GRAPHS IN HIGHER DIMENSIONS

by

TSVETOMIR MIHAYLOV

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

School of Mathematics
College of Engineering and Physical Sciences
The University of Birmingham
February 10, 2024

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

The overarching theme of this thesis is exploring exciting results in graph theory and discovering suitable ways to extend them to exciting results in higher dimensions.

The first subtopic of the thesis is extending two-dimensional results about graphs into three dimensions via 2-complexes. In Chapter 3, we extend the outerplanarity criterion through forbidden minors of Chartrand and Harary [20], in Chapter 5, we extend the Four Colour Theorem [4] and in Chapter 6, we extend Mac Lane's Theorem [49].

The second subtopic of the thesis is annulus graphs, which serve as extension of both unit distance graphs and unit disc graphs in higher dimensions. The paper in Chapter 4 concerns annulus graphs in the d-dimensional spaces $\mathbb{R}^d$ for $d > 0$. There, we introduce the relatively unstudied term annulus graph and show that they have a non-trivial structure. We also show that the ratio of the chromatic number to the independence number grows exponentially in the dimension of the underlying $\mathbb{R}$-space.

Contents of the thesis

Chapter 3 is based on the paper "Outerspatial 2-complexes: Extending the class of outerplanar graphs to three dimensions" by J. Carmesin and T. Mihaylov. It is published in The Electronic Journal of Combinatorics. The ideas in this paper were contributed to equally by both authors. The writing has been done mainly by T. Mihaylov with the help of J. Carmesin.

Chapter 4 is based on the paper "Annulus graphs in $\mathbb{R}^d$" by L. Lichev and T. Mihaylov. It is submitted for minor revisions to the Journal of Discreet and Computational Geometry. The ideas in this paper were contributed to equally by both authors. The writing has been shared equally by both authors. The last section of this chapter shows some additional works done by T. Mihaylov with the help of A. Jung and is written by T. Mihaylov.

Chapter 5 is based on a paper by J. Carmesin, J. Kurkofka, T. Mihaylov, E. Nevinson, which is in preparation for submission. The ideas in this paper were contributed to equally by all authors. As presented in the thesis, the paper has been mostly written by T. Mihaylov except for Section 5.5 which was mainly written by E. Nevinson.

Chapter 6 shows initial work related to the paper "Dual matroids of 2-complexes – revisited" by J. Carmesin. The ideas in this chapter were contributed to equally by T. Mihaylov and J. Carmesin. The writing has been done mainly by T. Mihaylov.

Chapter 7 shows initial work related to the material in Chapter 3. The ideas in this chapter were contributed to equally by T. Mihaylov and J. Carmesin. The writing has been done mainly by T. Mihaylov.

Acknowledgements

I am incredibly grateful to my supervisor Johannes Carmesin for introducing me to the captivating field of 2-complexes. The three projects we undertook together have been some of the most exhilarating experiences in mathematics that I have encountered. Johannes has not only invested considerable time and effort in working with me but has also nurtured my curiosity, granting me the freedom to explore my mathematical interests. His guidance has imparted valuable mathematical, research, and life skills, empowering me to continue expanding these abilities independently.

I extend my heartfelt thanks to Johannes' research group, including Emily, Jan, and Will, for the open exchange of ideas that we had during our lunch meetings. I am especially grateful to Jan and Emily for dedicating their precious time to collaborating with me on our paper. The interactions I have had with them have profoundly shaped my identity as a mathematician and researcher.

A special mention goes to a coauthor of mine, Lyben Lichev, the collaboration with whom brought both joy and enlightenment. Our work together was immensely rewarding, and I gained invaluable insights through our correspondence. Beyond that, he has also been a great colleague who has offered me valuable advice on numerous occasions in the last few years.

I would like to thank my fellow PhD colleagues, Alistair, Amarja, Soodeh, Vincent and Yahya, for making my time at (and off) the university significantly more enjoyable. I particularly appreciate the engaging post-seminar discussions and dinners with my Combinatorics colleagues. Additionally, I am very happy to have developed a great friendship group among my fellow Bulgarians, including the four Dimitars, Ivan, Chris, and Mimi. All these friends and colleagues not only shared countless memorable moments but also gave me a sense of belonging in the city and the University.

One person from the University of Birmingham holding a special place in my heart is Eoin Long. He was my teacher during my studies in Oxford and continued to be an incredible mentor when he joined the University of Birmingham. Throughout my degree, he provided unwavering support, offering immense emotional, practical, and technical assistance. Our bond transformed into a deep and meaningful friendship.

I want to express my profound gratitude to my friends and family, including my parents, my girlfriend Tanya and my dear friend Nikola. Their steadfast love and support were instrumental in my completion of this thesis. Their belief in me and the help they provided were invaluable. I sincerely appreciate their presence in my life.

Contents

Contents **5**

1 Introduction **8**

1.1 Outline of the thesis 12

2 Basic results and definitions **14**

2.1 Graphs 14

2.2 Complexes 15

2.3 Some topological properties 18

2.4 Space minors 23

2.5 Matroids 24

3 Characterisation of outerspatial 2-complexes **29**

3.1 Introduction 29

3.2 Basic definitions and initial approaches 33

3.3 Core Lemma 42

3.4 Main techniques 43

3.4.1 Proof of Lemma 3.4.8 45

3.4.2 Proof of Lemma 3.4.9 47

3.4.3 Proof of Lemma 3.3.7 and Theorem 3.1.1. 49

3.5 Further remarks regarding outerspatial locally 2-connected 2-complexes 50

3.6 Concluding remarks 54

4 Annulus graphs in $\mathbb{R}^d$ **56**

4.1 Introduction 56

4.1.1 Our results 58

4.2 Proof of Theorem 4.1.1 60

4.3 Proof of Theorem 4.1.2 74

4.3.1 Proof of the upper bound of Theorem 4.1.2 74

4.3.2 Proof of the lower bound 76

4.4 Discussion 82

4.5 The set of annulus graphs in relation to their dimension 84

5 Chromatic number of 2-complexes **89**

5.1 Introduction 89

5.2 Preliminaries 90

5.3 Link graphs 92

5.4 Upper bound of the chromatic number of embeddable 2-complexes 96

5.5 Lower bound of the chromatic number of embeddable 2-complexes 98

5.6 Alternative proof of the lower bound of the chromatic number of embeddable 2-complexes 104

5.7 Concluding remarks 108

6 Mac Lane's embeddability criterion in three dimensional space **113**

6.1 Introduction 113

6.2 Mac Lane's embeddability criterion in three dimensions using Mac Lane's planarity criterion 114

6.2.1 Preliminaries 114

6.2.2 Proof of the 'if' implication of Theorem 6.2.3 117

6.2.3 Proof of the ‘only if’ implication of Theorem 6.2.3 and completing the proof of Theorem 6.2.2 122
6.3 Mac Lane’s embeddability criterion in three dimensions using a Whitney type theorem 124
6.3.1 Preliminaries 124
6.3.2 Proof of Theorem 6.3.2 126

7 Two additional results regarding outerspatial 2-complexes 129
7.1 Introduction 129
7.2 Forbidden minor characterisation of outerspatial 2-complexes 130
7.3 Characterisation of weakly outerspatial 2- complexes via their dual graphs 132

Bibliography 135

Chapter 1

Introduction

There are a number of exciting results about graphs concerning their geometrical properties and in particular their embeddability in surfaces. One of the earliest and most prominent examples in this line of research is Kuratowski's planarity characterisation of graphs in terms of forbidden substructures [45]. Other notable results that will be a topic of discussion in this thesis are Mac Lane's theorem [49], Whitney's theorem [69] and the four colour theorem [4]. The main focus of this thesis is exploring such results and extending them to three dimensions. A non-trivial part of doing this is finding the most appropriate formulation of a definition/result and extending it in the best way possible. The most notorious example of this is outerplanarity/outerspatiality. There are multiple equivalent ways to define outerplanarity and we used two of them as candidates to extend to three dimensions. However, we found out that they give rise to two different definitions – outerspatial and weakly outerspatial. In this thesis we predominantly study the term outerspatial as it is the stronger of the two and gives rise to more interesting structural results.

There has been some research into this topic previously but the first solid series of results in it is quite recent. This series concerns the embeddability characterisation of 2-complexes in terms of forbidden space minors by Carmesin [12, 13, 14, 15, 17]. There are previous works regarding simplicial complexes and cell complexes, however this series is the one that introduces the particular term 2-complexes in order to examine their embeddability properties. In [12], the author introduces space minors and gives a Kuratowski type characterisation of embeddable locally 3-connected 2-dimensional simplicial complexes through those space minors. In [13] rotation systems are presented in detail and then an embeddability characterisation using them as its main ingredient is given. [14] includes some graph theoretical prerequisites for [15], the latter containing two main results. One of them is a three dimensional version of Whitney's theorem and the other tells us a connection between a simplicial complex and its split complex, which enables dealing with more general types of complexes. The main task in [17] is to remove the locally 3-connectedness condition from the embeddability condition in [12]. This is done by introducing the 'stretching' operation which handles the obstructions coming from the non-locally 3-connected simplicial complexes. Many of the results in this thesis are inspired by the ideas in this series – both by using some particular results from there as a stepping stone and by following the same methodology of handling 2-complexes as three-dimensional analogues of graphs.

The first graph theoretical result that we have extended to three dimensions is Mac Lane's planarity criterion [49]. It uses the theory of cycle spaces and finds a condition on the cycle space of a graph G to determine when a graph is planar. To extend Mac Lane's theorem to three dimensions, first we need three-dimensional analogues of the terms cycle and planar. For the former, we can use the definition of 2-cycle from homology theory and for the latter we will use the notion of embeddability. A 2-complex is embeddable if there exists an injective map from its related topological space to $\mathbb{R}^3$. In Chapter 6, we present two different proofs of the same result – an embeddability characterisation of simplicial complexes depending on properties of its cycle space. The first proof relies directly on Mac Lane's planarity criterion to build a planar rotation system on the link graphs, which we can extend to a planar rotation system on the whole simplicial complex. For the second proof, we use a different criterion for the embeddability of 2-complexes – a Whitney type characterisation [15]. This characterisation is a three dimensional analogue of Whitney's theorem [69], the latter stating that a graph is planar if and only if its dual matroid is graphic. Even though these two approaches look very different at first glance, they are actually closely related. They achieve quite similar things, one through geometry and the other through algebra.

The next graph theoretical term that we want to expand on is outerplanarity – a graph is outerplanar if and only if it can be embedded in the plane in such a way that a single face covers all of the graph's vertices. The relevant result that we extended to three dimensions is the forbidden minor characterisation of outerplanar graphs given by Chartrand and Harary [20]. We inspect two equivalent definitions of outerplanar graphs as candidates to be extended to three dimensions. However, these extensions give rise to two different and non-equivalent definitions – we call one of them weakly outerspatial and the other outerspatial. As the name suggests weakly outerspatial is the weaker definition and all outerpsatial 2-complexes are weakly outerspatial but not the other way around. This is why we direct our attention mostly to outerspatiality and most of our results use this definition. In Chapter 7 we have two substantial results. One of them uses the three-dimensional Kuratowski theorem as a black box to give a list of forbidden minors of outerspatial 2-complexes. The other one gives a dual graph characterisation of weakly outerspatial 2-complexes and can be seen as an extension of the outerplanarity criterion in [29]. In Chapter 3, we use some of the ideas from Chapter 7 and improve on both of these results by giving more concise characterisations of outerspatial complexes. The main result in Chapter 3 gives a complete forbidden space minor characterisation of outerspatial 2-complexes and has a very nice corollary concerning nested plane embeddings of cycles of graphs (alternatively called laminar cycles).

Vertex colouring of planar graphs is a very old and very popular topic in graph theory. The four colour theorem [4] is probably the biggest result in this field and is the next result that we want to extend to three dimensions. The first thing we need to do, is to define what is a colouring of a 2-complex. Traditional vertex and edge colourings are not good candidates because every complete graph K_n can be embedded in an orientable surface of high enough genus. Since we know that both the chromatic number and chromatic index of a graph K_n are at least $n-1$, with such definition of a colouring, we can find a 2-complex that requires arbitrary number of colours. A meaningful way to colour 2-complexes is to colour its edges in such manner that no two edges that are adjacent in a face are of the same colour. This type of colouring is motivated by the fact that when we localise it at a vertex, it is equivalent to vertex colouring of the link graph. The main result of Chapter 5 is showing that the chromatic number of embeddable 2-complexes is 12. To prove the upper bound, we define and use the total link graph of a 2-complex, which is a generalisation of the standard link graph. Given a 2-complex C we show that the (graph) chromatic number of its total link graph is equal to the (2-complex) chromatic number of C, which reduces our problem to a graph colouring one. For the upper bound of our theorem, we show that any graph that is constructed as the total link graph of an embeddable 2-complex has a chromatic number at most 12. In proving this bound we use another result concerning colouring planar graphs, namely the m-pire problem [33]. For the lower bound, our strategy is to find an appropriate 12-chromatic graph, from which we can build an embeddable 2-complex that has this graph as its total link graph. This will result in an embeddable 2-complex with chromatic number 12. We present two proofs of this following this same strategy, one of them being more compact, the other one being more explicit in the example of a 2-complex.

The last topic that is subject of this dissertation is the topic of annulus graphs. The set of d-dimensional annulus graphs $\mathcal{A}_d(R_1, R_2)$ is defined as follows. Its vertices are a set of points in d-space and two vertices are connected by an edge if the distance between their corresponding points is in the interval (R_1, R_2) for some real numbers $0 \leq R_1 \leq R_2$. The idea behind considering such graphs is that they serve as a bridge between two well-known types of graphs – unit distance graphs and unit disc graphs. The family of d-dimensional unit distance graphs coincides with the family of annulus graphs $\mathcal{A}_d(R, R)$, for any $R > 0$. One of the most notable problems concerning unit distance graphs is the Hadwiger-Nelson problem asking for the smallest number of colours in which the points of the plane may be coloured so that no two points at a distance 1 are monochromatic. It has long been known that the answer is between 4 and 7, and de Grey [21] improved the lower bound to 5 (result reproved independently by Exoo and Ismailescu [26]). The family of d–dimensional unit disc graphs coincides with the family $\mathcal{A}_d(0, R)$ for every $R > 0$. Unit disc graphs were introduced in 1971 by Gilbert [34] to model telecommunication networks. Since then, the most significant developments of the theory of unit disc graphs were made in the framework of random unit disc graph also known as random geometric graphs, see Penrose [54] for a detailed account. The model was generalised by Waxman [68]: he worked in a setting where two vertices in positions x and y are connected with probability $\beta \exp(-|x - y|/r)$ where β and r are parameters of the model. Penrose [55] introduced and studied a percolated version of the model. In most situations, both unit distance graphs and unit disc graphs are considered in two dimensions, i.e. $d = 2$. In Chapter 4 we prove two substantial results. The first one says that the set of annulus graphs is not monotone – i.e. that increasing the interval of possible distances does not yield a superset of annulus graphs. The second one gives a bound of the ratio between the chromatic and clique numbers of annulus graphs in terms of an exponential function in the dimension of their underlying $\mathbb{R}$-space. In Section 4.5, we prove two additional results concerning the growth of the set of annulus graphs relative to their dimension.

1.1 Outline of the thesis

The thesis is organised as follows. Chapter 2 is devoted to familiarising the reader with some of the most essential definitions and previous results used in the thesis. Chapter 3 extends the notion of outerplanarity to three dimensions by introducing the term ‘outerspatial’. The main result gives a spatial minor characterisation of outerspatial 2-complexes based on the outerplanarity characterisation given in [20]. The topic of Chapter 4 is annulus graphs – graphs defined by selecting edges according to the metric distance between their endpoints. It has two main results – one of them shows that annulus graphs are chi-bounded under certain conditions and the other one proves that different classes of annulus graphs are meaningfully different. Chapter 5 extends the notion of chromatic number to 2-complexes and proves that the chromatic number of embeddable 2-complexes is 12. This can be viewed as a three-dimensional version of the four colour theorem. Chapter 6 gives two Mac Lane type characterisations of embeddability – one using a Kuratowski type embeddability result [13, Theorem 1.1] and one using a Whitney type theorem in three dimensions [15, Theorem 1.2]. Chapter 7 serves as a predecessor of Chapter 3. It gives two 2-complex characterisations – one basic characterisation of outerspatial 2-complexes which follows directly from [12, Theorem 1.1] and one characterisation of weakly outerspatial 2-complexes using dual graphs of an embedding.

Chapter 2

Basic results and definitions

The purpose of this chapter is to introduce the reader to the main topics of this thesis as well as to its most common elements. It serves as a starting point, which aims to build a general picture of the contents of this thesis and to make the rest of the material clearer and easier to understand. Most notably, the definitions of central objects such as graph and 2-complex will vary slightly in a lot of the chapters. Some of the content in this chapter comes from [12, 13, 15, 46, 52].

2.1 Graphs

As graphs are one of the main objects of this thesis, we need to take a little care with them. First of all, the graphs in all chapters apart from Chapter 4, are going to be directed multigraphs. Hence, in these parts of the thesis, when we say graph, we will mean a directed multigraph. Furthermore, we are going to need a geometric definition of a graph as we are mainly concerned with embedding graphs in spaces. The main task of this subsection is clarifying these two definitions.

Definition 2.1.1. A *directed multigraph* is a quadruple $G = (V, E, \varphi^-, \varphi^+)$, where V and E are sets and their elements are called vertices and edges respectively. The function $\varphi^- \colon E \to V$ assigns a start vertex to each edge and the function $\varphi^+ \colon E \to V$ assigns an end vertex to $\varphi^+ \colon E \to V$ each edge.

Definition 2.1.2. A (topological) path of a topological space X is the image of a continuous map $p : [0,1] \to X$ which is injective.

Definition 2.1.3. A (topological) loop of a topological space X is the image of a continuous map $\ell : [0,1] \to X$ which is injective on $[0,1)$ and $\ell(0) = \ell(1)$.

Definition 2.1.4. The *geometric realisation* of a graph G is the topological space obtained by taking a set of disjoint points – one for each vertex and for each edge we glue a topological path at its endpoints to the points corresponding to the endvertices of the respective edge (or gluing a topological loop at a vertex if the edge is a loop).

Definition 2.1.5. An embedding of a topological space X into a topological space Y is an injective continuous map from X to Y.

Example 2.1.6. A graph G is planar if there exists an embedding of its geometric realisation in $\mathbb{R}^2$.

Definition 2.1.7. Consider a graph G with a vertex set V and let $\{V_1, V_2\}$ partition V. Then the set of edges between V_1 and V_2 is denoted by $E(V_1, V_2)$ and is called a *cut*.

Definition 2.1.8. A minimal non-empty cut in a graph is called a *bond*.

Definition 2.1.9. Fix $0 \leq R_1 \leq R_2$ and $d > 0$ – an integer. Then the *d-dimensional annulus graph* G with radii R_1 and R_2 is built in the following way. Its vertex set is a finite set of points in $\mathbb{R}^d$. Two points are connected by an edge if they are at a distance at least R_1 and at most R_2.

2.2 Complexes

As we said, we are interested in extending two-dimensional results about graphs to three dimensions. To this end we need a higher dimensional counterpart of graphs. This role is served by the 2-complex. It is the other main object of this thesis along with the graph. Roughly speaking, a 2-complex is obtained by gluing discs to cycles of some graph. Below, we will define the particular kind of cycle that we need for our purposes.

Definition 2.2.1. In a graph G, an *oriented closed walk* is a finite sequence $v_0, e_0, \ldots, v_n, e_n$ of vertices v_i and edges e_i together with a family $(\sigma_i \mid i \in \mathbb{Z}_{n+1})$ of *traversals* $\sigma_i \in \{-, +\}$ such that $\varphi^{\sigma_i}(e_i) = v_{i+1} = \varphi^{-\sigma_{i+1}}(e_{i+1})$ for all i.

In the rest of this subsection, we will give an introduction of the term 2-complex in a similar manner to the way we did with graphs.

Definition 2.2.2. A 2-*complex* is a pair $C = (G, F)$, where G is a graph and F is a family of oriented closed walks in G.

One of our main concerns in this thesis is embeddability of 2-complexes. Since 2-complexes are originally defined in an abstract way, we need to also give them geometric definition in order to be able to explore their geometric properties.

Definition 2.2.3. The *geometric realisation* of a 2-complex $C = (G, F)$ is the topological space obtained from $G \bigoplus (\mathbb{D}^2 \times F)$, where F carries the discrete topology, G is the 1-complex of the graph G and $\mathbb{D}^2$ is the two-dimensional unit disc, by identifying $\partial\mathbb{D}^2 \times \{f\}$ with the closed walk f in G for every $f \in F$. The gluing map is defined as the composition $\partial\mathbb{D}^2 \times \{f\} \to \mathbb{S}^1 \to f$. The first map in the composition is a homeomorphism and the second one comes from the fact that the closed walk f in the geometric realisation of G is a topological loop.

Definition 2.2.4. A *topological disc* is a topological space homeomorphic to the tw-dimensional unit disc.

Definition 2.2.5. Consider the geometric realisation of a 2-complex $C = (G, F)$. For every face $f \in F$, we define its *interior* $\mathring{f}$ to be the image of the interior of the corresponding glued disc.

Definition 2.2.6. A 2-complex C is *embeddable* in a topological space T if there exists an embedding of the geometric realisation of C into T.

If the ambient topological space is clear from the context we will use *embeddable* 2-complex as a shorthand. In this thesis, the space will usually be $\mathbb{R}^3$.

The notion of embeddability of 2-complexes can be thought of as a three dimensional analogue of planarity for graphs. That is because a 2-complex being embeddable means that its derived topological space can be embedded in $\mathbb{R}^3$ while a graph being planar means that its derived topological space is embeddable in $\mathbb{R}^2$. It is a very important term because all of the 2-dimensional results that we are extending, are related to planarity.

Definition 2.2.7. Given a 2-complex C with a vertex set V and a subset U of V, we denote by $C \restriction U$ the 2-complex obtained from C by deleting the vertex set $V - U$ (together with all incident edges and faces).

We will mostly work with 2-complexes, but at times, we will want to restrict our attention to 2-dimensional simplicial complexes. Let us give a definition for this topological object as well.

Definition 2.2.8. The standard *n-simplex* is the set
$\Delta^n = \{(x_0, x_1, \dots, x_n) \in \mathbb{R}^{n+1} : x_i \geq 0 \ \forall i \text{ and } \sum_i x_i = 1\}$. The points $(x_0, \dots, x_n)$ where $x_i = 1$ for some i are called the *vertices* of Δ^n and are denoted by $V(\Delta^n)$. The number n is referred to as the *dimension* of the simplex.

