) from the planes in the configuration $\mathcal{H}_m$; hence

$$\frac{1 - \epsilon_m}{n} |P \cap O_\alpha^{\mathcal{H}_\infty}| \leq \mu_m(O_\alpha^{\mathcal{H}_m}) = \frac{1}{8}.$$

Taking the limit in m we obtain the desired result. □

Corollary B.1.1. *Let X and P as above. If $\mathcal{H}_\infty = (H_1, H_2, H_3)$ is the configuration constructed in the proof of Lemma B.1.1, then any plane H_i in $\mathcal{H}_\infty$ bisects P and any pair (H_i, H_j) four-partitions the points.*

Proof. For simplicity we show the result for the first plane H_1 and the pair (H_1, H_2), all the other cases are identical. First, construct μ_i and $\mathcal{H}_i = (H_{i,1}, H_{i,2}, H_{i,3})$ converging to the limit $\mathcal{H}_\infty$ as in the proof of Lemma B.1.1.

Again, by choosing i_0 big enough we obtain that, for any $m \geq i_0$ and any sign $\alpha \in \mathbb{Z}_2$ every point $p \in P \cap H_1^\alpha$ is sufficiently far form $H_{m,1}$ hence

$$\frac{1-\epsilon_m}{n}|P \cap H_1^\alpha| \leq \mu_m(H_{m,1}^\alpha) = \frac{1}{2}.$$

Similarly, for any pair of signs $(\alpha_1, \alpha_2) \in \mathbb{Z}_2^2$, any point $p \in P \cap H_1^{\alpha_1} \cap H_2^{\alpha_2}$ is sufficiently far from both $H_{m,1}$ and $H_{m,2}$, therefore

$$\frac{1-\epsilon_m}{n}|P \cap H_1^{\alpha_1} \cap H_2^{\alpha_2}| \leq \mu_m(H_{m,1}^{\alpha_1} \cap H_{m,2}^{\alpha_2}) = \frac{1}{4}.$$

By taking the limit we obtain the desired result. □

Lemma B.1.2 (Limit argument for mass distributions with possibly disconnected support). *Let X be a compact set in $(S^3)^3$ such that, for any mass distribution with connected support there is a configuration $\mathcal{H} \in X$ that eight-partitions the measure. Let μ be a "general" mass distribution. Then there is a plane arrangement $\mathcal{H}_\infty \in X$ that eight-partitions μ.*

Proof. Define $\epsilon_i := \frac{1}{2^i}$ and let μ_i the measure defined, on a measurable set $A \subseteq \mathbb{R}^3$, as

$$\mu_i(A) := (1-\epsilon_i)\,\mu(A) + \epsilon_i \mathcal{N}(A),$$

where $\mathcal{N}$ is a normal Gaussian distribution on $\mathbb{R}^3$. Then, μ_i is a mass distribution with connected support hence there is a configuration $\mathcal{H}_i$ that eight-partitions μ_i. By compactness, up to taking a subsequence, $\mathcal{H}_i$ converges to a configuration $\mathcal{H}_\infty$.

Now, for any $\alpha \in \mathbb{Z}_2^3$, we have that

$$(1-\epsilon_i)\,\mu(O_\alpha^{\mathcal{H}_i}) \leq \mu_i(O_\alpha^{\mathcal{H}_i}) = \frac{1}{8}.$$

For any fixed measure $\tilde{\mu}$, the map $\mathcal{H} \to \tilde{\mu}(O_\alpha^{\mathcal{H}})$ is continuous; hence by taking the limit in i we obtain that, for all $\alpha \in \mathbb{Z}_2^3$,

$$\mu(O_\alpha^{\mathcal{H}_\infty}) \leq \frac{1}{8}.$$

However, since

$$\sum_{\alpha \in \mathbb{Z}_2^3} \mu(O_\alpha^{\mathcal{H}_\infty}) = \mu(\mathbb{R}^3) = 1,$$

it follows that all the previous inequalities are equalities. □

B.1.2 Four-partitioning in the plane with a prescribed bisector

This section is devoted to the proof of Lemma 7.1.1. For convenience, we restate the lemma here. Both the lemma and proof are due to [BK16].

Lemma 7.1.1 (Four-partitioning a mass distribution in $\mathbb{R}^2$ [BK16]). *Let $\mu^\#$ be a mass distribution (with connected support) on $\mathbb{R}^2$ and $v \in S^1$. Then there exists a pair (ℓ_1, ℓ_2) of lines in $\mathbb{R}^2$ that four-partitions $\mu^\#$ and such that v bisects the angle between ℓ_1 and ℓ_2.*

Moreover, if we orient ℓ_1 and ℓ_2 so that ℓ_1 is in the first direction clockwise from v, and ℓ_2 is in the first direction counterclockwise, the oriented pair is unique and the lines depend continuously on v.

Proof. Suppose, without loss of generality, that $v = (0, 1)$. We first prove existence. Let $\alpha \in [0, \pi/2]$, and rotate v counterclockwise and clockwise by angle α to obtain u_α and w_α, respectively. Then, since the measure has connected support, there exist unique lines ℓ_α and m_α that are perpendicular to u_α and w_α, respectively, and bisect μ. Note that v bisects the angle between ℓ_α and m_α.

The (oriented) lines ℓ_α and m_α partition the plane into four octants, which we label P_N, P_E, P_S, P_W (north, east, south, west) in the obvious manner. Since both lines are bisecting, we have

$$\mu(P_N) = \mu(P_S) = x, \quad \mu(P_W) = \mu(P_E) = \mu(\mathbb{R}^2)/2 - x = y.$$

When α tends to 0, P_W and P_E tend to empty sets and evidently $x > y$ for α sufficiently close to 0. When α tends to $\pi/2$, P_N and P_S tend to empty sets and then $x < y$ for α sufficiently close to $\pi/2$. Since x depends continuously on α, we must have $x = y$ somewhere in between, by the intermediate value theorem. Thus, we have existence.

We now show uniqueness. Assume we have a partition P_N, P_E, P_S, P_W with angle α and another partition Q_N, Q_E, Q_S, Q_W with angle α'. Assume without loss of generality that $\alpha' \leq \alpha$. Moreover, since for $\alpha' = \alpha$ the partition is defined uniquely, we may assume $\alpha' < \alpha$. Let $p = \ell_\alpha \cap m_\alpha$ and consider the following cases:

1. $p \in Q_N$: In this case $P_N \subset Q_N$ and both sets have the same measure. This contradicts connectivity of the set where the density is positive, since the density is positive in Q_S and in P_N, it must be positive somewhere in $Q_N \setminus P_N$, implying $\mu(Q_N) > \mu(P_N)$.

2. $p \in Q_E$: In this case $P_W \subset Q_W$ and we obtain a similar contradiction.

3. $p \in Q_S$ and $p \in Q_W$ are similar to considered cases.

Since the lines $\ell_\alpha, m_\alpha, \ell_{\alpha'}$, and $m_{\alpha'}$ are distinct, this covers all cases. In each case, we obtain a contradiction, hence, we have uniqueness.

As for continuity, it follows from the standard fact that a map with compact codomain and closed graph must in fact be continuous. The codomain is compact since we are only interested in directions of the halving lines that afterwards produce halving lines continuously, thus working with $S^1 \times S^1$ as the space of parameters. □