Definition 2.2.9. For each non-empty subset A of $\{0, \dots, n\}$ there is a corresponding *face* of Δ^n which is

$$\{(x_0, \dots, x_n) \in \Delta^n : x_i = 0 \ \ \forall i \notin A\}.$$

Definition 2.2.10. A *simplicial complex* is a pair (V, Σ), where V is a set, whose elements are called vertices and Σ is a set of non-empty finite subsets of V, called simplices, such that

- for each $v \in V$, the 1-element set $\{v\}$ is in Σ;

- if σ is an element of Σ, so is any non-empty subset of σ.

We say that the simplicial complex is finite if V is a finite set.

Definition 2.2.11. The *geometric realisation* of a simplicial complex $C = (V, \Sigma)$ is the space obtained by the following procedure:

- For each $\sigma \in \Sigma$ take a copy of the standard n-simplex, where $n + 1$ is the number of elements of σ. Denote this simplex by Δ_σ. Label its vertices with the elements of σ.

- Whenever $\sigma \subset \tau \in \Sigma$, identify Δ_σ with a subset of Δ_τ, via the face inclusion which sends the elements of σ to the corresponding elements of τ.

As with 2-complexes, a simplicial complex is embeddable in some topological space T if there is an embedding of its geometric realisation into T.

Definition 2.2.12. The dimension of a simplicial complex C is the highest dimension of a simplex in C.

Remark 2.2.13. Since all simplicial complexes that we are going to use in this thesis will be 2-dimensional, if the dimension is not specified, it is assumed to be 2.

2.3 Some topological properties

Before we start, let us recall the definition of a 2-complex – Definition 2.2.2 and note the following example, which serves to demonstrate its difference with 2-dimensional simplicial complexes.

Example 2.3.1. We can think of 2-complexes as 2-dimensional simplicial complexes with the following relaxations.

- The 1-skeleton can have loops and parallel edges.

- Faces can be any family of closed walks of edges, not just triangles.

Definition 2.3.2. Consider an embedding of a 2-complex C in $\mathbb{R}^3$. The connected components of $\mathbb{R}^3 - C$ are called *chambers of the embedding* of C. When the embedding is clear from the context, we will just say the *chambers* of C.

Consider an embedding of a 2-complex C in $\mathbb{R}^3$. We have that C is bounded, therefore it fits in some ball B. The complement of B is contained within a single chamber, which is unbounded. All the other chambers are within B, so they are bounded. Hence, there is a unique unbounded chamber.

Definition 2.3.3. We will call the unbounded chamber *outer chamber* and the bounded chambers *inner chambers.*

Consider an embedding in $\mathbb{R}^3$ of the 2-complex C and a 2-sphere S, centered at v, which is small enough so that it intersects C only in edges and faces that are incident to v. This sphere intersects each non-loop edge of C in a point and each loop in two points. For every occurrence of the vertex v in a face, there are two edges that meet at this occurrence of v. The sphere intersects C in one arc for every occurrence of v and this arc connects the two points corresponding to the two edges meeting at this occurrence of v. Thus, the resulting intersection is a graph, which we call the link graph of C at the vertex v.

This geometric definition gives us a recipe to define an abstract link graph by considering the incidence of edges and faces containing the vertex v.

Consider a 2-complex C with a vertex v. Its *link graph* $L(v)$ is defined as follows. For each edge e of C, incident with v we have a vertex e^+ in $L(v)$ if v is the start vertex of e and a vertex e^- in $L(v)$ if v is the end vertex of e. This way if an edge is not a loop it has only one vertex in $L(v)$ and if it is a loop it has two vertices because for loops the start vertex and the end vertex are the same. This agrees with the geometric picture where the sphere intersects a non-loop in one vertex and a loop in two vertices. A face f of C is an oriented closed walk of its underlying graph G. For every occurrence of v in this oriented closed walk there is a pair of edges e_i, e_{i+1} of f such that $\varphi^{\sigma_i}(e_i) = v = \varphi^{-\sigma_{i+1}}(e_{i+1})$. For every such occurrence we get one edge $e_i^{\sigma_i} e_{i+1}^{-\sigma_{i+1}}$ in the link graph $L(v)$.

Example 2.3.4. Consider the following 2-complex. It has the vertex set $\{W, X, Y, Z\}$. It also has an edge x, connecting W and X, an edge y, connecting X and Y, an edge z, connecting Y and Z, an edge w, connecting Z and W, an edge t, connecting W and Y, and a loop ℓ at the vertex W. Its faces are the walks $f_1 = \ell$, $f_2 = tz\ell w$, $f_3 = tzw$ and $f_4 = wxyz$.

The link graph at W has vertices ℓ^-, ℓ^+, z and w. From f_1 we have the edge $\ell^-\ell^+$. From f_2 we have the edges $z\ell^-$ and ℓ^+w. From f_3 we have the edge zw. From f_4, we have the edge zw. Overall, we get that $L(W)$ is a cycle of length 4, where one of the edges is doubled with a parallel edge.

Lemma 2.3.5. *All link graphs of an embeddable 2-complex are planar.*

Proof. Let C be an embeddable 2-complex and consider a fixed embedding of C in $\mathbb{S}^3$. In this embedding consider an arbitrary vertex v. Let $G(v)$ be the intersection of C and a sufficiently small 2-sphere centered at v. Recalling the geometric definition of a link graph, we can conclude that $G(v)$ is isomorphic to the link graph $L(v)$. Since $G(v)$ is embedded in the 2-sphere by construction, it follows that $L(v)$ is planar. □

Definition 2.3.6. Given a 2-complex C without loops, the *(2-dimensional) cone* over C is the following 2-complex. It is obtained from C by adding a single vertex (referred to as the *top* of the cone), one edge for every vertex of C from that vertex to the top, and one triangular face for every edge e of C whose endvertices are the endvertices of e and the top.

We denote the 2-dimensional cone over a 2-complex C by $\widehat{C}$.

Observation 2.3.7. *The link graph at the top of a cone of a 2-complex C is equal to the 1-skeleton of C.*

Proof. Let us denote the 1-skeleton of C by G. For each vertex v in G there is an edge $e_v = tv$ in the cone $\widehat{C}$. The edge e_v corresponds to a vertex $t_v \in L(t)$. Similarly for each edge uv in G there is a face tuv in $\widehat{C}$, which in turn corresponds to an edge t_ut_v in

$L(t)$. Taking this into consideration, we see that the map $v \to t_v$ is a graph isomorphism between the 1-skeleton G and the link graph $L(t)$ which completes the proof. $\square$

From the above, we can derive the most basic embeddability obstructions as follows.

Corollary 2.3.8. *The cone over $K_{3.3}$ and the cone over K_5 are not embeddable.*

Proof. Let C be the cone over $K_{3,3}$. By Observation 2.3.7, the link graph at the top of C is $K_{3,3}$. Then, by Lemma 2.3.5, it follows that C is not embeddable. The result follows similarly for K_5. $\square$

Let us recall a basic topological definition.

Definition 2.3.9. A topological space X is *simply connected* if it is path-connected and every loop in X is homotopic to a point.

Observation 2.3.10. 2*-dimensional cones are always simply connected.*

Proof. Firstly, any point on the cone $\widehat{C}$ is on a face containing the top t, therefore there is a path from t to any point of $\widehat{C}$. This means that $\widehat{C}$ is path connected. Now consider a loop in the cone $\widehat{C}$ based at t. Any such loop is homotopy equivalent to a loop ℓ restricted to the 1-skeleton of $\widehat{C}$. A homotopy equivalence between ℓ and the constant loop is given by the function $H(x, s) = \ell(x(1 - s))$, where x is the variable of ℓ and s is the homotopy variable. This function stays within $\widehat{C}$ because the edges in C run along their respective faces in the cone and the edges not in C run along themselves. Thus, we proved that $\widehat{C}$ is path-connected and loops are equivalent to constant loops or in other words that $\widehat{C}$ is simply-connected as claimed. $\square$

Definition 2.3.11. A 2-complex C is locally k-connected if at each vertex v of C, the link graph $L(v)$ is k-connected.

We will mostly use this definition with $k = 2$ and $k = 3$.

A rotation system of a graph G is a family $(\sigma(v)|v \in V(G))$ of cyclic orientations $\sigma(v)$ of the edges incident with the vertices v of G. The orientations $\sigma(v)$ are called rotators. Any rotation system of a graph G induces an embedding of G in an oriented (2-dimensional) surface S. To be precise, we obtain S from G by gluing faces onto (the geometric realisation of) G along closed walks of G as follows. Each directed edge of G is in one of these walks. Here the direction $\vec{a}$ is directly before the direction $\vec{b}$ in a face f if the endvertex v of $\vec{a}$ is equal to the starting vertex of $\vec{b}$ and b is just after a in the rotator at v. The rotation system is planar if that surface S is a disjoint union of 2-spheres. Note that if the graph G is connected, then for any rotation system of G, also the surface S is connected. A rotation system of a (directed) 2-complex C is a family $(\sigma(e)|e \in E(C))$ of cyclic orientations $\sigma(e)$ of the faces incident with the edge e. A rotation system of a 2-complex C induces a rotation system at each of its link graphs $L(v)$ by restricting to the edges that are vertices of the link graph $L(v)$; here we take $\sigma(e)$ if e is directed towards v and the reverse of $\sigma(e)$ otherwise. A rotation system of a 2-complex is planar if all induced rotation systems of link graphs are planar. The existence of a planar rotation system for a 2-complex does not ensure embeddability in $\mathbb{S}^3$ because embeddings in unorientable 3-manifolds can also induce planar rotation systems. However as we will later see, if we restrict our attention to simply connected 2-complexes, then a planar rotation system is sufficient to ensure embeddability in $\mathbb{S}^3$.

Planar rotation systems of graphs and 2-complexes naturally arise in the context of embeddability. To show this, let G be a graph embedded in $\mathbb{S}^2$ and fix one of the two possible rotation directions ρ. For each vertex orient its rotator according to ρ to obtain a rotation system σ. To see that σ is planar, note that if we apply the construction from above to it, we recreate the original embedding of G on $\mathbb{S}^2$. Given an embedding of a 2-complex C in $\mathbb{S}^3$, we can similarly construct a rotation system Σ of C. We do this by choosing a rotation direction ρ and we order the rotator at each edge e in the direction of ρ relative to the orientation of e. To see that Σ is planar, note that the rotation systems of the link graphs of C induced by it are planar by Lemma 2.3.5.

2.4 Space minors

Definition 2.4.1. A *space minor* of a 2-complex is obtained by successively performing one of these two operations.

1. contracting an edge that is not a loop;

2. deleting a face (and all edges or vertices only incident with that face);

This definition is going to be used as a tool for the classification of outerspatial 2-complexes in Chapter 3. It is based on the space minor definition of Carmesin in [12] where he defines it as a tool for a similar classification of embeddable 2-complexes. Since embeddability is more general than outerspatiality, the space minor definition in [12] is stronger – it has 3 additional operations apart from the two already given in Definition 2.4.1. In this thesis, when we use the term space minor, we will always mean Definition 2.4.1 and we mention the other definition only for reference.

The edge contraction operation can be easily generalised to contraction of a set of edges spanning a forest. We do that by giving an arbitrary ordering on this set of edges and then contracting the edges one by one according to the given order.

It is easy to see that the space minor operations preserve embeddability in $\mathbb{S}^3$ as well as in any other 3-manifold. Furthermore, we can show that the space minor relation is well-founded.

Lemma 2.4.2. *The space minor operations can be (non-trivially) applied only a finite amount of times.*

Proof. The face degree of an edge e is the number of faces incident with e. We consider the sum S of all face degrees ranging over all edges. The two operations always strictly decrease S, hence we can apply 1 or 2 only a finite number of times. $\square$

2.5 Matroids

There exist numerous ways to define, interpret and use matroids, some of which we will show below. The perspective of matroids that will interest us the most is that of graphic matroids, or matroids which come from the cycle space of a graph. This perspective comes in useful in the context of duality in planar graphs. If a graph G is planar, we can construct a dual graph from any embedding of G. However, this dual is not always unique as it depends on the embedding and it exists only when G is planar. This issue is resolved by introducing the dual matroid of a graph. It coincides with the cycle matroid of any dual graph of G when G is planar. It is unique and extends to cases where G is not planar. The main ways that we will utilise matroids will be in defining 'locality' of a 2-complex and the dual matroid of a 2-complex, both of which will be shown later in the section. We start by introducing one of the most standard definitions of a matroid.

Definition 2.5.1. A matroid M is a pair $(E, \mathcal{C})$, where E is a finite set, called the edges of $\mathcal{C}$ and $\mathcal{C}$ is a family of subsets of E, called the *circuits* of M, with the following properties.

(C1) The empty set is not a circuit, or in other words, $\emptyset \notin \mathcal{C}$.

(C2) $\mathcal{C}$ is an antichain, or in other words, for each pair of sets $C' \subseteq C \subseteq E$, we have that if $C, C' \in \mathcal{C}$, then $C = C'$.

(C3) $\mathcal{C}$ satisfies circuit elimination, or in other words, for any two distinct $C, C' \in C$ containing the common element $e \in E$, there is some $C'' \in \mathcal{C}$ with $C'' \subseteq (C \cup C') - e$.

The most basic type of matroids is as follows.

Definition 2.5.2. The uniform matroid U_n^r is defined over set of n elements. A subset of elements is a circuit in U_n^r if and only if it has exactly $r + 1$ elements.

There are a number of (equivalent) alternative definitions of matroids using different objects to define them, for example using bases or independent sets instead of circuits. Next, we define a matroid over a finite field. Before we start, let us note that we call a matrix or a matroid *p-ary* if it is defined over the field $\mathbb{F}_p$.

Definition 2.5.3. Take an arbitrary prime number p. A p-ary matroid is defined using a matrix by its set of edges and its set of circuits in the following way. Start with the p-ary matrix A. The edges of the matroid are the columns of A. The circuits of the matroid are the sets of edges corresponding to minimal non-empty linearly dependent sets of column vectors. We will denote a matroid obtained from the matrix A in such a way by $M(A)$. The dual matroid $M^*(A)$ of the matroid $M(A)$ can be defined by the same matrix in the following way. The edges of $M^*(A)$ will be again the columns, while the circuits of $M^*(A)$ are the non-empty minimal supports of the row space of A. The circuits of $M^*(A)$ are also called *cocircuits* of $M(A)$.

Definition 2.5.4. Given a graph G its cycle matroid can be defined over any field $\mathbb{F}_p$ using a matrix as in Definition 2.5.3. We label the rows of the matrix A with the vertices of G and the columns by the edges of G. At an entry that intersects the vertex v with the edge e, put a value 0 if they are not incident, a value 1 if e comes out of v in the graph G and a value -1 if e goes in v in the graph G.

Definition 2.5.5. A *directed simplicial complex* is a simplicial complex C together with an assignment of a direction to each edge of C and together with an assignment of a cyclic orientation to each face of C. A *signed incidence vector* of an edge e of C has one entry for every face f; this entry is zero if e is not incident with f, it is plus one if f traverses e positively and minus one otherwise.

The matrix given by all signed incidence vectors is called the *(signed) edge/ face incidence matrix* and we denote it by $A(C)$. The *cycle matroid of the simplicial complex* C is the matroid $M(A(C))$ defined over some field $\mathbb{F}_p$ and is denoted by $M(C)$. Similarly, the *dual matroid* of C is the matroid $M^*(A(C))$ and is denoted by $M^*(C)$.

The fields that we are usually going to use are $\mathbb{F}_2$ and $\mathbb{F}_3$.

As we have already mentioned, the term 'dual matroid' of a graph is motivated by the fact that if a graph G is planar and has some dual graph G^* then the dual matroid of G is equal to the cycle matroid of G^*. There is a similar motivation for the term dual matroid of a 2-complex and it comes from the following theorem.

Theorem 2.5.6. *[15, Theorem 1.1] Let C be a directed 2-dimensional simplicial complex embedded into $\mathbb{S}^3$. Then the edge/face incidence matrix of C represents over the integers a matroid M which is equal to the cycle matroid of the dual graph of the embedding.*

Where the dual graph of an embedding of a 2-complex is defined in the following way. It has a vertex for each chamber of C and two vertices are connected by one edge for each face their corresponding chambers share.

Example 2.5.7. Consider a tetrahedron with faces f_1, f_2, f_3, f_4. Let the edges of f_1 be e_1, e_2, e_5, of f_2 be e_2, e_3, e_6, of f_3 be e_1, e_3, e_4 and of f_4 be e_4, e_5, e_6. Orient all faces clockwise as seen from outside. Orient e_1 against the orientation of f_1, e_2 against f_2, e_3 against f_3 and orient e_4, e_5, e_6 along the orientation of f_4. It can be checked that this construction is possible and uniquely determined. The signed edge face incidence matrix A of this directed simplicial complex is

$$\begin{pmatrix} * & f_1 & f_2 & f_3 & f_4 \\ e_1 & -1 & 0 & -1 & 0 \\ e_2 & 1 & -1 & 0 & 0 \\ e_3 & 0 & 1 & 1 & 0 \\ e_4 & 0 & 0 & -1 & 1 \\ e_5 & 1 & 0 & 0 & 1 \\ e_6 & 0 & 1 & 0 & 1 \end{pmatrix}$$

.

There is no proper subset of column vectors that is linearly dependent, while the set $\{f_1, f_2, f_3, f_4\}$ is linearly dependent as $f_1 - f_3 + f_2 - f_4 = 0$, so the only circuit of $M(A)$ is $\{f_1, f_2, f_3, f_4\}$. Thus, the dual matroid of the tetrahedron is the dual matroid of $M(A)$, in other words the matroid with edges f_1, f_2, f_3, f_4 and each pair of edges is in a circuit.

Definition 2.5.8. The *cycle matroid of a graph* G is denoted by $M(G)$ and consists of the pair $(E, \mathcal{C})$, where E is the edge set of G and $\mathcal{C}$ is the set of edge sets of cycles of G.

Definition 2.5.9. The *dual matroid of a graph* G is denoted by $M^*(G)$ and consists of the pair $(E, \mathcal{C})$, where E is the edge set of G and $\mathcal{C}$ is the set of bonds of G.

Remark 2.5.10. A matroid M is called *graphic*, there exists a graph G such that $M = M(G)$.

It is easy to check that for a fixed graph G, the matroids $M(G)$ and $M^*(G)$ are dual to each other.

Definition 2.5.11. Consider a matroid M with a set of edges E and a set of circuits $\mathcal{C}$. We define the matroid $M - e$ to be the pair $(E - e, \{c \in \mathcal{C} | e \notin c\})$ and we call this operation deleting an edge from M. We define the matroid M/e to be the pair $(E - e, \{c - e | c \in \mathcal{C}\})$ and we call this operation contracting an edge from M. A minor of the matroid M is a matroid N obtained by M by deleting and contracting some of its edges.

It can be checked that contracting and deleting edges of a graph G is equivalent to contracting and deleting edges in its cycle matroid. Algebraically this states $M(G)/e = M(G/e)$ and $M(G - e) = M(G) - e$.

Definition 2.5.12. Given a matroid $M = (E, \mathcal{C})$ and a subset $X \subseteq E$ we define the restriction $M \upharpoonright X$ of the matroid M to the set X as follows.

$$M \upharpoonright_X = M/(E - X)$$

Definition 2.5.13. Given a vertex v in a 2-complex C, we define $F(v)$ to be the set of faces incident to v.

Definition 2.5.14. Consider a 2-complex C with a dual matroid $M^*(C)$. For an arbitrary vertex v of the 2-complex C and its link graph $L(v)$, we define $M^*(v)$ to be the dual matroid of $L(v)$.

A 2-complex is called *local* if for every vertex v, the dual matroid satisfies $M^*(v) = M^*(C) \upharpoonright F(v)$.

Example 2.5.15. Consider a tetrahedron and a vertex v. Then $L(v)$ is a triangle so $M^*(v)$ consists of 3 edges with each pair in a circuit. The dual matroid consist of 4 edges with each pair in a circuit. When we restrict $M^*(C)$ to the faces in $F(v)$ we get 3 edges with each pair in a circuit. Thus, $M^*(C) \restriction F(v) = M^*(v)$. So a tetrahedron is local.

If we remove the face opposite v, the link graph $L(v)$ stays unchanged and so the dual matroid $M^*(v)$ also stays unchanged. However the dual matroid of C consists of loops in this case, so $M^*(C) \restriction F(V) = M^*(v)$ cannot hold. So a tetrahedron with a face removed is not local.

Definition 2.5.16. Let k be a finite field. A set $\mathcal{S}$ of vectors in k^E is *sparse* if for each coordinate $e \in E$ there are precisely zero or precisely two vectors $v \in \mathcal{S}$ with $v(e) \neq 0$.

There is also a definition of sparse in which it is allowed to have at most two non-zero vectors in each coordinate instead of exactly zero or exactly two. Even though these definitions are different at face value, for our purposes they are interchangeable and we choose to use Definition 2.5.16.

Chapter 3

Characterisation of outerspatial 2-complexes

3.1 Introduction

An important class of planar graphs is the class of *outerplanar graphs*, those graphs with plane embeddings having a face containing all vertices. Kuratowski's[1] characterisation of planar graphs in terms of excluded minors implies a characterisation of outerplanar graphs in terms of excluded minors.

This paper is part of a project aiming to extend theorems from planar graph theory to three dimensions. The starting point of this project is [12]. In there a three-dimensional analogue of Kuratowski's theorem was proved: embeddability of simply connected 2-complexes in the 3-sphere was characterised by excluded 'space minors'. O-joung Kwon asked[2] whether a similar result is true for a natural higher dimensional analogue of outerplanar graphs.

One of the equivalent definitions for a graph to be outerplanar is that the 1-dimensional cone over it is planar. We will take this definition one dimension higher and will call a 2-complex *outerspatial* if the 2-dimensional cone over it is embeddable in $\mathbb{R}^3$. Hence, outerspatial 2-complexes are the natural generalisation of outerplanar graphs to three

[1]Wagner's characterisation is stated in terms of minors while Kuratowski characterisation is stated in terms of subdivisions. They are equivalent, but since Kuratowski result comes first, we refer to either statement as Kuratowski's theorem.

[2]In private communication.

dimensions. The main result of this paper answers the above mentioned question of O-joung Kwon affirmatively, and is the following.

Theorem 3.1.1. *A locally* 2*-connected*[1] *simple* 2*-complex* C *is outerspatial if and only if it does not contain a surface of positive genus or a space minor with a link graph that is not outerplanar.*

The obstructions given in the theorem above are necessary: firstly, the torus has genus one and any 2-complex homeomorphic to a torus is not outerspatial by Lemma 3.3.3 below. Secondly, cones over non-outerplanar graphs have link graphs that are not outerplanar and thus are not outerspatial 2-complexes by Lemma 3.3.5 below.

Using the three-dimensional Kuratowski characterisation [12, Theorem 1.3] one can obtain a characterisation of the class of outerspatial 2-complexes in terms of excluded minors. However, the set in question cannot be defined clearly and is too large to be of practical use. Our main result, Theorem 3.1.1, provides a simple and short forbidden structures characterisation.

In order to have such a short list as in Theorem 3.1.1, the assumption of local 2-connectedness is also necessary. Indeed, triangulations of the torus with a single topological disc removed (and various slight modifications of higher genus surfaces) are excluded minors of the class of outerspatial 2-complexes. Hence, for this super-class the list of excluded minors is much more complicated than that of Theorem 3.1.1.

Example 3.1.2. Consider a 2-complex built by gluing a set of triangulated 2-spheres step by step such that at each step the gluing set is a single face. Any 2-complex built this way is an example of a locally 2-connected outerspatial 2-complex. We will show in Section 3.5 that all such 2-complexes can be constructed in this way.

Example 3.1.3. The topological space Bing's House, as described in [39, Chapter 0], has only one chamber but it is not outerspatial. It is not outerspatial because it has a link graph isomorphic to K_4.

[1]For a definition of locally 2-connected look at Definition 3.2.12

Proof that the Bing's House has a link graph isomorphic to K_4. Consider the link graph at the vertex v as shown in the representation of the Bing's house in Figure 3.1. There are 4 edges incident to v, so $L(v)$ has 4 vertices. It is easy to check that all pairs of edges are adjacent in some face, hence all pairs of vertices in $L(v)$ are connected by an edge. Hence $L(v)$ is isomorphic to K_4. □

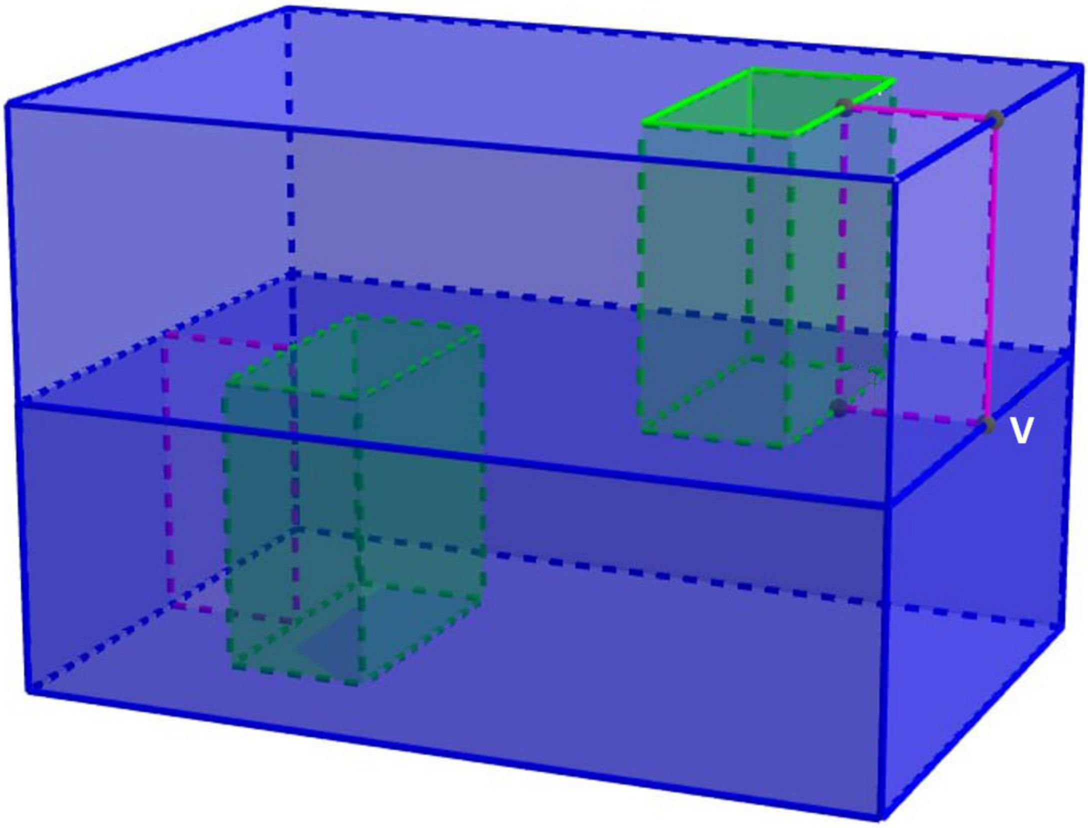

Figure 3.1: A 2-complex representation of the Bing's house.

Remark 3.1.4. Throughout this paper, faces of 2-complexes are bounded by genuine cycles (of arbitrary length), as restricting to faces of size three makes the question posed in Theorem 3.1.1 one-dimensional. See Proposition 3.2.34 for details.

In this paper we find a correspondence between particular plane embeddings and outerspatial embeddability in 3-space, as follows. Given a set of cycles $\mathcal{C}$ in a plane graph G, we say that the pair $(G, \mathcal{C})$ has a *nested plane embedding* if G has an embedding in the plane such that any two cycles in $\mathcal{C}$ do not intersect internally (for a more precise definition, look at Definition 3.2.26). The 2-complex *associated to* $(G, \mathcal{C})$ is the 2-complex whose 1-skeleton is G and whose set of faces is $\mathcal{C}$. We prove the following connection between outerspatial 2-complexes and nested plane embeddings.

Corollary 3.1.5. *A graph G together with a set of cycles $\mathcal{C}$ has a nested plane embedding if and only if its associated 2-complex is outerspatial.*

This result follows from Lemma 3.2.31. Given this corollary, Theorem 3.1.1 can be applied directly to characterise the existence of nested plane embeddings of graphs.

Related Results. The two most important concepts of this paper are embeddings of 2-complexes in 3-space and nested plane embeddings of graphs. Our methods for embedding 2-complexes is related to and based on the series of papers on this topic [12, 13, 14, 15, 17]. Some previous works related to nested plane embeddings focus on the triangle case. Such special types of nested plane embeddings are studied in papers [63, 36, 44]. The first one explores properties of graphs in relation to structural information on these nested triangles. The other two papers use nestedness as a tool to find an example of minimal area straight line drawings of planar graphs. The term 'laminar' is a general notion relating to sets, but is also used with the same meaning as our definition in terms of nested plane embeddings. Its usage in the context of cycles is motivated by the fact that the interiors of the faces bounded by a set of laminar cycles form a family of laminar subsets of $\mathbb{R}^2$. Laminar cycles play central part in the papers [35, 28, 5, 24]. In [35] the main problem is finding a minimum-weight set of vertices that meets all cycles in the subset. There the authors optimise an algorithm that they have found over laminar sets of cycles. In [28] the aim is to bound the number of odd cycle vertex packings by the number of odd cycle vertex transversal. A main idea in proving this is considering laminar sets of odd cycles. In [5] laminar cycles are used to count 3-colourings of triangle-free planar graphs. In [24] laminar cycles are used to find maximal sets of laminar 3-separators in 3-connected planar graphs. In [31] it was shown that if a set of nested cycles in a graph satisfies some further properties, then this graph has a packing of k odd cycles if and only if $G - v$ does for some specific v. There is also the notion of simply nested k-outerplanar graphs, which is somewhat related to our project; see [3] for definitions.

The structure of this paper is as follows. In the second section we give some basic definitions and prove some initial results. In the third section we build up to and state the Core Lemma – Lemma 3.3.7, which is the key component of the proof of Theorem 3.1.1. In the fourth section we prove the main techniques needed for the Core Lemma and we complete its proof and consequently prove Theorem 3.1.1. The fifth section is devoted to deriving some properties of locally 2-connected simple outerspatial 2-complexes following from our results.

3.2 Basic definitions and initial approaches

We start this section by giving basic definitions related to 2-complexes. Next, we will explore various ways of defining the concept of outerspatiality and provide a brief explanation as to why we believe our definition to be the most effective. Then we start building the theory needed to prove the main theorem. At the end we show a proposition which on its own proves a rudimentary version of the main result, but the idea behind it is also quite useful later on.

Let us note that in this paper when we talk about a graph, it is assumed that it can have parallel edges and loops. We will now define what a 2-complex is.

Definition 3.2.1. A 2-*complex* is a graph $G = (V, E)$ together with a family[1] F of cycles, called its *faces*.

Remark 3.2.2. In this paper, we will assume that all edges of a 2-complex C lie on some face of C.

Definition 3.2.3. We will call a 2-complex *simple* if it (that is, its underlying graph) does not have loops or parallel edges.

[1] Parallel faces are not relevant to the question of embeddability. That is why in this paper we will assume that this family is a set.

Example 3.2.4. A 2-dimensional simplicial complex is a simple 2-complex where all faces have three edges.

Definition 3.2.5. The 1-skeleton of a complex $C = (V, E, F)$ is the graph $G = (V, E)$.

A notion that underlies this paper is that of space minors, the 3-dimensional analogue of graph minors[1].

Definition 3.2.6. A *space minor* of a 2-complex is obtained by successively performing one of these two operations.

1. contracting an edge that is not a loop;

2. deleting a face (and all edges or vertices only incident with that face);

Remark 3.2.7. For detailed discussion on these operations and a proof that they are well-founded and preserve embeddability in 3-space, see [12].

The aim of this paper is to extend the notion of outerplanarity of graphs to three dimensions. Before we do this, we need a definition of outerplanarity that translates well to 2-complexes. There are two ways of defining outerplanar graphs that suit our purposes. One is to find an 'outer' face containing all vertices and the other is through planarity of the cone. These two definitions are shown below.

Definition 3.2.8. Let G be a graph. Take the disjoint union of G and an additional vertex t and connect t to all vertices of G by an edge. The resulting graph is called the *(1-dimensional) cone* over G and the vertex t is called the top of the cone.

Definition 3.2.9. (Outerplanarity criterion 1) A graph G is outerplanar if it can be embedded in the plane in such a way that there is a face of the embedding containing all vertices of G.

Definition 3.2.10. (Outerplanarity criterion 2) A graph G is outerplanar if the 1-dimensional cone over G is planar.

[1] The systematic study of the minor relation was initiated by Klaus Wagner. This is attributed to the fact that a fundamental result in graph minors theory was conjectured by him – for reference, check [57].

The fact that these two definitions are equivalent is a well-known result. A proof sketch goes as follows.

If a graph G is outerplanar by criterion 1, we can embed it in the plane so that there exists a face containing all vertices of G. Then we can add a vertex on the interior of this face and connect it to all vertices of this face and hence all vertices of the graph. Thus, we embedded the cone over G in the plane, which shows that criterion 1 implies criterion 2.

If G is outerplanar by criterion 2, we can embed the cone over G in the plane. When we delete the top of the cone, the connected component of the point corresponding to the deleted vertex is the interior of a face that contains all vertices of G. So there is a face containing all vertices of G, which shows that criterion 2 implies criterion 1.

Definition 3.2.11. Consider a vertex v of the 2-complex C and define the following graph. Its vertices are the edges of C incident to v and two vertices of $L(v)$ are connected by an edge if their corresponding edges in C lie on the same face. This graph is called the *link graph* of v and is denoted by $L(v)$.

Definition 3.2.12. A 2-complex whose link graphs are all k-connected simple graphs is called *locally k-connected.*

Definition 3.2.13. Given a 2-complex C without loops, the *(2-dimensional) cone* over C is the following 2-complex. It is obtained from C by adding a single vertex (referred to as the *top* of the cone), one edge for every vertex of C from that vertex to the top, and one triangular face for every edge e of C whose endvertices are the endvertices of e and the top.

We denote the 2-dimensional cone over a 2-complex C by $\widehat{C}$.

Observation 3.2.14. *The link graph at the top of a cone of a 2-complex C is equal to the 1-skeleton of C.*

Proof. For a proof, see Observation 2.3.7. □

Observation 3.2.15. *2-dimensional cones are always simply connected.*

Proof. For a proof, see Observation 2.3.10. □

Definition 3.2.16. The *geometric realisation* of a 2-complex $C = (V, E, F)$ is the topological space obtained by gluing discs to the geometric realisation of the graph $G = (V, E)$ along the face boundaries.

Definition 3.2.17. A (topological) embedding of a simplicial complex C into a topological space X is an injective continuous map from (the geometric realisation of) C into X. We say that a 2-complex is *embeddable* in $\mathbb{R}^3$ if its geometric realisation is embeddable in $\mathbb{R}^3$ as a topological space.

Remark 3.2.18. If a 2-complex is embeddable in $\mathbb{R}^3$ we will say as a shorthand that the 2-complex is *embeddable*.

Now we are ready to give the definition of an outerspatial 2-complex. We have the two definitions of outerplanarity, Definition 3.2.9 and Definition 3.2.10. We can make two different definitions for outerspatial 2-complexes based on them.

Definition 3.2.19. (Outerspatiality criterion 1) A 2-complex C is *weakly outerspatial* if there is an embedding of C in $\mathbb{R}^3$ such that some chamber of this embedding is incident to all of the edges of the 2-complex.

Definition 3.2.20. (Outerspatiality criterion 2) A 2-complex is *outerspatial* if its 2-dimensional cone embeds in 3-space.

It would be best if these two definitions were equivalent in the same way the two outerplanarity definitions are. However, this is not the case. It turns out that outerspatiality criterion 2 is a stronger definition as shown by the lemma below.

Lemma 3.2.21. *If a 2-complex C is outerspatial, then it is also weakly outerspatial.*

Proof. Consider an embedding of the cone over C in $\mathbb{R}^3$; assume that the cone is embedded in the outer face. Delete the top of the cone with all incident edges and faces. The chamber where the top was includes all faces incident with the top and thus has all edges of C in its boundary. Thus, this defines a weakly outerspatial embedding of C. □

We showed that outerspatiality implies weak outerspatiality. To show that it is a strictly stronger definition we need an example of 2-complex which is weakly outerspatial but not outerspatial. This is the Bing's house – Example 3.1.3.

For this paper we have chosen the second definition of outerspatial (Definition 3.2.20) as it is more specific and yields more interesting and relevant characterisations.

Below, we will need to use a notion of inside and outside of a sphere. The following theorem provides the definition that we want.

Theorem 3.2.22. *(Jordan–Brouwer separation theorem [59]) Any compact, connected hypersurface X in $\mathbb{R}^n$ will divide $\mathbb{R}^n$ into two connected regions; the 'outside' D_0 and the 'inside' D_1. Furthermore, $\bar{D}_1$ is itself a compact manifold with boundary $\partial\bar{D}_1 = X$.*

Definition 3.2.23. The *interior of a cycle* in a plane embedding of a graph is the inside of its image as defined in Theorem 3.2.22.

Let us recall the following basic lemma (Lemma 2.3.5).

Lemma 3.2.24. *All the link graphs of an embeddable 2-complex are planar.*

Definition 3.2.25. Consider two cycles embedded in the plane. We say that they *intersect internally* if their interiors have a proper non-empty intersection.

Definition 3.2.26. Consider a planar graph G with a set of cycles $\mathcal{C}$. We say that a plane embedding of G is *nested* if no two cycles in $\mathcal{C}$ intersect internally.

Definition 3.2.27. Take a graph embedded in a 2-sphere S which is in turn embedded in the Euclidean space $\mathbb{R}^3$. Glue discs to cycles of this graph inside the sphere so that they do not intersect each other in interior points, and they intersect the sphere S precisely in the gluing cycles. Call a topological space that can be obtained in this way *outerspherical.*

Definition 3.2.28. We will call a 2-complex C *outerspherical*, if it has an embedding in $\mathbb{R}^3$ that is an outerspherical topological space, where the 1-skeleton of C is mapped to the graph in the sphere and the faces of C are mapped to the discs glued to the graph as described in Definition 3.2.27.

Definition 3.2.29. Consider a graph G with a set of vertices V. We will call a cyclic orientation of the edges incident to a vertex $v \in V$ a *rotator* at v.

Before we start with the next lemma, we need a definition which will help us differentiate between the faces of a graph and the faces of a 2-complex.

Definition 3.2.30. Consider an embedding of a planar graph G on the sphere $\mathbb{S}^2$. We will call the connected components of $\mathbb{S}^2 - G$ the *facets* of this embedding.

Lemma 3.2.31. *Let C be a 2-complex. Then the following are equivalent*

(1) C is outerspherical

(2) C is outerspatial

(3) The 1-skeleton of C together with the set of face boundaries of C has a nested plane embedding.

Proof. For the (1) $\implies$ (2) implication, consider an outerspherical 2-complex C. There exists an embedding of it in $\mathbb{R}^3$ such that the unit sphere intersects C in its 1-skeleton and everything else of C is embedded in the interior of the unit ball. Let G be the plane embedding of the 1-skeleton of C of this embedding in the unit ball. The cone over $\mathbb{S}^2$ is the full 3-dimensional unit ball. Assume that G is drawn onto $\mathbb{S}^2$. This way we obtain an embedding of the cone over G in the full unit ball that takes the embedding G into $\mathbb{S}^2$ at the boundary. Glue these two full unit balls by gluing at G to obtain an embedding of the cone over C in $\mathbb{R}^3$. This shows that C is outerspatial.

For the (2) $\implies$ (3) implication, consider a 2-complex C and suppose that it is outerspatial. In other words, its cone $\widehat{C}$ is embeddable. Consider an embedding of $\widehat{C}$

and let the top of this cone be t. The embedding of $\widehat{C}$ induces an embedding of the link graph at the top in the 2-sphere. As the link graph at the top is the 1-skeleton of C, the embedding of $\widehat{C}$ induces an embedding of the 1-skeleton of C in the plane. Call that embedding ι. Denote the 1-skeleton of C by G and the set of face boundaries of C by F.

Sublemma 3.2.32. *The embedding ι of G is a nested plane embedding for the set F of cycles.*

Proof. Suppose not for a contradiction. As the embedding ι is plane by construction, there are two cycles in the set F that are not nested; that is, they intersect internally. So by Theorem 3.2.22, there is a vertex v of the link graph $L(t)$ such that the cycles intersect internally at v. That is, there are edges $(e_i|i \in \mathbb{Z}_4)$ of $L(t)$ incident with v that appear in that order at the rotator at v such that e_1 and e_3 are in one of the cycles of F and e_2 and e_4 are in the other of these cycles.

Now we find a vertex w of C such that the embedding of the link graph $L(w)$ induced by the embedding of $\widehat{C}$ is not planar; this will be the desired contradiction. Let w be the endvertex of v, considered as an edge of the cone $\widehat{C}$, aside from the top t. Consider the link graph $L(w)$. Now v is an edge of $L(w)$, and also the edges e_i are edges of $L(w)$ incident with v; and the rotator at v is the same as in $L(t)$ (up to reversing). Here, however, as the edges e_1 and e_3 are in a common face of C, their endvertices in $L(w)$ aside from v are joined by an edge, call it x_1. Similarly, the endvertices of the edges e_2 and e_4 aside from v are joined by an edge; call it x_2. The subgraph of $L(w)$ with the six edges $(e_i|i \in \mathbb{Z}_4)$ and $(x_i|i = 1, 2)$ is not planar with the specified rotator at v, see Figure 3.2. This is the desired contradiction. □

We started with the assumption that C is outerspatial and we obtained the result from Sublemma 3.2.32. This completes the (2) $\implies$ (3) implication.

For the (3) $\implies$ (1) implication suppose that we have a complex C such that its 1-skeleton G together with the set F of face boundaries has a nested plane embedding. We prove by induction on the number of elements of F that do not bound a facet of the

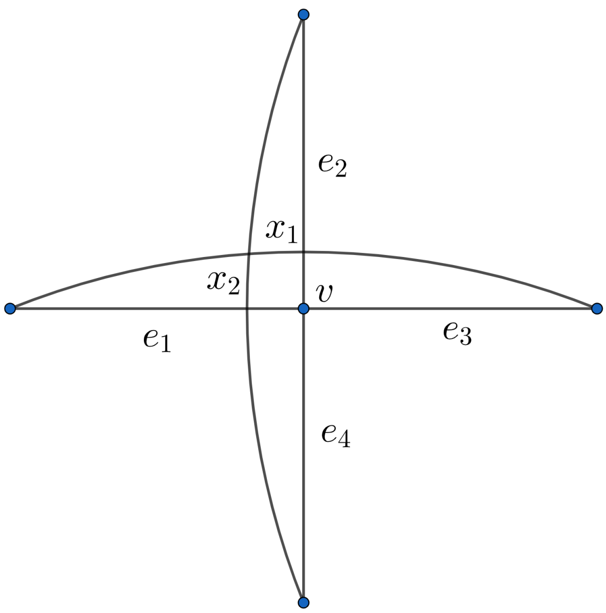

Figure 3.2: A subgraph of the link graph $L(w)$ that is not planar for the induced rotator at the vertex v.

nested plane embedding that C has an outerspherical embedding. If the number is zero, we get that C is an embedding of G on the sphere, which covers the base case. Now, let $f \in F$ that does not bound a facet be given. By Theorem 3.2.22, f divides the plane embedding of G into two subgraphs G_1 and G_2 intersecting at the cycle f. Let F_i be the set of faces in F attaching at G_i for $i = 1, 2$. As f does not bound a facet, each G_i is strictly smaller than G. Hence, by induction, each 2-complex C_i obtained from G_i by adding the faces from F_i has an outerspherical embedding, which includes the face f. Now glue these two embeddings at the face f to obtain an outerspherical embedding of G.

We proved that implications (1) $\implies$ (2) $\implies$ (3) $\implies$ (1), therefore all these three statements are equivalent as claimed. □

The equivalence of a 2-complex being outerspatial and having a nested plane embedding of the 1-skeleton is an interesting connection between graphs and 2-complexes which will also be important for the proof of the main result.

The equivalence of a 2-complex being outerspatial and being outerspherical is a good geometric characterisation of outerspatial 2-complexes, that does not require further assumptions, which is also going to be useful to this paper. However, we want to find a more concrete characterisation, particularly with forbidden minors, which we will do in Section 3.3.

Next, we prepare to give the details for Remark 3.1.4 from the introduction.

Lemma 3.2.33. *(Folklore) Let G be a planar graph without parallel edges and $\mathcal{C}$ be any set of triangles of G. Then any plane embedding of G is nested.*

Proof. Since any two cycles which are triangles cannot intersect internally, the result immediately follows from Definition 3.2.26. □

Proposition 3.2.34. *A 2-dimensional simplicial complex is outerspatial if and only if its 1-skeleton is planar.*

Proof. For the 'only if' implication consider an outerspatial 2-dimensional simplicial complex C. Since C is outerspatial, its cone $\widehat{C}$ is embeddable. The link graph $L(t)$ at the top t of the cone is equal to the 1-skeleton of C. By Lemma 3.2.24 and the fact that $\widehat{C}$ is embeddable, the graph $L(t)$ is planar. Since $L(t)$ is planar and is equal to the 1-skeleton of C, it follows that the 1-skeleton of C is planar.

For the 'if' implication, suppose that the 1-skeleton of C is planar. Since C is a simplicial complex, its 1-skeleton has no parallel edges, and all face boundaries are triangles. Thus, by Lemma 3.2.33 the 1-skeleton has a nested plane embedding. So, the 2-dimensional simplicial complex is outerspatial by Lemma 3.2.31. □

Remark 3.2.35. For 2-dimensional simplicial complexes, the problem of outerspatiality is simple, however, this changes when one allows larger faces. That is because Proposition 3.2.34 follows easily from Lemma 3.2.33 and Lemma 3.2.31 but for the 2-complex analogue of Proposition 3.2.34 we do not have an analogue of Lemma 3.2.33 that we can use so we cannot apply Lemma 3.2.31 as easily. Checking for outerspatiality becomes significantly more complex when moving from the class of 2-dimensional simplicial complexes to general 2-complexes. This is because the associated nested plane embedding problem becomes nontrivial. This motivates why we work with 2-complexes in the rest of the paper.

3.3 Core Lemma

In this section we are going to introduce the Core Lemma, which is the main ingredient in proving Theorem 3.1.1. We will prove it in the next section.

Definition 3.3.1. In this paper we will call a compact connected 2-dimensional topological manifold without boundary a *surface*.

Definition 3.3.2. A 2-complex is *aspherical* if its geometric realisation is a surface which is not homeomorphic to the 2-sphere.

Lemma 3.3.3. *Aspherical 2-complexes are not outerspatial.*

Proof. Consider an aspherical 2-complex C with a 1-skeleton G. Then the link graph at the top of the cone is equal to G. Since G is a triangulation of a surface of positive genus, it cannot be planar because it does not have the required Euler characteristic. But the link graphs of an embeddable 2-complexes are planar by Lemma 3.2.24. Hence, $\widehat{C}$ does not embed in $\mathbb{R}^3$ which means that C is not outerspatial. □

Lemma 3.3.4. *Contraction of non-loop edges preserves being outerspatial.*

Proof. Clearly being outerspherical is preserved by contracting non-loop edges. So, this follows from Lemma 3.2.31. □

Lemma 3.3.5. *All link graphs of an outerspatial 2-complex are outerplanar.*

Proof. Consider an outerspatial 2-complex C and a link graph $L(v)$ at an arbitrary vertex v of C. Let $\widehat{C}$ be the cone over C with a top t and let $\widehat{L(v)}$ be the link graph at v as a vertex of $\widehat{C}$.

Sublemma 3.3.6. $\widehat{L(v)}$ *is the* 1*-dimensional cone over* $L(v)$*.*

Proof. Recall that the link graph $L(v)$ has a vertex for each edge of C and two vertices in $L(v)$ are connected by an edge if they share a face in C. To build the cone $\widehat{C}$ from the 2-complex C, we add one new edge incident to v, namely tv, so $L(v)$ has one new vertex. For each edge uv incident to v in C we add one face tuv, incident to the edge tv. Therefore, for any vertex uv of $L(v)$ we add one edge between tv and uv. Thus, we showed that to obtain $\widehat{L(v)}$ from $L(v)$ we add one new vertex and connect it to all old vertices by an edge. So $\widehat{L(v)}$ is the cone over $L(v)$ as claimed. □

Since C is outerspatial, $\widehat{C}$ is embeddable in 3-space by definition. Therefore, by Lemma 3.2.24, we know that $\widehat{L(v)}$ is planar. By Sublemma 3.3.6, we know that $\widehat{L(v)}$ is the cone over $L(v)$. In other words, the cone over $L(v)$ is planar. This is precisely the definition of an outerplanar graph. Since v was arbitrary, we showed that all link graphs of C are outerplanar, as desired. □

Lemma 3.3.7. *(Core Lemma) A simple locally* 2*-connected* 2*-complex C is outerspatial if and only if it does not contain an aspherical 2-complex as a subcomplex, and it does not contain a path P such that the link graph at the vertex P of C/P is not outerplanar.*

3.4 Main techniques

In this section we are going to prove two lemmas that are needed for one of the implications of Lemma 3.3.7. At the end of the section we will show the proof of Lemma 3.3.7 given these two lemmas and consequently show the proof of Theorem 3.1.1 given Lemma 3.3.7 . We will start the section with a definition that will be used in the context of both lemmas.

Definition 3.4.1. We are going to call a 2-connected simple outerplanar graph a *bi-outerplanar* graph.

The class of bi-outerplanar graphs will be important for the proofs of these lemmas. Thus, before we get to proving them, we are going to need to explore bi-outerplanar graphs through definitions and a few small lemmas.

Theorem 3.4.2. *([20]) The set of excluded minors for the class of outerplanar graphs consists of the graphs K_4 and $K_{2,3}$.*

Observation 3.4.3. *The link graphs of an outerspatial 2-complex do not have K_4 or $K_{2,3}$ minors.*

Proof. The link graphs of an outerspatial 2-complex are all outerplanar by Lemma 3.3.5. The class of outerplanar graphs is characterised by its forbidden minors K_4 and $K_{2,3}$. From this, the conclusion follows. □

Lemma 3.4.4. *(Folklore) Every bi-outerplanar graph G has a unique Hamiltonian cycle which bounds its outer face.*

Proof. Consider the face that is adjacent to all vertices. By 2-connectedness, this face is bounded by a cycle. So, this cycle is Hamiltonian. □

Using the notions from the previous lemma we have the following definitions.

Definition 3.4.5. In a bi-outerplanar graph, we pick a cycle as in Lemma 3.4.4 and refer to it as the *boundary cycle*. We call an edge *diagonal* if it connects two non-consecutive edges of the boundary cycle.

Definition 3.4.6. A face of a 2-complex C is called *diagonal* if it is a diagonal edge in some of the link graphs of C.

Definition 3.4.7. A diagonal face of a 2-complex is called *perfectly diagonal* if it is diagonal in the link graphs of all vertices it contains.

Lemma 3.4.8. *Consider a simple 2-complex C. If the link graphs of the complexes C/P for paths P of C (possibly trivial) are all 2-connected outerplanar graphs, then all of its diagonal faces are perfectly diagonal.*

Lemma 3.4.9. *Suppose that the simple 2-complex C does not have an aspherical subcomplex. If the diagonal faces of C are all perfectly diagonal and the link graphs of the complexes C/P for paths P of C (possibly trivial) are all bi-outerplanar, then C is outerspatial.*

The next two subsections are devoted to proving these two lemmas.

3.4.1 Proof of Lemma 3.4.8

Definition 3.4.10. Let H_1 and H_2 be two graphs with a common vertex v and a bijection ι between the edges incident with v in H_1 and H_2. The *vertex sum* of H_1 and H_2 over v given ι is the graph obtained from the disjoint union of H_1 and H_2 by deleting v in both H_i and adding an edge between any pair $(v_1; v_2)$ of vertices $v_1 \in V(H_1)$ and $v_2 \in V(H_2)$ such that v_1v and v_2v are mapped to one another by ι.

Proof of Lemma 3.4.8. If there are no diagonal faces, the claim is vacuously true. Suppose that there is a diagonal face f bounded by the cycle with vertices $u, x, x_1, \ldots, x_n$ and edges $ux, xx_1, x_1x_2, \ldots, x_nu$, which is a chord in the link graph $L(u)$ at the vertex u. The edge $e_x = ux$ is a vertex in the link graphs $L(u)$ and $L(x)$ and is of the same degree in both. The degree d of e_x in $L(u)$ is at least three as it is the endvertex of a chord so the degree of e_x in $L(x)$ is also at least three, hence it is also the endvertex of a chord. We claim that the face f is a chord in $L(x)$. We will prove this using the following.

Sublemma 3.4.11. *If the face f is not a chord in $L(x)$, then the link graph $L(e_x)$ at the vertex e_x of the 2-complex C/e_x has a $K_{2,3}$ minor.*

Proof. Suppose for a contradiction that the face f is not a chord in $L(x)$. By [12, Observation 3.1.] the link graph $L(e_x)$ at the vertex e_x in the 2-complex C/e_x is equal to the vertex sum of the graphs $L(u)$ and $L(x)$ in C at their common vertex e_x. Denote H to be this vertex sum. We shall see how the edges incident to e_x in $L(u)$ and $L(x)$ get identified to obtain H. Recall that e_x is the endvertex of at least one chord in $L(x)$. Since the chord f in $L(u)$ is a non-chord in $L(x)$, by the pigeonhole principle one of the chords in $L(x)$ incident to e_x is a non-chord in $L(u)$, call that chord g. Let the other end of f in $L(u)$ be e'_x and the other end of g in $L(x)$ be e''_x. Since $L(x)$ is bi-outerplanar, there are two non-chord edges incident to e_x, one of which is g, let the other one be k.

Now, there is a path of length at least two from e'_x to e_x through g in $L(u)$ and a path of length at least one from e_x to e''_x through g in $L(u)$ and the same for k. There is a path of length at least one from e'_x to e_x through f in $L(u)$ and a path of length at least two from e_x to e''_x through f in $L(x)$. All the paths mentioned above are pairwise internally vertex disjoint. Therefore, in the vertex sum H there are three internally vertex disjoint paths between the same pair of endvertices containing f, g and k respectively each of length at least two. This yields a subdivision of $K_{2,3}$. So $L(e_x) = H$ has a subgraph that is a subdivision of $K_{2,3}$. This means that $L(e_x)$ has a $K_{2,3}$ minor as claimed. □

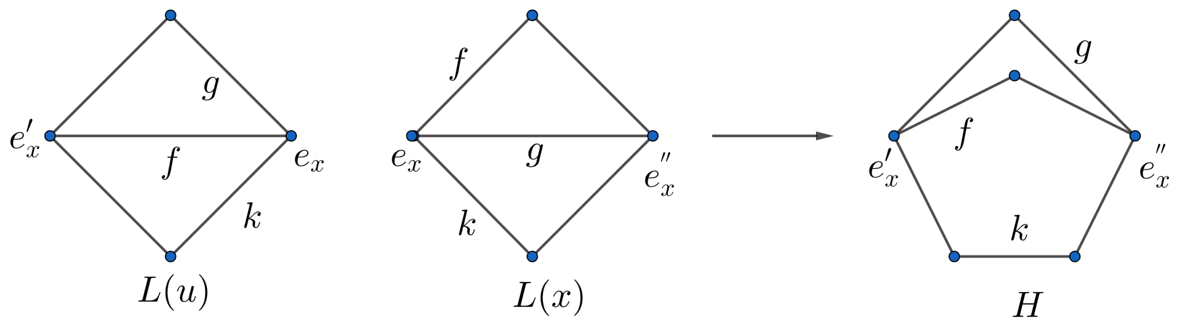

Figure 3.3: H is the vertex sum of $L(u)$ and $L(x)$ at the vertex e_x and has a $K_{2,3}$ minor.

If the graph $L(e_x)$ has a $K_{2,3}$ minor, then C/e_x has a link graph that is not outerplanar. Using Observation 3.4.3, this means that C/e_x is not outerspatial. Now, from Lemma 3.3.4, we conclude that C is also not outerspatial, which is a contradcition with our assumptions. This yields that f is a chord in $L(x)$. Similarly, looking at the link graphs $L(x)$ and $L(x_1)$, we obtain that f is a chord in $L(x_1)$. Repeating this argument inductively, we obtain that f is a chord in each $L(u)$, $L(x)$ and $L(x_i)$, $1 \leq i \leq n$. Therefore, f is a chord in the link graph of all of its endvertices. Since f was arbitrary, we proved that every diagonal face is perfectly diagonal as claimed. □

3.4.2 Proof of Lemma 3.4.9

Proof of Lemma 3.4.9. If C has no diagonal faces, then all link graphs are cycles, since they are all bi-outerplanar. This means that the geometric realisation of C is homeomorphic to a surface. By the assumption of the lemma, in such a case the geometric realisation of C must be homeomorphic to a sphere. A sphere is outerspatial and homeomorphism preserves outerspatiality, so we are done in this case.

Now, suppose that C has a diagonal face and consider one such face f. The link graph at each endvertex is bi-outerplanar and f is a chord in each of these link graphs. Therefore, if we remove f from the 2-complex C, we only remove chords from link graphs of C and they stay bi-outerplanar. Thus, we can remove the diagonal faces of C one by one to arrive at a 2-complex D whose link graphs are all cycles. As seen above, the geometric realisation of D must be homeomorphic to a sphere. Because of that and the fact that the 1-skeleton of D is naturally embedded in D, we can view the 1-skeleton of D as a plane graph.

Sublemma 3.4.12. *The 1-skeleton of D together with the boundaries of the removed diagonal faces form a nested plane embedding.*

Proof. Suppose for a contradiction, that there are two diagonal faces f_1 and f_2 with face boundaries c_1 and c_2 which are not nested. By Theorem 3.2.22, c_1 divides D into two connected components. Since c_1 and c_2 are not nested, there are edges of c_2 in both connected components of $D - c_1$. Therefore, there exists a subpath of c_2 that starts in one of the connected components of $D - c_1$ and ends in the other. Choose a minimal such path and call it p. Contract the subpath p' of p that consists of p with the first and last edge removed. Let the complex obtained from this contraction be $D' = D/p'$. The path p' is a subpath of c_1 by minimality of p, and is obviously a subpath of c_2. Let $c_1' = c_1/p'$

and $c_2' = c_2/p'$.

In the 2-complex D' we find two consecutive edges of c_2', each in a different component of $D' - c_1'$. Call the two edges a_2 and b_2 and notice that p' is the vertex that these edges share. Let the two edges of c_1' incident with p' be a_1 and b_1. By assumption, the link graph $L(p')$ is bi-outerplanar and thus it is a cycle C together with a set of edges between the vertices of C by Lemma 3.4.4. In $L(p')$, the vertices a_1 and b_1 are connected by the edge f_1 and the vertices a_2 and b_2 are connected by the edge f_2. Since f_1 and f_2 are perfectly diagonal, they are chords in $L(p')$. The above shows that $L(p')$ contains a cycle together with two non-parallel chords as a subgraph. The latter has a K_4 minor and so $L(p')$ also has a K_4 minor. This means that the 2-complex D/p' for the path p' has a non-outerplanar link graph at the vertex p'. The link graph $L(p')$ in D/p' is a subgraph of the link graph $L(p')$ in C/p'. Therefore, C/p' also has a non-outerplanar link graph at the vertex p'. This is a contradiction with the assumptions of the lemma which tells us that the boundaries of any two removed diagonal faces are nested, from which the sublemma follows. □

Let $\mathcal{C}_1$ denote the set of boundaries of the faces of D. Let $\mathcal{C}_2$ denote the set of boundaries of the removed diagonal faces. Then the set $\mathcal{C}$ of boundaries of the faces of C satisfies $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2$. Sublemma 3.4.12 gives us that the 1-skeleton of D together with $\mathcal{C}_2$ form a nested plane embedding. The 2-complexes C and D have the same 1-skeleton and the elements of $\mathcal{C}_1$ are nested with any other element in $\mathcal{C}$, because they are boundaries of faces on the sphere D. These two facts and Sublemma 3.4.12 together give us that the 1-skeleton of C, together with the boundaries of the faces of C form a nested plane embedding. From this it follows that C is outerspatial by Lemma 3.2.31 as claimed, which finishes the proof. □

3.4.3 Proof of Lemma 3.3.7 and Theorem 3.1.1.

Proof of Lemma 3.3.7: For the 'only if' direction, assume that the 2-complex C has a subcomplex that is an aspherical 2-complex. As being outerspatial is closed under deletion of faces, it follows from Lemma 3.3.3 that the 2-complex C cannot be outerspatial.

Next, assume that the 2-complex C contains a path P such that the link graph at the vertex P of C/P is not outerplanar. Since $L(P)$ is not outerplanar, it follows from Lemma 3.3.5 that C/P is not outerspatial. Then by Lemma 3.3.4 it follows that C is also not outerspatial.

We proved that, if a 2-complex C contains an aspherical subcomplex or a path P such that the link graph at the vertex P of C/P is not outerplanar, then C is not outerspatial. This proves the 'only if' direction as required.

For the 'if' implication, consider a locally 2-connected 2-complex C which does not contain an aspherical 2-complex as a subcomplex and does not contain a path P such that the link graph at the vertex P of C/P is not outerplanar. Then we can first apply Lemma 3.4.8 to obtain that all of its diagonal faces are perfectly diagonal. Next, since the result of Lemma 3.4.8 completes the assumptions of Lemma 3.4.9, we can apply the latter to obtain the final result. □

Proof of Theorem 3.1.1: Theorem 3.1.1 states that a simple locally 2-connected 2-complex C is outerspatial if and only if it does not contain a surface of positive genus as a subcomplex or a 2-complex with non-outerplanar link graph as a space minor. Lemma 3.3.7 states that a simple locally 2-connected 2-complex C is outerspatial if and only if it does not contain an aspherical 2-complex as a subcomplex and it does not contain a path P such that the link graph at the vertex P of C/P is not outerplanar.

Looking at both of these statements we can see that they are of the form 'a locally 2-connected 2-complex C is outerspatial if and only if it does not contain some forbidden structures'. So, to prove the implication we need to prove that the forbidden structures in Lemma 3.3.7 are a subset of the forbidden structures in Theorem 3.1.1.

Aspherical 2-complexes are homeomorphic to surfaces of positive genus. Since contracting a path is a space minor operation, if a 2-complex C contains a path P such that C/P has a non-outerplanar link graph, then it has a space minor with a non-outerplanar link graph.

The previous paragraph proves that the set of forbidden structures in Theorem 3.1.1 contains the set of forbidden structures in Lemma 3.3.7 and thus Lemma 3.3.7 implies Theorem 3.1.1 as claimed. □

3.5 Further remarks regarding outerspatial locally 2-connected 2-complexes

We start this section by recalling the definition of an outerspherical topological space.

Definition 3.5.1. Take a graph embedded in a 2-sphere S embedded as a 2-dimensional unit sphere in the Euclidean space $\mathbb{R}^3$. Glue discs to cycles of this graph inside the sphere so that they do not intersect each other in interior points, and they intersect the sphere S precisely in the gluing cycles. Call a topological space that can be obtained in this way *outerspherical*.

Now we elaborate a bit further on this definition.

Definition 3.5.2. We call the discs glued to faces of the 1-skeleton of the sphere *outer* discs and all other discs we call *inner* discs.

Definition 3.5.3. The *closure* of an outerspherical topological space T is obtained by adding all missing outer discs.

Definition 3.5.4. We will call an outerspherical topological space *maximal* if it is equal to its closure.

Remark 3.5.5. An outerspherical topological space induces a unique outerspatial 2-complex by taking the 1-skeleton of the 2-complex to be the graph of the topological space and the faces of the 2-complex to be the discs glued to cycles of this graph. Let a 2-complex induced in such a way by an outerspherical topological space T be denoted by $C(T)$.

Lemma 3.5.6. *If T is maximal, then $C(T)$ is locally 2-connected.*

Proof. A 2-complex C, which can be embedded as a maximal outerspherical topological space is a subdivision of a sphere together with some additional faces. Since the link graphs of a subdivided sphere are cycles, which are 2-connected, it follows that C is locally 2-connected. □

Lemma 3.5.7. *A simple outerspatial 2-complex C is locally 2-connected if and only if it has an embedding that is a maximal outerspherical topological space.*

Proof. For the 'if' direction, let T_C be an embedding of C that is a maximal outerspherical topological space. The result follows from Lemma 3.5.6 and the fact that $C = C(T_C)$.

For the 'only if' direction, consider an outerspatial 2-complex C with an outerspherical embedding T and note that $C = C(T)$. Let the closure of T be denoted by $\overline{T}$. Suppose for a contradiction that T is not maximal, but C is locally 2-connected. This means that $\overline{T} - T$ contains an outer disc, let one such disc be f. We have that f is a face in $C(\overline{T})$. Let v be one of the vertices of f and let $L(v)$ be the link graph of v with respect to $C(\overline{T})$, we know by Lemma 3.3.5 that $L(v)$ is outerplanar. A rotator at v induces a Hamiltonian cycle H in $L(v)$ which is unique and bounds the outer face of $L(v)$ by Lemma 3.4.4. In the graph $L(v)$, let $f = ab$ and suppose that there are two paths P_1 and P_2 between a and b different from ab. The vertices of both of these paths are all vertices of H due to the fact that H is Hamiltonian. Since H bounds the outer face of $L(v)$, we have that all edges of P_1 are on the inside of H, thus the path P_2 lies on the inside of the face bounded by the cycle $P_1 \cup f$, which leads to P_1 and P_2 intersecting internally. Therefore, the connectivity between a and b in the graph $L(v) - ab$ is 1 and so $L(v) - ab$ is not 2-connected. Hence, $C(\overline{T} - f)$ is not locally 2-connected and consequently $C(T)$ is also not locally 2-connected. Since $C = C(T)$, it follows that C is not locally 2-connected. This is a contradiction with our assumption, which shows that if any outerspherical embedding T of C is not maximal, then C cannot be locally 2-connected. This is the contrapositive of the 'only if' direction which finishes the proof. □

Definition 3.5.8. The *dual graph* of an outerspherical topological space is constructed in the following way. We have a vertex for each chamber apart from the outer chamber and two vertices are connected by an edge for each disc their respective chambers share.

Lemma 3.5.9. *The dual graph of a maximal outerspherical topological space T is a tree.*

Proof. We prove this by induction on the number of inner discs of T. When the number of inner discs is zero, we have a sphere, so the dual graph is a vertex which is a tree and thus the base case is true. Consider a maximal outerspherical topological space T and suppose that all maximal outerspherical topological spaces with less inner discs have dual graphs that are trees.

Let G be the dual graph of T and suppose that we remove some inner disc e. Then the dual graph of $T - e$ is G/e. Furthermore, since e is an inner disc of T, we have that $T - e$ is also maximal outerspherical topological space. Hence, by the inductive hypothesis G/e is a tree. Therefore, G is also a tree, which completes the inductive step and thus completes the proof. □

Proposition 3.5.10. *Let C be a locally 2-connected simple outerspatial 2-complex. Then the dual graph of an embedding of C is a tree.*

Proof. Firstly note that the term dual graph of an embedding of C is well-defined by Lemma 3.2.31. From here, the result follows from Lemma 3.5.7 and Lemma 3.5.9. □

Corollary 3.5.11. *A simple outerspatial locally 2-connected 2-complex can be constructed by starting from a sphere and then gluing a sequence of spheres one by one at an already existing face.*

Proposition 3.5.12. *Every locally 2-connected simple 2-complex has a unique embedding up to combinatorial equivalence.*

Proof. By Corollary 3.5.11, such a 2-complex can be built by gluing a sequence of spheres at some faces. If we remove all the gluing faces we obtain a subdivided sphere. There is a unique way to embed the sphere and then we can embed back the gluing faces uniquely on the 1-skeleton of this sphere. This shows that there exists a unique embedding of C (up to combinatorial equivalence). □

Proposition 3.5.13. *Every n-vertex locally 2-connected simple outerspatial 2-complex has at most $3n - 6$ edges and at most $3n - 8$ faces.*

Proof. Consider a locally 2-connected simple outerspatial 2-complex with n vertices. That it has at most $3n - 6$ edges follows from the fact that its 1-skeleton is planar and from Euler's formula.

We will prove by induction on the number of spheres glued in Corollary 3.5.11 that there are at most $3n - 8$ faces.

For one sphere there are n vertices and $2n - 4$ faces. Since $n \geq 4$, we have that $3n - 8 \geq 2n - 4$ so the base case is true. Suppose that we have a 2-complex with n vertices at most $3n - 8$ faces and we glue a sphere with m vertices and $2m - 4$ faces at some face. The new number of vertices is $n + m - 3$ and the new number of faces is at most $3n - 8 + 2m - 4 - 1 = 3n + 2m - 13$. Since $m \geq 4$, we have that $3n + 2m - 13 \leq 3n + 3m - 17 = 3(n + m - 3) - 8$. This completes the inductive step and thus we proved that an n-vertex 2-complex has at most $3n - 8$ faces. □

Lemma 3.5.14. *Consider a 2-complex C that is locally 2-connected. Then the cone over it is locally 3-connected.*

Proof. For any vertex in C, its link graph in $\widehat{C}$ is the (1-skeleton of the) cone over its link graph in C. We know that if G is 2-connected, then (the 1-skeleton of) its cone is 3-connected, so the link graphs at the vertices in C are 3-connected. Suppose that the link graph $L(t)$ of the top of the cone t has a 2-separator $\{tu, tv\}$. Then, restricting to C, consider the link graph $L(u)$. If $uv \in E(C)$, then v is a cutvertex in $L(u)$. If not, then $L(u)$ is disconnected. In either case, we have a contradiction with C being locally

2-connected. If $L(t)$ has a 1-separator tu, then $L(u)$ is disconnected when restricted to C. This is again contradiction with C being locally 2-connected. If $L(t)$ is disconnected, then so is C which is contradiction to C being simply-connected. We proved that $L(t)$ has no 0-, 1- or 2-separators. Thus, the link graph at the top is also 3-connected. All link graphs of $\widehat{C}$ are 3-connected, therefore the cone over C is locally 3-connected as required. □

Proposition 3.5.15. *Let C be a locally 2-connected 2-complex with F faces. Then, there exists an algorithm that checks in time linear in F whether C is outerspatial.*

Proof. Given the locally 2-connected 2-complex C we can construct its cone $\widehat{C}$ in a linear time. $\widehat{C}$ is locally 3-connected by Lemma 3.5.14. The methods of [12] give an algorithm that checks in linear time whether a locally 3-connected 2-complex is embeddable. Given that algorithm we can check whether $\widehat{C}$ is embeddable in linear time. Since C is outerspatial if and only if $\widehat{C}$ is embeddable, this gives a linear algorithm that checks whether C is outerspatial. □

3.6 Concluding remarks

The scope of this project was to prove Theorem 3.1.1 with the given assumptions. The result that we think could be eventually proved is as follows.

Conjecture 3.6.1. *A 2-complex is outerspatial if and only if it has no minor in the set Z.*

For a Z to be determined (the potentially necessary elements of Z will be discussed below).

To go from Theorem 3.1.1 to Conjecture 3.6.1 we need to go through the following steps.

1. Remove the requirement that the link graphs are simple in the locally 2-connected definition.

2. Remove the locally 2-connected requirement altogether.

3. Remove the requirement that the 2-complex is simple.

4. Tidy up the final forbidden space minor set and present it in a digestible manner.

To deal with point 1 we need to find a way to remove parallel cyclic faces without breaking the 2-outerplanar structure of the link graphs of the 2-complex the same way we removed chordal faces. In [12], we have local connectedness assumptions, as we have for Theorem 3.1.1. Removing the assumptions on Theorem 3.1.1 would be a similar endeavor to what is done in [17], which is proving the main result in [12] without the local connectedness assumptions. Such an idea could be used to deal with point 2. For point 3 we would need to develop machinery to deal with loops and parallel edges in the 1-skeleton. A potential strategy for the last point would be to redefine certain sets of forbidden/necessary substructures into a more succinct form similarly to the characterisation of locally planar 2-complexes in terms of 'strict marked minors' in [12, Section 5].

From the proof of Lemma 3.3.7 we can extract the following very insightful corollary.

Corollary 3.6.2. *Consider a graph G with a set of cycles $\mathcal{C}$ such that the* 2*-complex with* 1*-skeleton G and faces the set $\mathcal{C}$ is locally* 2*-connected. Then, G has a nested plane embedding if and only if this complex satisfies the conditions from Theorem 3.1.1.*

Chapter 4

Annulus graphs in $\mathbb{R}^d$

4.1 Introduction

In this paper, we are interested in graphs constructed as follows. Fix two real numbers $R_1, R_2 \geq 0$ satisfying $R_2 \geq R_1$. The family of *d-dimensional annulus graphs with radii R_1 and R_2*, denoted by $\mathcal{A}_d(R_1, R_2)$, consists of the graphs G whose vertex set can be embedded in $\mathbb{R}^d$ so that, for all pairs of different vertices $u, v \in V(G)$, uv is an edge of G if and only if the distance between u and v in the embedding is in the interval $[R_1, R_2]$. We call any such embedding an (R_1, R_2)-*annulus embedding*, or just an annulus embedding, of the graph G. In the sequel annulus graphs will be assumed finite unless explicitly defined as infinite.

The motivation for studying annulus graphs is twofold. To begin with, the notion interpolates between two classical models in graph theory: unit disc graphs and unit distance graphs. The family of d-dimensional unit disc graphs coincides with the family $\mathcal{A}_d(0, R)$ for every $R > 0$. Unit disc graphs were introduced in 1971 by Gilbert [34] to model telecommunication networks. Since then, the most significant developments of the theory of unit disc graphs were made in the framework of random unit disc graph also known as random geometric graphs, see Penrose [54] for a detailed account. The model was generalized by Waxman [68]: he worked in a setting where two vertices in positions

x and y are connected with probability $\beta \exp(-|x-y|/r)$ where β and r are parameters of the model. Penrose [55] introduced and studied a percolated version of the model.

Another line of research was initiated in 1946 by Erdős who was interested in the largest number of edges in a graph whose vertices may be embedded in $\mathbb{R}^2$ so that two vertices are at a distance 1 in the embedding if and only if they form an edge in the graph. Erdős [25] himself was able to show a lower bound of the form $n^{1+c/\log\log n}$ for some constant $c > 0$ and offered a 500 dollar prize for a proof whether or not there is an upper bound of the same form. To our knowledge, the best currently known upper bound is proportional to $n^{4/3}$ and was provided by Spencer, Szemerédi and Trotter [64]. Another related problem is the Hadwiger-Nelson problem asking for the smallest number of colours in which the points of the plane may be coloured so that no two points at a distance 1 are monochromatic. It has long been known that the answer is between 4 and 7, and de Grey [21] improved the lower bound to 5 (result reproved independently by Exoo and Ismailescu [26]).

Another major motivation of the project is the (rather general) Goldilock's principle, which roughly states that "objects that interact with each other should be neither too close nor too far from each other". The principle appears in many areas such as cognitive science [43] (infants prefer to occupy themselves with tasks that are neither too complex nor too simple), astronomy (the habitable zone around a star must be neither too close nor too far from it), economy (balancing high economic growth with low inflation) and machine learning [10] (concerning the learning rate of an algorithm) among others.

Despite the interest attracted by the subject in many fields of science, to our knowledge rigorous mathematical analysis of annulus graphs was conducted only by Galhotra, Mazumdar, Pal and Saha [32]. There, the authors studied the threshold of connectivity of the d-dimensional random annulus graph $\mathcal{G}(a(\log n)^{1/d}, b(\log n)^{1/d})$ with $a, b = O(1)$, which is obtained as follows. Consider a d-dimensional cube of side length $n^{1/d}$ and embed n vertices uniformly at random. Then, connect two vertices if the distance between them is in the interval $[a(\log n)^{1/d}, b(\log n)^{1/d}]$. They show that there is a function $\varphi : \mathbb{N} \to \mathbb{R}_+$

such that, if $b^d - a^d < \varphi(d)$, then the graph is disconnected whp, while if $b^d - a^d > \varphi(d)$, then the graph is connected whp. A wider class of random intersection graphs with more general connection functions was studied in [22].

4.1.1 Our results

While previous research concentrates on random annulus graphs, our results are deterministic in nature. Before presenting our first theorem, let us observe that, for any $d \in \mathbb{N}$ and $R_2 \geq R_1 > 0$, $\mathcal{A}_d(0,1) \neq \mathcal{A}_d(R_1, R_2)$. To see this, note that the family $\mathcal{A}_d(0,1)$ contains the family of all finite complete graphs while the clique number of every graph in $\mathcal{A}_d(R_1, R_2)$ is bounded.[1] Another easy remark is that, for all $d \in \mathbb{N}$ and $R_2 \geq R_1 > 0$, $\mathcal{A}_d(R_1, R_2) = \mathcal{A}_d(1, R_2/R_1)$. The next result points in the direction of distinguishing the families $\mathcal{A}_d(1, R)$.

Theorem 4.1.1. *For every $d \in \mathbb{N}$ there is a constant $C = C(d) > 0$ such that, for every pair of distinct real numbers $x, y \geq C$, $\mathcal{A}_d(1, x) \not\subseteq \mathcal{A}_d(1, y)$.*

To push our investigation of the families $\mathcal{A}_d(1, R)$ further, we study the question whether or not these families are χ-bounded. A family of graphs $\mathcal{F}$ is *χ-bounded* if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that, for every graph $G \in \mathcal{F}, \chi(G) \leq f(\omega(G))$, where $\chi(G)$ denotes the chromatic number of G and $\omega(G)$ denotes the size of the largest clique in G. Identifying χ-bounded classes has become a hot trend in recent years, for a complete account we direct the reader to the outstanding survey of Scott and Seymour [62] on the subject. Particular attention was paid to *intersection graphs*: given a collection $\mathcal{F}$ of sets, the intersection graph of $\mathcal{F}$ has vertex set $\mathcal{F}$ and edge set $\{(X, Y) : X, Y \in \mathcal{F}, X \cap Y \neq \emptyset\}$. Asplund and Grünbaum [6] showed that the family of intersection graphs of axis-parallel rectangles in the plane $\mathbb{R}^2$ is χ-bounded. Surprisingly, this is not the case for 3-dimensional

[1] Indeed, consider an (R_1, R_2)-*annulus embedding* of a complete graph in $\mathbb{R}^d$: then, the vertices of the graph must be embedded at the centers of disjoint balls of radii $R_1/2$ which themselves are contained in a ball of radius R_2 (and center any of the embedded vertices, say). Thus, the smaller balls of radii $R_1/2$ pack a larger ball of radius $R_2 + R_1/2$, and consequently there are at most $\left(\frac{R_2+R_1/2}{R_1/2}\right)^d$ smaller balls.

boxes as observed by Burling [11] - he provided an explicit construction of a sequence of intersection graphs of boxes with a bounded clique number and chromatic number that tends to infinity. Intersection graphs of discs were shown to form a χ-bounded class [9, 37] (in particular, for every unit disc graph G one has $\chi(G) \leq 6\omega(G) - 5$). Moreover, for every unit disc graph G, Peeters [53] showed that $\chi(G) \leq 3\omega(G) - 2$. His proof is based on an algorithm that, given an embedding of the vertex set of G witnessing that G is a unit disc graph, sweeps the points from bottom to top and greedily attributes a colour to a vertex in the moment when it is met. The proof of the upper bound in the next theorem uses a similar idea although the fact that an annulus is not a convex shape in general leads to some complications and requires additional ideas. For the lower bound, we construct a concrete embedding of a graph based upon a discretisation of the unit sphere $\mathbb{S}^{d-1}$.

Theorem 4.1.2. *There exist constants $M, m > 1$ such that, for every $d \geq 2$:*

(i) for every $x \geq 1$ we have

$$\sup_{G \in \mathcal{A}_d(1,x)} \frac{\chi(G)}{\omega(G)} \leq M^d;$$

(ii) for every $x \geq 1.2$ we have

$$m^d \leq \sup_{G \in \mathcal{A}_d(1,x)} \frac{\chi(G)}{\omega(G)}.$$

Both bounds hold for unit disc graphs as well. Moreover, for $d = 1$ and every $x > 1$ one may ensure that $3/2 \leq \sup_{G \in \mathcal{A}_1(1,x)} \chi(G)/\omega(G) \leq M$.[1]

Outline of the proofs. In the proof of Theorem 4.1.1 we consider a particular graph defined via a $(1, x)$-annulus embedding in $\mathbb{R}^d$ and show that it does not admit a $(1, y)$-annulus embedding. The proof is by contradiction and is divided into two cases: $d \geq 2$ and $d = 1$. In the more substantial case $d \geq 2$, we first show that "most pairs of vertices"

[1]The graphs in $\mathcal{A}_1(0, 1)$ are also called unit interval graphs. It is well-known that the clique number and the chromatic number of interval graphs are equal, see e.g. Section 5.5 in [23]. Also, the connected unit distance graphs in $\mathbb{R}$ are paths so the same conclusion holds for them as well.

at a distance at most 1 in the $(1, x)$-annulus embedding must be at a distance at most 1 in every $(1, y)$-annulus embedding of the graph. Then, we show that "most pairs of vertices" at a distance more than x in the $(1, x)$-annulus embedding must be at a distance more than y in every $(1, y)$-annulus embedding of the graph. On the basis of these results we distinguish the cases $x < y$ and $x > y$ and reach a contradiction in each of them.

The proof of the upper bound of Theorem 4.1.2 is based on the analysis of a geometric exploration algorithm. On the other hand, the proof of the lower bound is primarily based on an upper bound on maximal packings of the unit sphere with spherical caps.

Plan of the paper. In Section 4.2 we prove Theorem 4.1.1. In Section 4.3 we prove Theorem 4.1.2. Section 4.4 is dedicated to a related discussion.

4.2 Proof of Theorem 4.1.1

For any real numbers $x \geq 1$ and $\varepsilon > 0$, denote by $G_{\infty,d}(x, \varepsilon)$ the infinite graph that admits the following $(1, x)$-annulus embedding: its vertex set is embedded at the points $\{(\varepsilon n_i)_{i=1}^d : n_1, n_2, \dots, n_d \in \mathbb{Z}\}$ (a set which we denote $\varepsilon\mathbb{Z}^d$ in the sequel) and its edge set consists of the pairs $\{\{v_1, v_2\} : |v_1 - v_2|_d \in [1, x]\}$ (here, $|\cdot|_d$ denotes the Euclidean distance). We call this the *natural* embedding of $G_{\infty,d}(x, \varepsilon)$. We will show that for all sufficiently large x, y satisfying $x \neq y$ one may choose a small enough $\varepsilon > 0$ so that $G_{\infty,d}(x, \varepsilon)$ contains a finite subgraph outside $\mathcal{A}_d(1, y)$. Note that the case $d = 1$ requires certain modifications, and is therefore treated in a simpler way in the end of the section. For now, we assume that $d \geq 2$. In the sequel, we simplify the notation $|\cdot|_d$ to $|\cdot|$, and denote by $\mathrm{vol}(\cdot)$ the volume function in $\mathbb{R}^d$.

For every $n \in \mathbb{N}$, denote by $G_{n,d}(x, \varepsilon)$ the graph defined by restricting the natural embedding of $G_{\infty,d}(x, \varepsilon)$ to the ball of center $(0, 0)$ and radius n. Fix sufficiently large x and y and a sufficiently small $\varepsilon > 0$ and suppose for contradiction that for all $n \in \mathbb{N}$ one has $G_{n,d}(x, \varepsilon) \in \mathcal{A}_d(1, y)$. We start with a couple of preliminary results. The first of them says that, roughly speaking, most pairs of vertices of $G_{n,d}(x, \varepsilon)$ at a distance at most 1 in the natural embedding are at a distance at most 1 in any $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$.

Lemma 4.2.1. *Fix a sufficiently large $C_0 = C_0(d) > 0$, $x, y \geq C_0$ satisfying $x \neq y$, an integer $n \geq x$ and any sufficiently small $\varepsilon > 0$. Then, for every pair of vertices v_1, v_2 of $G_{n,d}(x, \varepsilon)$ whose images p_1, p_2 in the natural embedding satisfy $|p_1 - p_2| \leq 1$ and $\max\{|p_1|, |p_2|\} \leq n - x$, the images of v_1 and v_2 in every $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ in $\mathbb{R}^d$ are at a distance at most 1.*

We will need some preparation before presenting the proof of Lemma 4.2.1.

Lemma 4.2.2. *Fix $C_1 > 0$, $x > 2C_1 + 1$, any $\gamma \in ((C_1 + 1)/x, 1)$ sufficiently close to 1, an integer $n \geq x$ and any sufficiently small $\varepsilon > 0$. Also, fix a vertex v of $G_{n,d}(x, \varepsilon)$ such that its image p in the natural embedding satisfies $|p| \leq n - x$, and fix a point q on the boundary of $B(p, \gamma x)$. Denote by S the set of points of $\varepsilon\mathbb{Z}^d$ in the intersection of $B(p, \gamma x)$ and $B(q, C_1)$. Then, for every $c_1 > 0$ there is $c_2 = c_2(d, \gamma, c_1) > 0$ such that every subset of S of size $c_1|S|$ contains a complete graph on $\lfloor c_2 C_1^d \rfloor$ vertices.*

Proof. Denote $\mathcal{R} = B(p, \gamma x) \cap B(q, C_1)$. Since $\gamma x > C_1 + 1$, $\mathcal{R}$ contains a cone $\mathcal{C}$ with center q, radius C_1 and aperture $2\pi/3$. Also, note that $\mathcal{R}$ is disjoint from $B(p, 1)$ as seen in Figure 4.1. Thus, for every sufficiently small ε the number of points in $\mathcal{R}$ is at least $cC_1^d \varepsilon^{-d}$, where $c = c(d)$ is any constant smaller than the ratio of $\mathrm{vol}(\mathcal{C})$ and $\mathrm{vol}(B(0, C_1))$. Note that c is chosen independently of C_1.

Now, fix any $c_1 \in (0, 1]$ and let $\widehat{S} \subset S$ be a subset of S of size at least $c_1|S|$ points. Note that, since the diameter of $\mathcal{R}$ is bounded from above by $2C_1 \leq x$, every point in $\widehat{S}$ must be connected to all points in $\widehat{S}$ but at most $(\varepsilon^{-1} + 1)^d \mathrm{vol}(B(0, 1))$, which is at most $(2\varepsilon^{-1})^d \mathrm{vol}(B(0, 1))$ for all $\varepsilon \in (0, 1/2]$. Thus, for all such ε one may construct greedily a set of at least

$$\frac{|\widehat{S}|}{(2\varepsilon^{-1})^d \mathrm{vol}(B(0, 1))} \geq \frac{c_1 c C_1^d}{2^d \mathrm{vol}(B(0, 1))}$$

vertices in $\widehat{S}$ which form a clique. This shows the lemma for $c_2 = c_1 c/(2^d \,\mathrm{vol}(B(0, 1))$. □

In the sequel, for every $\gamma \in (0,1)$ denote by $N_\gamma \geq 1$ the maximal number of points in $B(0,\gamma)$ that are pairwise at a distance strictly greater than 1.

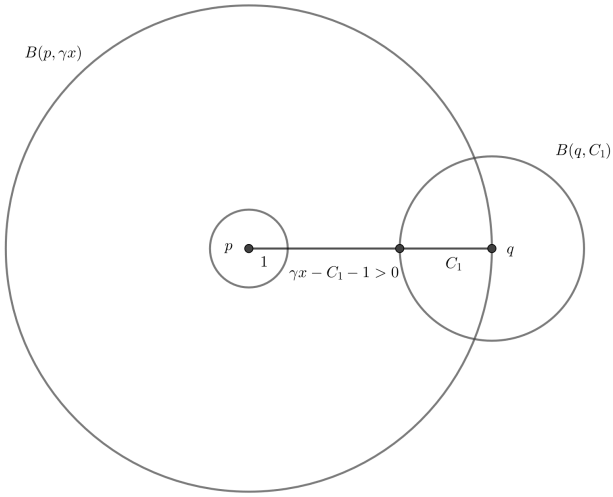

Figure 4.1: The configuration from the proof of Lemma 4.2.2.

Lemma 4.2.3. *There is a sufficiently large constant $C_2 > 0$ such that the following holds. Fix $x \geq C_2$, an integer $n \geq x$, $\gamma \in (0,1)$ sufficiently close to 1 and any sufficiently small $\varepsilon > 0$. Also, fix a vertex v of $G_{n,d}(x,\varepsilon)$ such that its image p_v in the natural embedding satisfies $|p_v| \leq n-x$. Then, for every $(1,y)$-annulus embedding of $G_{n,d}(x,\varepsilon)$ there exist N_γ vertices in the neighbourhood of v in the graph $G_{n,d}(x,\varepsilon)$ which are pairwise at a distance more than y in this embedding and are pairwise at a distance more than x in the natural embedding of $G_{n,d}(x,\varepsilon)$.*

Proof. First, notice that one may cover a ball with radius 1 with a finite number of $M = M(d)$ balls of radius $1/2$. Moreover, if the images of two vertices of $G_{n,d}(x,\varepsilon)$ in any $(1,y)$-annulus embedding of this graph are in a ball of radius $1/2$, they are not connected by an edge. Hence, the vertices in no ball of radius 1 in any $(1,y)$-annulus embedding of $G_{n,d}(x,\varepsilon)$ contain a clique of size $M+1$.

Next, consider a set of points $(q_i)_{i\in[N_\gamma]}$ in $B(p_v,\gamma x)$ for which

$$\min_{i,j\in[N_\gamma],i\neq j} |q_i - q_j|$$

is maximised. Since N_γ is the maximal number of points in a ball of radius γx that are pairwise at a distance strictly greater than x, we know that this maximum is more than x, and even that by choosing $\gamma \in (0,1)$ sufficiently close to 1 one may ensure that the distance between any two points is at least $(2-\gamma)x$ (indeed, this way one may ensure that $2-\gamma$ is arbitrarily close to 1).

Now, choosing C_2 so that $(1-\gamma)C_2 > 2\sqrt{C_2}$ means that any points in two different balls among $(B(q_i, \sqrt{C_2}))_{i=1}^{N_\gamma}$ are at a distance more than x. Also, for all $i \in [N_\gamma]$, denote by S_i the set of vertices with image in $\varepsilon\mathbb{Z}^2 \cap (B(p_v, \gamma x) \cap B(q_i, \sqrt{C_2}))$ in the natural embedding. Note that each S_i is disjoint from $B(p_v, 1)$ as seen in Figure 4.2. Now, choosing $c_1 = N_\gamma^{-1}$ and $c_2 = c_2(d, \gamma, c_1)$ as in Lemma 4.2.2, choosing $C_2 > c_2^{-1}(M+1)$ and using Lemma 4.2.2 for any $i \in [N_\gamma]$ and any subset $\widehat{S}_i \subset S_i$ satisfying $|\widehat{S}_i| \geq |S_i|/N_\gamma$ gives that there are $M+1$ vertices in $\widehat{S}_i$ that induce a complete graph. Fix any vertex u with image $p_u \in \bigcup_{i=1}^{N_\gamma}(B(p_v, \gamma x) \setminus B(p_v, 1)) \cap B(q_i, \sqrt{C_2})$ in the natural embedding of $G_{n,d}(x, \varepsilon)$. By the preceding observation, for any $i \in [N_\gamma]$ and in any $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$, $B(p_u, 1)$ contains less than $|S_i|/N_\gamma$ vertices of S_i.

Finally, fix any $(1, y)$-annulus embedding of the neighbourhood of v. Now, for all $j \in [N_\gamma - 1]$, choosing any vertices $v_1, \ldots, v_j$ from $S_1, \ldots, S_j$ ensures that there remain at least $(1 - j/N_\gamma)|S_{j+1}|$ vertices in S_{j+1}, which are at a distance at least y from each of $v_1, \ldots, v_j$. This allows us to choose greedily a set of N_γ vertices, one from each of $(S_j)_{j=1}^{N_\gamma}$, whose images in the $(1, y)$-annulus embedding are pairwise at a distance more than y, which completes the proof. □

Lemma 4.2.4. *Consider two points $a, b \in \mathbb{R}^d$ which are at a distance more than* 1. *Then, if $d \geq 3$ (respectively $d = 2$), for every sufficiently large $\gamma \in (0,1)$ there is no set of N_γ (respectively* 3*) points in $B(a, 1) \cap B(b, 1)$ that are pairwise at a distance larger than* 1.

Proof. Let S be a subset of points in $B(a, 1) \cap B(b, 1)$ such that all pairs of points in S are at a distance strictly greater than 1. Suppose first that $d \geq 3$ (see Figure 4.3). Then, there is $\gamma_1 \in (0,1)$ sufficiently close to 1 such that, up to slight modifications of the positions of the points in S if necessary, one may assume that $S \subseteq B(a, \gamma_1) \cap B(b, \gamma_1)$. Set

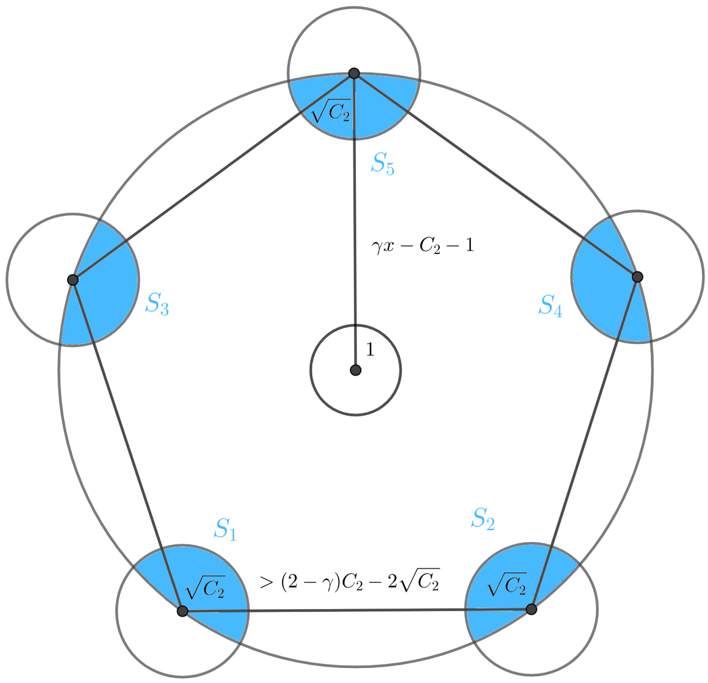

Figure 4.2: The configuration from the proof of Lemma 4.2.3.

$\gamma = (1+\gamma_1)/2$. Then, on the one hand, $B(a, \gamma_1) \subseteq B(a, \gamma)$, and moreover the intersection point of the ray $ba^{\rightarrow}$ and the sphere $\partial B(a, \gamma)$, which is further from b, is at a distance at least $\gamma + (1 - \gamma_1) > 1$ from S. Hence, $B(a, \gamma)$ contains a point at a distance more than 1 from every point in S, and by definition of N_γ we conclude that $N_\gamma > |S|$, which concludes the proof when $d \geq 3$.

Now, suppose that $d = 2$ (see Figure 4.4). Let the circles with centers a and b intersect at points c and c', and let the segment ab intersect these circles at points e and f. Then, for each point p in the curved triangle cef (where ef is a segment and ce and cf are arcs of the unit circles), the ball $B(p, 1)$ covers the curved triangle cef. Hence, there can be at most one point in S in the curved triangle cef. Similarly, there can be at most one point in S in the curved triangle $c'ef$. Since $B(a, 1) \cap B(b, 1)$ is covered by these two triangles, there can be at most two points at a distance larger than 1 there, which completes the proof. □

We are ready to come back to the proof of Lemma 4.2.1.

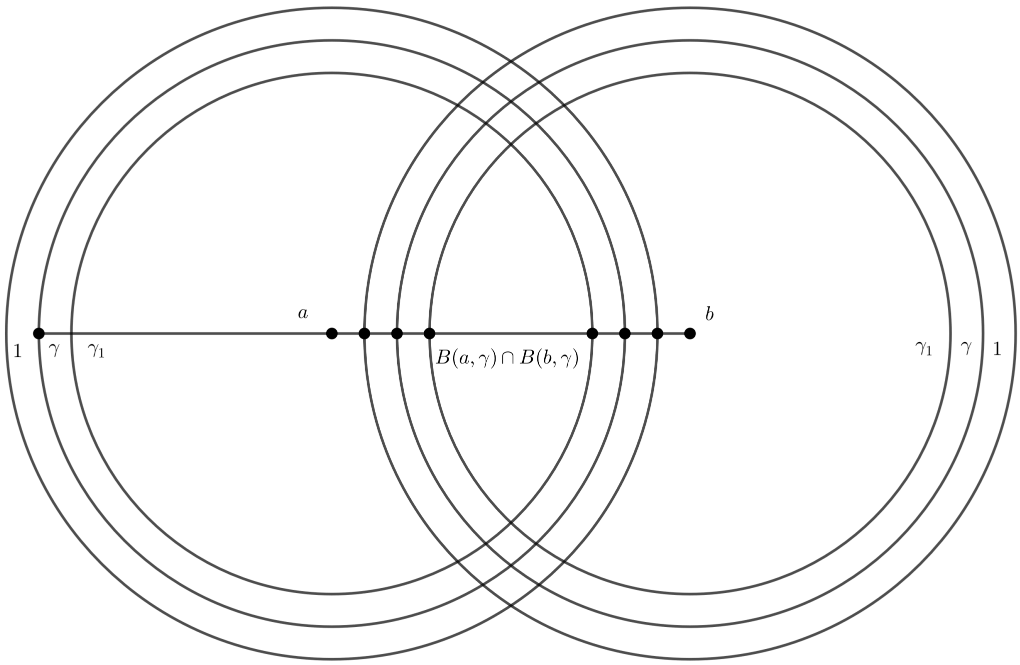

Figure 4.3: The configuration from the proof of Lemma 4.2.4 for $d \geq 3$ (depicted in two dimensions for simplicity); γ_1, γ and 1 indicate the radii of the spheres in the figure.

Proof of Lemma 4.2.1. Fix $\gamma \in (0, 1)$ as in Lemma 4.2.4 and $C_0 > (1-\gamma)^{-1}$. This ensures that $B(p_2, x)$ contains $B(p_1, \gamma x)$ (indeed, every point at a distance at most γx to p_1 is at a distance no more than $\gamma x + 1 < x$ from p_2). Thus, up to choosing ε sufficiently small and C sufficiently large, by Lemma 4.2.3 in any $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ there are N_γ common neighbours to v_1 and v_2 that are embedded into points, which are pairwise at a distance more than y.

Now, assume for contradiction that in some $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ the vertices v_1 and v_2 are embedded at a distance more than y. Then, the intersection of the balls with radius y, centered at the images of v_1 and v_2 in the $(1, y)$-annulus embedding (which we denote by q_1 and q_2), cannot contain N_γ points at pairwise distances more than y by Lemma 4.2.4, which finishes the proof. □

We showed that, roughly speaking, "most pairs of vertices that are close" in the natural embedding of $G_n(x, \varepsilon)$ are still "close" in any $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$. A second step is to show that "most" pairs of vertices that are far in the natural embedding of $G_{n,d}(x, \varepsilon)$ are far in any $(1, y)$-annulus embedding of $G_{n.d}(x, \varepsilon)$.

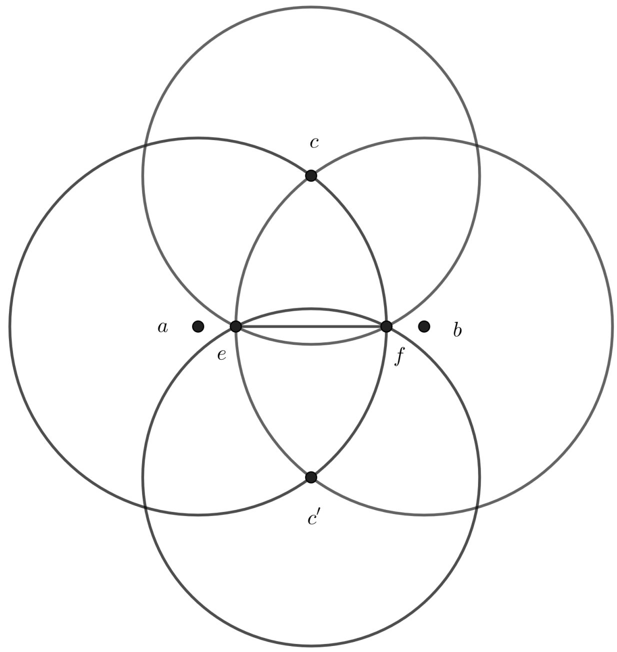

Figure 4.4: The configuration from the proof of Lemma 4.2.4 for $d = 2$.

Lemma 4.2.5. *Fix a sufficiently large $C_3 > C_0$ (with C_0 defined in Lemma 4.2.1), $x, y \geq C_3$ satisfying $x \neq y$, an integer $n \geq x+1$ and any sufficiently small $\varepsilon > 0$. Then, for every pair of vertices v_1, v_2 of $G_{n,d}(x, \varepsilon)$ whose images p_1, p_2 in the natural embedding satisfy $|p_1 - p_2| > x$ and $\max\{|p_1|, |p_2|\} \leq n - x - 1$, the images of v_1 and v_2 in any $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ in $\mathbb{R}^d$ are at a distance more than y.*

Before proceeding with the proof, we need to prepare the ground with one definition and two technical results.

Definition 4.2.6. The *unit half-ball* centered at 0 is obtained by taking the unit ball centered at 0 and removing the points with negative first coordinate.

Lemma 4.2.7. *For every $\gamma \in [0.99, 1)$, we have $N_\gamma \geq 5$ if $d = 2$ and $N_\gamma \geq 2d + 2$ if $d \geq 3$.*

Proof. If $d = 2$, inscribing a regular pentagon in the sphere $\partial B(0, 0.99)$ is sufficient since the side of this pentagon is $0.99 \cdot 2\sin(\pi/5) > 1$. Suppose that $d \geq 3$. Then, consider the set S of the points

$$a_1 = (x, y, 0, \ldots, 0),\ a_2 = (x, -y, 0, \ldots, 0),\ a_3 = (-x, y, 0, \ldots, 0),\ a_4 = (-x, -y, 0, \ldots, 0),$$
$$b_1 = (x, 0, y, \ldots, 0),\ b_2 = (x, 0, -y, \ldots, 0),\ b_3 = (-x, 0, y, \ldots, 0),\ b_4 = (-x, 0, -y, \ldots, 0),$$
$$\forall i \in [4, d], z_{i,1} = -z_{i,2} = (0, \ldots, 0, 1, 0, \ldots, 0),$$

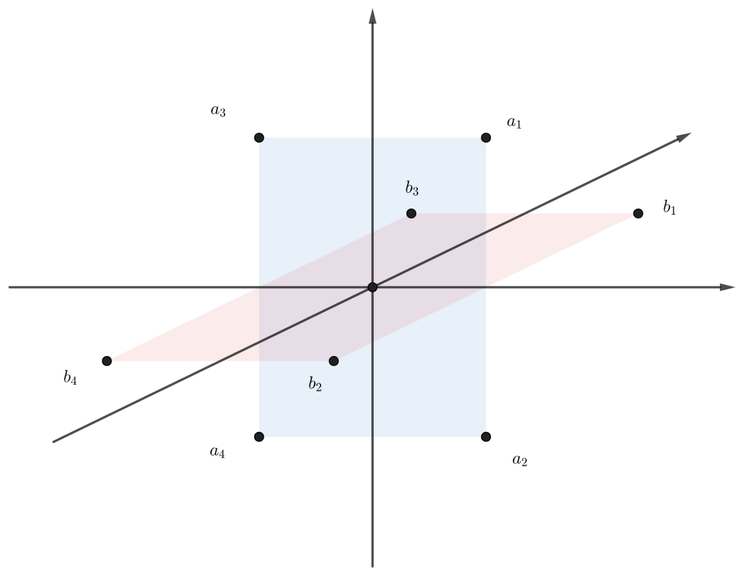

Figure 4.5: The configuration from the proof of Lemma 4.2.7.

where the unique 1 in $z_{i,1}$ is in the i-th position. We choose positive x and y so that $x^2 + y^2 = 0.99^2$ and $x = 0.6$. Then, $y = \sqrt{0.99^2 - 0.6^2} \approx 0.79 > x$ and so

$$\min_{i,j \in [4], i \neq j} |a_i a_j| = 2|x| = 1.2 \text{ and } \min_{i,j \in [4]} |a_i b_j| = \sqrt{2y^2} \approx 1.24 > 1.$$

Moreover, since for all $i \in [4, d]$ and $j \in \{1, 2\}$, the minimum to the distances from $z_{i,j}$ to any other point in S is more than 1, we conclude that $N_\gamma \geq |S| = 2d + 2$, which proves the lemma. □

The second technical result we need appears as Theorem 2 in [50] (which is formulated in terms of a graph parameter called *sphericity* of the complete bipartite graph $K_{d,d}$); see also Section 5 in [56] for the related notion of spherical dimension.

Lemma 4.2.8 (see Theorem 2 in [50])**.** *For every $d \geq 1$, there is no set of $2d+2$ points $\{p_1, \ldots, p_{2d+2}\}$ in $\mathbb{R}^d$ such that:*

- $|p_i - p_j| > 1$ *for all i, j such that either both $i, j \geq d+2$ or both $i, j \leq d+1$;*
- $|p_i - p_j| \leq 1$ *for all i, j such that $1 \leq i \leq d+1$ and $d+2 < j \leq 2d+2$.* □

We are ready to prove Lemma 4.2.5.

Proof of Lemma 4.2.5. Choosing $C_3 > \max\{C_0, 3\}$, where C_0 was chosen in Lemma 4.2.1, we have that $B(p_1, 1) \cap B(p_2, 1) = \emptyset$ and moreover one may find unit half-balls $B^1_{1/2} \subset B(p_1, 1)$ and $B^2_{1/2} \subset B(p_2, 1)$ satisfying $B^1_{1/2} \cap B(p_2, x) = \emptyset$ and $B^2_{1/2} \cap B(p_1, x) = \emptyset$. Also, choose $\gamma \in [0.99, 1)$ as in Lemma 4.2.4, and fix a set S of N_γ points at a distance strictly greater than 1 in $B(0,1) \setminus \{0\}$. Now, for ε sufficiently small there is an injective map $\phi : B(0,1) \to B^1_{1/2} \cup B^2_{1/2}$ for which $\phi(S)$ consists of $|S|$ points at a distance strictly greater than 1, from which $k := \lfloor N_\gamma/2 \rfloor$ are in $B^1_{1/2}$ and $N_\gamma - k$ are in $B^2_{1/2}$. Such a map indeed exists: for example, consider the hyperplane orthogonal to the vector $(1, 0, \ldots, 0) \in \mathbb{R}^d$, and then rotate it in the plane, generated by the first two coordinates. By discrete continuity of the difference of the number of points in S on the two sides there is a moment when:

- there in no point lying on the rotating hyperplane, and
- the difference between the numbers of points in S on the two sides is at most 1.

At this point, it is sufficient to "split" $B(0,1)$ into two halves and map them to $B^1_{1/2}$ and to $B^2_{1/2}$ in the natural way (the choice of an image for the points on the hyperplane itself is arbitrary).

Now, we argue by contradiction. Suppose that in some $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ the images of v_1 and v_2, which we denote by q_{v_1} and q_{v_2}, respectively, are at a distance at most 1. Denote by $(q_i)_{i=1}^k$ (respectively $(q_i)_{i=k+1}^{N_\gamma}$) the images of the vertices, corresponding to the points in $\phi(S) \cap B^1_{1/2}$ (respectively in $\phi(S) \cap B^2_{1/2}$) in the given $(1, y)$-annulus embedding. Then, since $B(q_{v_1}, 1) \cup B(q_{v_2}, 1)$ has diameter at most $3 \leq y$, we have that

- $|q_i - q_j| > 1$ for all i, j such that either both $i, j \geq k + 1$ or both $i, j \leq k$;
- $|q_i - q_j| \leq 1$ for all i, j such that $1 \leq i \leq k$ and $k + 1 \leq j \leq N_\gamma$.

If $d \geq 3$, then $k \geq d + 1$ by Lemma 4.2.7 and therefore such set of points does not exist in $\mathbb{R}^d$ by Lemma 4.2.8, which is a contradiction to our assumption. If $d = 2$, then $k \geq 2$ and $N_\gamma - k \geq 3$, which again leads to contradiction with the statement for $d = 2$ in Lemma 4.2.4. □

The next two corollaries show that we can apply Lemma 4.2.1 and Lemma 4.2.5 multiple times to obtain linearly multiplied version of them.

Corollary 4.2.9. *Fix $C_0 > 0$ as in Lemma 4.2.1, $x, y > C_0$ satisfying $x \neq y$ and integers $n \geq x$ and $k > 1$. Then, for every pair of vertices v_1, v_2 of $G_{n,d}(x, \varepsilon)$ whose images p_1, p_2 in the natural embedding satisfy $|p_1 - p_2| < k$ and $\max\{|p_1|, |p_2|\} \leq n - x - 1$, the images of v_1 and v_2 in every $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ in $\mathbb{R}^d$ are at a distance at most k for every sufficiently small $\varepsilon > 0$ (depending only on p_1, p_2 and k).*

Proof. Fix $\delta = k - |p_1 - p_2|$ and $\varepsilon \in (0, (5k)^{-1}\delta]$. Divide the segment between p_1 and p_2 into k segments $x_0x_1 = p_1x_1, x_1x_2, \dots, x_{k-2}x_{k-1}, x_{k-1}x_k = x_{k-1}p_2$ of equal length. Then, for all $i \in [k] \cup \{0\}$, associate to x_i a nearest vertex $x'_i \in \varepsilon\mathbb{Z}^d$ (so $x'_0 = x_0$ and $x'_k = x_k$). Clearly for all $i \in [k] \cup \{0\}$ one has $|x_i - x'_i| \leq 2\varepsilon$ and $|x'_i| \leq 2\varepsilon + x_i \leq 2\varepsilon + n - x - 1 < n - x$. By the triangle inequality $|x'_i - x'_{i-1}| \leq 1 - k^{-1}\delta + 2 \cdot 2\varepsilon < 1$, which means that $(x'_{i-1}x'_i)_{i\in[k]}$ are all of length at most 1, so we conclude by applying Lemma 4.2.1 k times. □

Corollary 4.2.10. *Fix $x, y > C_3$ satisfying $x \neq y$ and integers $n \geq x$ and $k > 1$. Then, for every pair of vertices v_1, v_2 of $G_{n,d}(x, \varepsilon)$ whose images p_1, p_2 in the natural embedding satisfy $|p_1 - p_2| < kx$, the images of v_1 and v_2 in every $(1, y)$-annulus embedding of*

$G_{n,d}(x, \varepsilon)$ in $\mathbb{R}^d$ are at a distance at most ky for every sufficiently small $\varepsilon > 0$ (depending only on p_1, p_2, x and k).

Proof. Consider two vertices v_1, v_2 of $G_{n,d}(x, \varepsilon)$ whose images p_1, p_2 in the natural embedding are at a distance at most x. If they are at a distance at most 1, then the images of v_1, v_2 in an $(1, y)$-annulus embedding of $G_{n,d}(x, \varepsilon)$ are also at a distance at most 1 by Lemma 4.2.1 (and the fact that $C_3 > C_0$, as stated in Lemma 4.2.5). If p_1 and p_2 are at a distance between 1 and x, then we have that v_1 and v_2 are connected by an edge in $G_{n,d}(x, \varepsilon)$. Therefore, in any $(1, y)$-annulus embedding they should be at a distance between 1 and y. We showed that if two vertices of $G_{n,d}(x, \varepsilon)$ have images in the natural embedding at a distance less then x, they should be at a distance less than y in any $(1, y)$-annulus embedding. To finish the proof we do the same argument as in Corollary 4.2.9 but with steps of length x and y instead of steps of length 1. □

Definition 4.2.11. For every $d \in \mathbb{N}$ and every pair of point sets $K, L \subseteq \mathbb{R}^d$, define $M_d(K, L)$ as the maximum number of disjoint congruent copies of L included in K, that is, the maximum size of an L-packing of K.

Proof of Theorem 4.1.1 for $d \geq 2$. We set $C = C_3$, where C_3 is as defined in Lemma 4.2.5, $x, y \geq C$ satisfying $x \neq y$, any sufficiently large n and any sufficiently small $\varepsilon > 0$ (these two last parameters will be specified later). We will prove that $G_{n,d}(x, \varepsilon)$ cannot be realised in any $(1, y)$-annulus embedding. Denote by $\mathcal{X}$ the natural embedding of $G_{n,d}(x, \varepsilon)$ and fix some $(1, y)$-annulus embedding $\mathcal{Y}$ of $G_{n,d}(x, \varepsilon)$. We split the proof over two cases.

Case 1: $x < y$

Let k be a positive integer. Consider a set of vertices V in $G_{n,d}(x, \varepsilon)$ with image V_x in $\mathcal{X}$, which is included in $B(0, k-1)$ and for every $p_1, p_2 \in V_x$ we have $|p_1 - p_2| > x$. For $n > k + x + 1$, by Lemma 4.2.5 the image V_y of V in $\mathcal{Y}$ is a set of points at a distance more than y from each other. Moreover, by Corollary 4.2.9 and by choosing ε sufficiently small, these points are contained in a closed ball of radius k.

Fix a sufficiently small $\delta > 0$ and consider a packing $\mathcal{B}$ of the ball $B(0, k - 1 + x/2)$ with balls of radius $x/2 + \delta$. By choosing a sufficiently small $\varepsilon = \varepsilon(\delta) > 0$, for each ball $B \in \mathcal{B}$ there is a point $u \in \mathcal{X}$ at a distance at most $\delta/2$ from the centre of B. Thus, from $\mathcal{B}$ one may construct a packing $\mathcal{B}_{\mathcal{X}}$ of $B(0, k-1+x/2)$ with balls of radius $x/2 + \delta/2$ and centres in $\mathcal{X}$. Consequently, one may choose a set V of size at least $M_d(B(0, k-1+x/2), B(0, x/2+\delta))$.

However, in $\mathcal{Y}$ the balls with radii $y/2$ and centers V_y pack the ball $B(0, k+y/2)$. Moreover, since $x < y$, by choosing $\delta \leq (y-x)/4$ we obtain that

$$M_d(B(0, k-1+x/2), B(0, (x+y)/4)) \leq M_d(B(0, k+y/2), B(0, y/2)) \tag{4.2.1}$$

for all $k \in \mathbb{N}$.

However, for all fixed $r > 0$ the limit of $M_d(B(0,k), B(0,r))\text{vol}(B(0,r))/\text{vol}(B(0,k))$ when $k \to +\infty$ exists and is given by the packing density of $\mathbb{R}^d$ with unit balls. Since

$$\lim_{k \to +\infty} \frac{\text{vol}(B(0, k-1+x/2))}{\text{vol}(B(0, k+y/2))} = 1,$$

we conclude that $\text{vol}(B(0, (x+y)/4)) \geq \text{vol}(B(0, y/2))$, which leads to a contradiction.

Case 2: $x > y$

Let k be a positive integer. Consider a set of vertices V in $G_{n,d}(x, \varepsilon)$ with image V_x in $\mathcal{X}$, which is included in $B(0, kx-1)$ and for every $p_1, p_2 \in V_x$ we have $|p_1 - p_2| > 1$. For $n > k+x+1$ Lemma 4.2.1 implies that the image V_y of V in the embedding $\mathcal{Y}$ is a set of points at a distance more than 1 from each other, and by Corollary 4.2.10 and by choosing ε sufficiently small these points are contained in a ball of radius ky.

In the exact same way as in the first case we find out that for every $\delta > 0$ one has

$$M_d(B(0, kx-1/2+\delta), B(0, 1/2+\delta)) \leq M_d(B(0, ky+1/2), B(0, 1/2)) \tag{4.2.2}$$

Hence,

$$\begin{aligned}&\lim_{k\to+\infty}\frac{M_d(B(0,kx-1/2+\delta),B(0,1/2+\delta))\mathrm{vol}(B(0,1/2+\delta))}{\mathrm{vol}(B(0,kx-1/2+\delta))}\\ \text{and}\quad&\lim_{k\to+\infty}\frac{M_d(B(0,ky+1/2),B(0,1/2))\mathrm{vol}(B(0,1/2))}{\mathrm{vol}(B(0,ky+1/2))}\end{aligned}\tag{4.2.3}$$

exist and are both equal to the packing density of $\mathbb{R}^d$ with unit balls. However,

$$\lim_{k\to+\infty}\frac{\mathrm{vol}(B(0,kx-1/2+\delta))}{\mathrm{vol}(B(0,ky+1/2))}=\left(\frac{x}{y}\right)^d,\tag{4.2.4}$$

so 4.2.2, 4.2.3 and 4.2.4 together imply that $x^{-d}\mathrm{vol}(B(0,1/2+\delta))\geq y^{-d}\mathrm{vol}(B(0,1/2))$, which leads to a contradiction by choosing δ sufficiently small. Thus, the proof of Theorem 4.1.1 is completed. □

It remains to deal with the case $d=1$. Although the main points of the proofs are the same, the proof is technically simpler in this case.

Proof of Theorem 4.1.1 for $d=1$. Again, consider the graph $G_{n,d}(x,\varepsilon)$.

Claim 4.2.12. *Lemma 4.2.1 holds for $d=1$ as well.*

Proof of Claim 4.2.12. We argue by contradiction. Let without loss of generality $p_1<p_2$. Suppose that in some $(1,y)$-embedding, the images q_1,q_2 of v_1,v_2, respectively, are at a distance more than y. Then, all common neighbours of v_1 and v_2 in $G_{n,1}(x,\varepsilon)$ have images in the $(1,y)$-embedding, contained in the intersection of the annuli around q_1 and q_2 with radii 1 and y, which is a segment of length less than y (denoted J in the sequel).

Now, denote by V_1 (resp. V_2) the common neighbours of v_1,v_2 in the interval $[p_2-x,p_1-x/2]$ (resp. $[p_2+x/2,p_1+x]$) in the natural $(1,x)$-embedding. On the one hand, since $p_2+x/2-(p_1-x/2)>x$, no vertex in V_1 is connected by an edge to a vertex in V_2. On the other hand, the images of the vertices in $V_1\cup V_2$ in the $(1,y)$-embedding are contained in J, and therefore, the image of every vertex in V_1 must be at a distance at most 1 from the image of every vertex in V_2. As a result, the distance between the images of any two vertices in $V_1 \cup V_2$ (in the $(1, y)$-embedding) must be at most 2, so all of them must be contained in some subinterval of J of length 2. By choosing $x \geq 10$, we have that $p_1 + x - (p_2 + x/2) \geq x/2 - 1 \geq 4$, which implies that V_2 contains a clique on four vertices. This contradicts the fact that all these vertices are contained in an interval of length 2 in the $(1, y)$-embedding, as desired. □

Figure 4.6: The configuration from the proof of Claim 4.2.12.

Claim 4.2.13. *Lemma 4.2.5 holds for $d = 1$ as well.*

Proof of Claim 4.2.13. We argue by contradiction. Let without loss of generality $p_1 < p_2$. Suppose that in some $(1, y)$-annulus embedding, the images q_1, q_2 of v_1, v_2, respectively, are at a distance less than 1. Then, if $y \geq 2$ and ε is sufficiently small, one may conclude by Claim 4.2.12 and an easy induction that for every vertex with position in the interval $(-\infty, p_1]$ in the $(1, x)$-annulus embedding, its position in the $(1, y)$-annulus embedding is at a distance at most 1 from q_2. At the same time, for all sufficiently large y and sufficiently small ε, this contradicts with the fact that the vertices with images in $(-\infty, p_1]$ in the $(1, x)$-annulus embedding induce a graph, containing a K_4. □

The remainder of the proof is analogous to the proof in the case $d \geq 2$. □

4.3 Proof of Theorem 4.1.2

4.3.1 Proof of the upper bound of Theorem 4.1.2

In the heart of the proof of the upper bound of Theorem 4.1.2 is the following algorithm, which colours the vertices of an annulus graph $G \in \mathcal{A}_d(r_1, r_2)$ properly (i.e. no two adjacent vertices share the same colour). First, given an annulus embedding of G, rotate it so that no two (images of) vertices in $V(G)$ have coinciding last coordinates. Then, fix an affine hyperplane orthogonal to the last coordinate axis of $\mathbb{R}^d$ and which is below the entire vertex set of G. Then, start moving this hyperplane continuously upwards. We colour the vertices of G in colours indexed by the positive integers. When (the image of a) vertex in G meets the moving hyperplane:

- if this vertex has already been coloured, do nothing;
- if this vertex has not been coloured before, consider the set of uncoloured vertices at a distance at most $r_1/2$ from it. Colour all of these vertices in the smallest colour which is still available for all of them.

At each step when the hyperplane meets a still uncoloured vertex v, the algorithm colours a set S_v of previously uncoloured vertices (of course, $v \in S_v$). For every vertex $u \in S_v$, call v the *token* of u and denote $v = t(u)$. Observe that, first, of all vertices coloured at the same moment the token has the smallest last coordinate, and second, vertices that have the same token also have the same colour. By construction the above algorithm produces a proper colouring of G, which we call c_G. Moreover, note that the particular case when $r_1 = 0$ coincides with Peeters' sweeping algorithm [53].

We proceed with a lemma that we will use in the proof of our theorem. In the sequel, we tacitly identify any (r_1, r_2)-annulus graph with an arbitrary (r_1, r_2)-annulus embedding of this graph in $\mathbb{R}^d$.

Lemma 4.3.1. *Consider a graph $G \in \mathcal{A}_d(r_1, r_2)$. For every vertex v in G, the ball $B(v, r_1)$ contains vertices in at most 7^d colours in c_G.*

Proof. If $r_1 = 0$, the statement is trivial. Assume that $r_1 > 0$. Note that each vertex $u \in B(v, r_1) \cap V(G)$ satisfies $|u - t(u)| \leq r_1/2$. Since u is in $B(v, r_1)$, we have that $t(u)$ is in $B(v, 3r_1/2)$. Moreover, a token t_2 coloured later than a token t_1 in the algorithm must be at a distance more than $r_1/2$ from t_1 since otherwise the vertex t_2 would itself have t_1 as a token. Hence, for each pair of tokens t_1 and t_2, the balls $B(t_1, r_1/4)$ and $B(t_2, r_1/4)$ do not intersect. We conclude that the number of tokens that may fit into $B(v, 3r_1/2)$ is at most $M_d(B(0, 7r_1/4), B(0, r_1/4)) \leq 7^d$. The last inequality holds because 7^d is the quotient of the volume of the two balls. Therefore, the ball $B(v, r_1)$ contains points in at most 7^d different colours. □

We also make use of the following result.

Theorem 4.3.2 (simplified version of Theorems 1.1 and 1.2 in [67], see also [58])**.** *Fix any $d \in \mathbb{N}$, $r, R > 0$ such that $R \geq r$, and let $T = R/r$. Let $\nu_{T,d}$ be the minimal number of closed balls of radius r which may cover a closed ball of radius R in $\mathbb{R}^d$. Then,*

$$1 \leq \nu_{T,d} \leq P(d)T^d,$$

where $P(d)$ is a polynomial function.

Proof of upper bound of Theorem 4.1.2. Consider $r_1 \geq 0$ and $r_2 > 0$ satisfying $r_2 \geq r_1$. Consider an (r_1, r_2)-annulus graph G together with the (proper) colouring c_G given by the colouring algorithm described above. Let $k = k(G)$ be the largest colour in c_G and t_k be a token with colour k. Note that the vertices which may forbid the colours in $[k-1]$ for t_k are at a distance at most r_2 from $B(t_k, r_1/2)$. Hence, all such vertices are contained in the lower half of $B(t_k, r_1/2 + r_2)$.

Since $r_1 \leq r_2$, by Theorem 4.3.2 the ball $B(t_k, r_1/2 + r_2)$ may be covered by at most $\nu_{T,d} \leq (3 + o_d(1))^d$ balls of radius $r_2/2$ where $T = 2 + r_1/r_2 \leq 3$. By the pigeonhole principle one of these balls of radius $r_2/2$, say B, contains vertices in at least $k/\nu_{T,d}$ colours of c_G. By Lemma 4.3.1 one may find a set $S \subseteq B \cap V(G)$ of at least $k/(\nu_{T,d}\, 7^d)$ vertices in B which are pairwise at a distance more than r_1. Thus, the distance between each pair of vertices in S is between r_1 and r_2, and hence S is a clique. Therefore,

$$\frac{k}{\nu_{T,d}\, 7^d} \le \omega(G). \tag{4.3.1}$$

Hence, $\chi(G) \le k \le (21 + o_d(1))^d \omega(G)$, where $o_d(1)$ signifies a quantity that goes to 0 as d goes to ∞. Thus, we have

$$M := \sup_{d\in\mathbb{N}} \sup_{r_1,r_2} \sup_{G\in\mathcal{A}_d(r_1,r_2)} \left(\frac{\chi(G)}{\omega(G)}\right)^{1/d} < +\infty,$$

which proves the theorem. Note that here we proved the result also for unit disc graphs. □

4.3.2 Proof of the lower bound

Before providing the proof of the lower bound of Theorem 4.1.2 in full generality, we show the case $d = 1$ as a warm-up. We recall that when $r_1 > 0$, we set $x = r_2/r_1$, and in particular $\mathcal{A}_d(r_1, r_2) = \mathcal{A}_d(1, x)$.

Proof of the lower bound of Theorem 4.1.2 for $d = 1$. Fix $x > 1$. We are going to show that $\mathcal{A}_1(1, x)$ contains a triangle-free graph with an odd cycle, which has chromatic number at least 3 and clique number 2. We consider two cases, illustrated in Figure 4.7 and Figure 4.8:

- if $x \ge 2$, then the following five points form an embedding of the 5–cycle with vertices, listed in consecutive order: $0, x, 2x, x + 0.99, x - 0.99$.
- if $x < 2$, then let $k \ge 2$ be the smallest integer such that $kx \ge k + 1$. The following $2k+1$ points form an embedding of a graph containing a $(2k+1)$-cycle with vertices, listed in consecutive order:

$$0, x, \dots, kx, (k - k(k+1)^{-1})x, (k - 2k(k+1)^{-1})x, \dots, (k - k^2(k+1)^{-1})x.$$

Clearly all pairs of consecutive points as well as the first and the last point form edges. At the same time, every graph in $\mathcal{A}_1(1, x)$ is triangle-free since the sum of two numbers in the interval $[1, x]$ is at least $2 > x$.

This completes the proof in the case $d = 1$. □

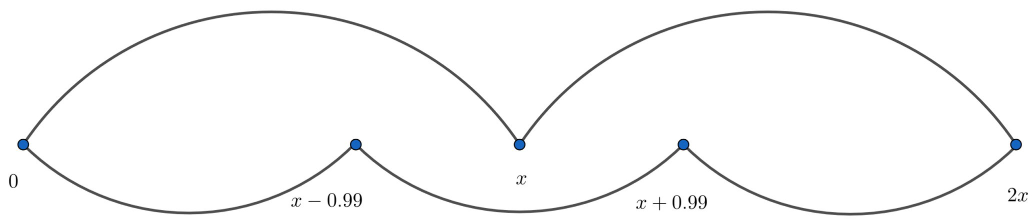

Figure 4.7: Illustration of the graph from the first case. Here, x was chosen to be 3.

Fix any $d \geq 2$ and $x \in [1.2, +\infty)$. We first provide a proof of the lower bound for the family $\mathcal{A}_d(1, x)$ and then come back to the case of unit disc graphs (that is, $\mathcal{A}_d(0, 1)$). The proof of the lower bound relies on a construction that approximates the infinite uncountable $(2/x, 2)$-annulus graph with vertex set $\mathbb{S}^{d-1}$. Our main goal is to provide an example of an annulus graph in $\mathcal{A}_d(1, x)$ with a large (multiplicative) gap between its chromatic number and its clique number based on discrete versions of the following two theorems.

For any $d \geq 2$ and a set $X \subseteq \mathbb{S}^{d-1}$, the *spherical diameter of* X, denoted ϕ_X, is the supremum over all pairs of points $x, y \in X$ of the spherical distance between x and y. The next theorem due to Schmidt [60, 61] relates the diameters of two compact sets X_1, X_2 with the maximum distance between a point in X_1 and a point in X_2.

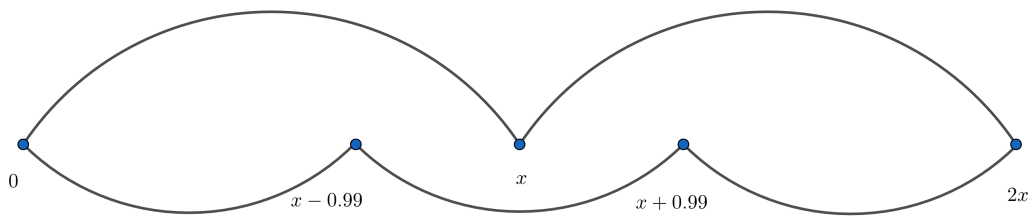

Figure 4.8: Illustration of the graph from the second case. Here, x was chosen to be $\frac{3}{2}$, which means $k = 2$.

Theorem 4.3.3 (isodiametric inequality, see Chapter II, Section 8 in [60])**.** *Let E be a metric space among $\mathbb{S}^{d-1}, \mathbb{R}^{d-1}$ and the $(d-1)$-dimensional hyperbolic space. Fix two compact sets $X_1, X_2 \subseteq E$ and define D as the supremum of the distance between x_1 and x_2 over all pairs of points $x_1 \in X_1, x_2 \in X_2$. Also, let d_1 (respectively d_2) be the diameter of a ball in the same space having volume equal to the volume of X_1 (respectively of X_2). Then, $d_1 + d_2 \le 2D$.*

We will use the following corollary of Theorem 4.3.3 that was proved independently by Böröczky and Sagmeister, see Theorem 1.2 in [7].

Corollary 4.3.4. *Let X be a measurable subset of $\mathbb{S}^{d-1}$ with spherical diameter $\phi_X < \pi$. Then, the Lebesgue measure of (the closure of) X is at most the volume of a spherical cap with spherical diameter ϕ_X.*

In the sequel, we denote by $\mathrm{Cap}^{d-1}(\phi)$ the spherical cap in $\mathbb{S}^{d-1}$ with center $(1, 0, \dots, 0)$ and diameter ϕ, and define $M(d, \phi)$ to be the maximum number of disjoint copies of $\mathrm{Cap}^{d-1}(\phi)$ that can be packed in $\mathbb{S}^{d-1}$, that is, $M(d, \phi) = M_d(\mathbb{S}^{d-1}, \mathrm{Cap}^{d-1}(\phi))$. The next theorem from [42] provides an upper bound on $M(d, \phi)$.

Theorem 4.3.5 ([42], see also Section 2.4 in [65]). *For all $\phi \in [0, \pi)$,*

$$\frac{1}{d} \ln M(d, \phi) \leq \frac{1 + \sin\phi}{2\sin\phi} \ln\left(\frac{1 + \sin\phi}{2\sin\phi}\right) - \frac{1 - \sin\phi}{2\sin\phi} \ln\left(\frac{1 - \sin\phi}{2\sin\phi}\right) + o_d(1).$$

Now, we outline the idea of the proof. Fix $x \in [1.2, +\infty)$. We will provide a graph $G \in \mathcal{A}_d(2/x, 2)$ for which $\chi(G)/\omega(G) \geq m^d$ for some constant $m > 1$ independent of d and x. For a graph G, we denote by $|G|$ the number of its vertices and by $\alpha(G)$ the size of its maximum independent set, that is, the size of the largest set of vertices inducing no edge. Recall that, for every non-empty graph G, $\chi(G) \geq |G|/\alpha(G)$. Therefore, we will provide a graph G satisfying $|G|/(\alpha(G)\omega(G)) \geq m^d$ and conclude by the previous observation.

We start by giving an upper bound on the clique number of any annulus graph on the sphere $\mathbb{S}^{d-1}$. Then, we find a set of points in $\mathbb{S}^{d-1}$ forming an annulus graph G_d with $|G_d|/\alpha(G_d)$ suitably bounded from below.

Lemma 4.3.6. *Consider a graph G with an $(2/x, 2)$-embedding in the sphere $\mathbb{S}^{d-1}$. Then, its clique number is at most $M(d, 2\arcsin(x^{-1}))$.*

Proof. Since the diameter of $\mathbb{S}^{d-1}$ is 2, a pair of vertices of G are connected by an edge if and only if they are at a distance at least $2/x$. Thus, a clique in G is a set of vertices which are pairwise at Euclidean distance at least $2/x$, and in particular the open balls with radii $1/x$ around these points are disjoint. Hence, the largest clique in such a graph has size at most $M(d, 2\arcsin(x^{-1}))$. □

For every integer $d \geq 2$, we define $\mathrm{svol}_{d-1}(X)$ as the $(d-1)$-dimensional Lebesgue measure on $\mathbb{S}^{d-1}$. The next lemma provides an upper bound for $M(d, 2\arcsin(x^{-1}))$.

Lemma 4.3.7. *For all $x \in [1.2, +\infty)$, we have*

$$M(d, 2\arcsin(x^{-1})) \leq \frac{(0.997 + o_d(1))^d \, \mathrm{svol}_{d-1}(\mathbb{S}^{d-1})}{\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\arcsin(x^{-1})))}$$

.

Proof. It is well-known (see e.g. page 67 in [48]) that for all $\theta \in (0, \pi/2]$,

$$\frac{\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\theta))}{\mathrm{svol}_{d-1}(\mathbb{S}^{d-1})} = \frac{\int_0^\theta (\sin t)^{d-2} dt}{\int_0^\pi (\sin t)^{d-2} dt} = (\sin\theta + o_d(1))^d. \tag{4.3.2}$$

Note that for the second equality in (4.3.2), we used Laplace's method to approximate $\int_0^b (\sin t)^{d-2} dt$ by $(\sin x_0 + o_d(1))^d$, where x_0 is the unique maximum of sin in $(0, b] \subseteq (0, \pi]$. Set $\theta = \arcsin(x^{-1})$. By Theorem 4.3.5, we have

$$(\sin\theta)^d M(d, 2\theta) \leq$$
$$\left(\sin\theta \exp\left(\frac{1+\sin(2\theta)}{2\sin(2\theta)} \ln\left(\frac{1+\sin(2\theta)}{2\sin(2\theta)}\right) - \frac{1-\sin(2\theta)}{2\sin(2\theta)} \ln\left(\frac{1-\sin(2\theta)}{2\sin(2\theta)}\right)\right) + o_d(1)\right)^d.$$

A study of the function of $\theta \in [0, \arcsin(1.2^{-1})]$

$$\theta \mapsto \sin\theta \exp\left(\frac{1+\sin(2\theta)}{2\sin(2\theta)} \ln\left(\frac{1+\sin(2\theta)}{2\sin(2\theta)}\right) - \frac{1-\sin(2\theta)}{2\sin(2\theta)} \ln\left(\frac{1-\sin(2\theta)}{2\sin(2\theta)}\right)\right)$$

shows that its maximum is attained at $\theta = \arcsin(1.2^{-1})$ and this maximum is smaller than 0.997. The reader may verify our claim by following the link in [1]. This shows that

$$\frac{\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(x^{-1}))\, M(d, 2\arcsin(x^{-1}))}{\mathrm{svol}_{d-1}(\mathbb{S}^{d-1})} \leq (0.997 + o_d(1))^d, \tag{4.3.3}$$

from which the lemma follows. □

A set $\mathcal{N}_\varepsilon \subseteq \mathbb{S}^{d-1}$ is called an ε-net if every point in $\mathbb{S}^{d-1}$ has a point in $\mathcal{N}_\varepsilon$ at spherical distance at most ε.

Lemma 4.3.8. *For any sufficiently small real number $\delta = \delta(x, d) > 0$, there is a $(2/x, 2)$-annulus graph G satisfying*

$$\frac{|G|}{\alpha(G)} \geq \frac{\mathrm{svol}_{d-1}(\mathbb{S}^{d-1})}{\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\arcsin(x^{-1}) + \delta))}.$$

Proof. We show that for every sufficiently small $\varepsilon = \varepsilon(\delta, d) > 0$, there is an ε-net $\mathcal{N}_\varepsilon$ in

$\mathbb{S}^{d-1}$ which is the vertex set of an annulus graph with the required property. Fix a small enough ε and consider a tessellation $\mathcal{T}$ of $\mathbb{S}^{d-1}$ into regions of spherical diameter at most ε and area at least $a = a(\delta, d, \varepsilon) > 0$. Consider a Poisson Point Process $\mathcal{P}$ with intensity $\lambda = \lambda(\delta, d, \varepsilon, a) > 0$ on the sphere satisfying that $\lambda a \geq (\log |\mathcal{T}|)^2$, where $|\mathcal{T}|$ stands for the number of regions in $\mathcal{T}$. Using well-known estimates for the tails of Poisson random variables (see e.g. Theorem A.1.15 in [2]) and the union bound, we get

$$\begin{aligned}
&\mathbb{P}(\exists R \in \mathcal{T} : |R \cap \mathcal{P}| - \mathbb{E}|R \cap \mathcal{P}| \notin (-\varepsilon\mathbb{E}|R \cap \mathcal{P}|, \varepsilon\mathbb{E}|R \cap \mathcal{P}|)) \\
\leq\ & |\mathcal{T}| \max_{R \in \mathcal{T}} \mathbb{P}(|R \cap \mathcal{P}| - \mathbb{E}|R \cap \mathcal{P}| \notin (-\varepsilon\mathbb{E}|R \cap \mathcal{P}|, \varepsilon\mathbb{E}|R \cap \mathcal{P}|)) \\
\leq\ & |\mathcal{T}| \exp(-\Omega_\varepsilon(\log |\mathcal{T}|)^2) = o_{|\mathcal{T}|}(1).
\end{aligned}$$

Thus, every region $R \in \mathcal{T}$ contains whp a number of points that is in the interval $((1 - \varepsilon)\mathbb{E}|R \cap \mathcal{P}|, (1 + \varepsilon)\mathbb{E}|R \cap \mathcal{P}|)$ and, in particular, at least one point. We condition on this event and set $\mathcal{N}_\varepsilon = \mathcal{P}$. Let G be the $(2/x, 2)$-annulus graph on vertex set $\mathcal{N}_\varepsilon$ and I be any maximum independent set of G. Also, let $\mathcal{T}_I$ be the union of all regions of $\mathcal{T}$ containing a vertex in I. Then, since every region in the tessellation $\mathcal{T}$ has spherical diameter at most ε, the spherical diameter of $\mathcal{T}_I$ is at most $\varepsilon + 2\arcsin(x^{-1}) + \varepsilon = 2(\arcsin(x^{-1}) + \varepsilon)$. Hence, by choosing $\varepsilon < \delta/2$ one may derive by Corollary 4.3.4 that $\mathcal{T}_I$ has $(d-1)$-dimensional volume at most $\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\arcsin(x^{-1}) + \delta/2))$. We conclude that every independent set of G contains at most

$$|\mathcal{N}_\varepsilon| \frac{(1+\varepsilon)\,\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\arcsin(x^{-1}) + \delta/2))}{(1-\varepsilon)\,\mathrm{svol}_{d-1}(\mathbb{S}^{d-1})}$$

points, which, up to choosing ε sufficiently small, is bounded from above by

$$|\mathcal{N}_\varepsilon| \frac{\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\arcsin(x^{-1}) + \delta))}{\mathrm{svol}_{d-1}(\mathbb{S}^{d-1})},$$

which proves the lemma. $\square$

Proof of Theorem 4.1.2 (ii). Fix a sufficiently small $\delta > 0$ and the corresponding graph G given in Lemma 4.3.8. Then,

$$\frac{\chi(G)}{\omega(G)} \geq \frac{|G|}{\alpha(G)\,\omega(G)} \geq \frac{\mathrm{svol}_{d-1}(\mathbb{S}^{d-1})}{\mathrm{svol}_{d-1}(\mathrm{Cap}^{d-1}(\arcsin(x^{-1})+\delta))\,M(d,2\arcsin(x^{-1}))}, \tag{4.3.4}$$

where the first inequality holds for any graph and the second inequality follows from Lemma 4.3.8. Choosing $\delta > 0$ small enough and using Lemma 4.3.7, we can ensure that the right hand side in (4.3.4) is bounded below by 1.003^d, which finishes the proof. □

It remains to deduce the result in the case of unit disc graphs. Fix $x \in [1.2, \infty]$. Consider a unit disc graph G which is constructed as the annulus graph with radii 0 and $2/x$ on a set of points lying on the sphere $\mathbb{S}^{d-1}$. Then, its complement G^c is a graph for which two vertices are connected if the distance between them is in the interval $(2/x, 2]$. Note that every largest clique in a graph G is a largest independent set in G^c (that is, the complement of G) and every largest independent set of G is a largest clique in G^c. Therefore,

$$\frac{\chi(G)}{\omega(G)} \geq \frac{|G|}{\alpha(G)\,\omega(G)} = \frac{|G^c|}{\alpha(G^c)\,\omega(G^c)}.$$

Moreover, the above proof of Theorem 4.1.2 works with minor modifications for (R_1, R_2)-annulus graphs for which two points are connected if they are at a distance $(R_1, R_2]$, which concludes the proof.

4.4 Discussion

We finish with a short discussion around the bounds on $\sup_{G\in\mathcal{A}_d(1,x)} \chi(G)/\omega(G)$ provided by Theorem 4.1.2. To begin with, showing that a similar lower bound holds for $x \in [1, 1.2)$ is an obvious open question. One construction similar to ours is to connect a point p (see it as "the north pole") with all points at a distance between x_1 and x_2 where $x_2/x_1 = x$ and $x_1^2 + x_2^2 = 4$ (this is a description of the points in $\mathbb{S}^{d-1}$ in a strip symmetric with respect to "the equator"). The missing piece to show that this more universal construction provides an exponential lower bound for every $x > 1$ is the following conjecture of ours, which bears close resemblance to Kalai's double-cap conjecture, see Conjecture 2.8 in [27].

Conjecture 4.4.1. *Fix some* $\theta \in [0, \pi/2)$*. Every measurable set* $X \subseteq \mathbb{S}^{d-1}$ *containing no two points at spherical distance in the interval* $[\pi/2 - \theta, \pi/2 + \theta]$ *has measure at most* $2|\mathrm{Cap}^{d-1}(\pi/4 - \theta/2)|$*, where the unique maximiser (up to rotation) is a pair of diametrically opposite caps with spherical diameters* $\pi/2 - \theta$*.*

In Lemma 4.3.7 we study the function $\theta \in [0, \arcsin(1.2^{-1})]$

$$\theta \mapsto \sin\theta \exp\left(\frac{1+\sin(2\theta)}{2\sin(2\theta)} \ln\left(\frac{1+\sin(2\theta)}{2\sin(2\theta)}\right) - \frac{1-\sin(2\theta)}{2\sin(2\theta)} \ln\left(\frac{1-\sin(2\theta)}{2\sin(2\theta)}\right)\right) \tag{4.4.1}$$

to obtain that it has its maximum at $\theta = arcsin(1.2^{-1})$ and this maximum is smaller than 0.997. The key takeaway here is that we chose $x = 1.2$ as the lower bound for which Theorem 4.1.2 works because we want the function in (4.4.1) to have a maximum smaller than 1 in the interval $[0, \arcsin(x^{-1})]$, so that the right-hand side of the final inequality in Lemma 4.3.7 is bounded by 1. As the maximum for $x = 1.2$ is less than 0.997, there exists an ε for which Theorem 4.1.2 works for $x \geq 1.2 - \varepsilon$, so this lower bound could be slightly improved.

We remark that an approach similar to ours (but providing a worse lower bound for x) may be applied when θ is "close" to $\pi/2$ but not when θ is "small". It is worth observing that when $x = 1$ (that is, in the case of the unit distance graphs), we have that the largest cliques have size $d+1$ while the chromatic number of a unit-distance graph may be larger than 1.2^d for all sufficiently large d by a result of Frankl and Wilson [30]. Nevertheless, a randomised construction embedding points on the d-dimensional sphere with radius slightly larger than $1/\sqrt{2}$ uniformly at random implies that $(1, x)$-annulus graphs may contain cliques with exponential (in d) number of vertices for all $x > 1$.

Theorem 4.1.2 shows that $\sup_{G \in \mathcal{A}_d(1,x)} \chi(G)/\omega(G)$ grows exponentially fast with d, which is somehow satisfactory for high dimensional annulus graphs. However, although it is easy to improve Lemma 4.3.1 by a factor of 2, our approach seems to provide an upper bound that is far from optimal for small values of d. In the particular case of $\mathcal{A}_2(0,1)$ a simplified version of our algorithm coincides with the one used by Peeters [53], so in particular every graph $G \in \mathcal{A}_2(0,1)$ is shown to satisfy $\chi(G) \leq 3\omega(G) - 2$. However, if $\omega(G) = 2$, we claim that $\chi(G) \leq 3$: indeed, for any embedding of G in $\mathbb{R}^2$ witnessing that $G \in \mathcal{A}_2(0,1)$, no two edges on four different vertices may intersect since otherwise by triangle inequality G must contain a triangle with three out of these four vertices. Thus, all triangle-free graphs in $\mathcal{A}_2(0,1)$ are planar, so also 3-colourable by a theorem of Grötzsch [38] – a bound attained by any cycle of odd length.

On the other hand, Malesińska, Piskorz and Weißenfels [51] showed that for every $\omega \in \mathbb{N}$ there is a graph $G \in \mathcal{A}_2(0,1)$ satisfying $\omega(G) = \omega$ and $\chi(G) \geq \lfloor 3\omega/2 \rfloor$. Closing the gap between the lower and the upper bounds is a long-standing open problem.

4.5 The set of annulus graphs in relation to their dimension

In this section we give two results about annulus graphs. The first result shows that for every $d > 0$, there exists a graph that is not a d-dimensional annulus graph. This tell us that the set of annulus graphs up to a fixed dimension is strictly growing. This gives us a crucial motivation for the use of dimension as a parameter of annulus graphs, because for each dimension there are graphs which are annulus graphs in this dimension but not for lower ones. The second result offers us in some sense the opposite, that if we fix a graph G, then there exists some dimension that can accommodate an annulus embedding of this graph. As a consequence, we have that all graphs are annulus graphs, given high enough dimension to embed them in.

Theorem 4.5.1. *(Jung's theorem) Given a compact set $K \subset R^n$ and let D be the diameter of K. Then there exists a closed ball of radius $r \leq D\sqrt{\frac{n}{2(n+1)}}$ that contains K.*

Theorem 4.5.2. *For a fixed $d > 0$ there exists a graph which is not a d-dimensional annulus graph for any pair of radii $R_1 \leq R_2$.*

Proof. To prove this, we will build a family of graphs $\{G_n \,|\, n \in \mathbb{N}\}$, such that for any fixed $d > 0$, if n is large enough, then G_n is not d-dimensional. For a given $n \in \mathbb{N}$ we build the graph G_n in the following way. Start with a copy of the complete graph on n vertices K_n and call this copy the *inner* block of G_n. Next, take n more copies of K_n and call them *outer blocks* of G_n. We will call each of the K_n that makes up G_n a *block*. Now, glue each of the n outer blocks at a unique vertex of the inner block. For a point A in the inner block, denote the outer block containing A by $\mathcal{B}_A$. The resulting graph is our desired G_n and has n^2 vertices and $\frac{1}{2}(n-1)n(n+1)$ edges. A picture of G_4 can be seen below.

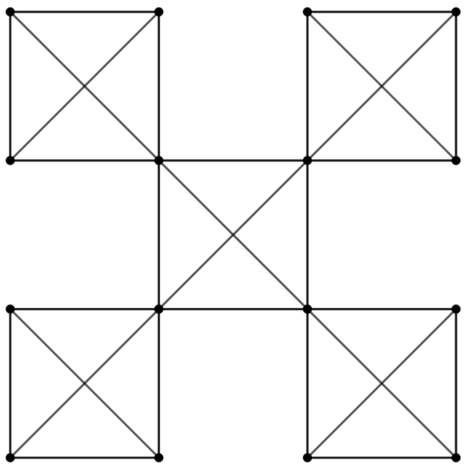

Figure 4.9: The K_4 in the center is the inner block and each of the four outer blocks share one vertex with it.

For a fixed $d > 0$ and a variable n, consider an annulus embedding of G_n in $\mathbb{R}^d$ with any radii $R_1 \leq R_2$. In this embedding the set of points in each of the $(n+1)$ blocks of G_n is compact and with diameter R_2. Hence, by Theorem 4.5.1, they should be in a ball of radius $R = R_2\sqrt{\frac{d}{2(d+1)}}$. Moreover each of the outer blocks share a vertex with the inner block of G_n, so the union of the inner block with an arbitrary outer block is contained in a ball of radius $2R$. Overall, the collection of all blocks and thus the whole embedding of G_n is contained in a ball of radius $3R$.

Sublemma 4.5.3. *Consider k points $A_1, A_2, \ldots, A_k$ in the inner block of G_n. Provided that n is large enough in terms of d and k, there exists a set of k points $a_i \in \mathcal{B}_{A_i} - A_i$,* *for $1 \leq i \leq k$, which are at a pairwise distance larger than R_2.*

Proof. We will prove this by induction on k. The base case $k = 1$ is trivial. Consider sets of points $A_1, A_2, \ldots, A_{k-1}$ and $a_1, a_2, \ldots, a_{k-1}$ satisfying the inductive hypothesis and consider a new point A_k in the inner block. Let the set of points in $\mathcal{B}_{A_k} - A_k$ which are at a distance less than R_1 from a_1 be denoted by $\mathcal{A}_1$. Since the points in $\mathcal{A}_1$ are in the same block they are at a distance at least R_1 from each other. Therefore, the points in $\mathcal{A}_1$ are contained in a ball of radius R_1 centered at a_1 and at a distance at least R_1 from each other. For a fixed d the size of such a set is bounded. We can also obtain that the similarly defined sets $\mathcal{A}_i$ for $1 \leq i \leq k-1$ are of bounded size. Hence, the set $\mathcal{A} = \bigcup_{i=1}^{k-1} \mathcal{A}_i$ is of bounded size for a fixed d. Thus, if n is large enough in terms of d, we can find a point $a_k \in \mathcal{B}_{A_k} - (A_k \cup \mathcal{A})$, which means that a_k is at a distance larger than R_1 from each a_i, for $1 \leq i \leq k-1$. If we add the fact that a_k is not connected by an edge to each a_i, for $1 \leq i \leq k-1$, since all k points are in different blocks, we get that a_k is at a distance larger than R_2 from each a_i, for $1 \leq i \leq k-1$. As such, we have found the desired a_k which completes the inductive step. □

From Sublemma 4.5.3 we obtain that for any k and fixed d, if $n = n(k, d)$ is large enough, there exist k points of G_n which are at a pairwise distance larger than R_2. Since R is linear in R_2 for a fixed d, there can be only finitely many points in the ball of radius $3R$, which are at a pairwise distance larger than R_2. Hence, for n large enough, we get that G_n cannot be embedded in $\mathbb{R}^d$ because all points in G_n are contained in a ball of radius $3R$. Since n depends only on d and R_1 and R_2 can vary, this completes the proof of the theorem. □

Theorem 4.5.4. *Fix $d > 0$. Every graph on $(d+1)$ vertices is a d-dimensional annulus graph for any radii $R_1 \leq R_2$.*

Proof. Let $E_{d,\varepsilon} = \{e_{(i,j)} \mid (i,j) \in [d+1]^{(2)},\ e_{(i,j)} \in \{1, 1+\varepsilon\}\}$ be a set of elements labeled by the pairs in a set of size $(d+1)$ such that each element is either 1 or $1+\varepsilon$.

We will prove by induction on d that for any set $E_{d,\varepsilon}$, given sufficiently small ε, there exists a K_{d+1} with vertex set $v_{[d+1]}$ such that $\text{dist}(v_i, v_j) = e_{(i,j)}$ for each pair $(i,j) \in [d+1]^{(2)}$. The base case $d = 0$ is trivial. Suppose that the claim is true for $d-1$ and consider any set $E_{d,\varepsilon}$. By the inductive hypothesis, picking ε small enough, we can build a K_d with vertex set $v_{[d]}$ and side lengths according to the elements of $E_{d-1,\varepsilon}$ as a subset of the set $E_{d,\varepsilon}$. To complete our construction we need to embed K_d in $\mathbb{R}^d$ and find a point v_{d+1} which will be at a distance $e_{(i,d+1)}$ from the v_i. Such a point has to be the intersection of the d spheres with center v_i and radius $e_{(i,d+1)}$ for $1 \leq i \leq d$. We are going to prove the following.

Sublemma 4.5.5. *For a sufficiently small ε, the intersection of any d $(d-1)$-spheres embedded in $\mathbb{R}^d$ with centres at a distance 1 or $1+\varepsilon$ from each other and radii 1 or $1+\varepsilon$ consists of exactly 2 points.*

Proof. Let the coordinates of the centers of the d spheres be $(a_1^i, a_2^i, \ldots, a_d^i)$ for $1 \leq i \leq d$ and let their radii be $r_1, r_2, \ldots, r_d$. Then the points on each of the spheres satisfy the following equations.

$$(x_1 - a_1^i)^2 + (x_2 - a_2^i)^2 + \cdots + (x_d - a_d^i)^2 = r_i^2, \text{ for } 1 \leq i \leq d. \tag{4.5.1}$$

Subtracting the first equation in (4.5.1) from each of the subsequent ones we obtain a system of $(d-1)$ linear equations in $x_1, x_2, \ldots, x_d$.

Now consider the case where the distances between all centers as well as all radii are 1. The intersection set of these spheres is governed by a system of quadratic equations similar to that in (4.5.1). This system of quadratic equations leads in the same way to a system of $(d-1)$ linear equations.

Let the first pair of systems be called the ε-systems and let the second pair of systems be called the 1-systems. The solution to the linear 1-system is a line. When we substitute the equation of this line in any equation in the quadratic 1-system we get a quadratic equation with two solutions. Overall, this shows that the quadratic 1-system of equations has 2 solutions. Hence, by continuity, we can apply the same argument to the ε-systems to obtain that the quadratic ε-system has two solutions. As the number of solutions of the ε-system is equal to the number of intersections of the d spheres, the proof is complete. □

Sublemma 4.5.5 tells us that there are two choices for v_{d+1} satisfying the required conditions. Pick an arbitrary one of them to find the desired K_{d+1} and thus complete the induction.

Now suppose that we have an arbitrary graph G on $(d+1)$ vertices and we want to prove that it has an annulus embedding in $\mathbb{R}^d$ with radii R_1 and R_2. By rescaling, we can assume that $R_2 = 1$. From the edge pairings of G we can construct a set $E_{d,\varepsilon}$ by setting $e_{(i,j)} = 1$ if (i, j) is an edge in G and $e_{(i,j)} = 1 + \varepsilon$ otherwise. As we proved above, we can choose ε such that there exists a K_{d+1} with side lengths according to the values of $E_{d,\varepsilon}$. It is easy to check that an embedding of this d-simplex is an annulus embedding of G in $\mathbb{R}^d$ with radii $R_1 \leq R_2 = 1$. □

The conclusion that we can derive from the two results in this section is that the set of all annulus graphs up to a given dimension d is strictly growing and that as d goes to infinity it coincides with the set of all graphs. Further possible work in this direction would be to find the steps at which the set of annulus graphs grows. That is, to answer the question 'For any $d > 0$, which graphs are $(d+1)$-dimensional annulus graphs but not d-dimensional?'. Or equivalently, given a graph G, what is the smallest d for which G is d-dimensional?

Chapter 5

Chromatic number of 2-complexes

5.1 Introduction

This paper is part of a project aiming to extend theorems from planar graph theory to three dimensions. The starting point of this project is the series of papers [12, 13, 14, 15, 17]. In there, a three-dimensional analogue of Kuratowski's theorem was proved: embeddability of simply connected 2-complexes in the 3-sphere was characterised by excluded 'space minors'. Other papers from this project are [18] where the forbidden minor characterisation of outerplanar graphs [20] is extended to three dimensions and [19] where the three colour theorem for even degree planar graphs [40] is extended to three dimensions.

The four colour theorem is a central result in graph theory, stating that we can properly colour the vertices of every planar graph with four colours. This theorem together with the example of the four-chromatic planar graph K_4 gives a sharp upper bound of the chromatic number of planar graphs. In this paper we propose and prove a three-dimensional analogue of this theorem.

The first thing we need to do, is to define what is a colouring of a 2-complex. Traditional vertex and edge colourings are not good candidates because every complete graph K_n can be embedded in an orientable surface of high enough genus. Since we know that both the chromatic number and chromatic index of a graph K_n are at least $n-1$, with such definition of a colouring, we can find a 2-complex that requires arbitrary number of colours. A meaningful way to colour 2-complexes is to colour its edges in such manner that no two edges that are adjacent in a face are of the same colour. This type of colouring is motivated by the fact that when we localise it at a vertex, it is equivalent to vertex colouring of the link graph.

A *proper edge-colouring* of a 2-complex is a colouring of its edges such that any two distinct edges that are adjacent in a face have different colours. The *chromatic number* $\chi(C)$ of a 2-complex C is the smallest number of colours for which C has a proper edge-colouring. The *chromatic number* $\chi(\mathcal{C})$ of a class of 2-complexes $\mathcal{C}$ is the maximum chromatic number of a member of $\mathcal{C}$. A similar definition has been given for colouring simplicial complexes as well, for more details, see [19].

Theorem 5.1.1. *The chromatic number of the class of* 2*-complexes embeddable in* $\mathbb{S}^3$ *is* 12.

The following variation of Theorem 5.1.1 remains unsolved.

Open Question 5.1.2. *What is the chromatic number of the class of* $\mathbb{S}^3$*-embeddable simplicial complexes?*

It is also natural to consider Open Question 5.1.2 in a context of 3-manifolds different than $\mathbb{S}^3$.

The reminder of this paper is structured as follows. In Section 5.2, we present some more technical details about 2-complexes and graphs specific to this paper. In Section 5.3 we define link graphs and 2-pire maps and prove some preliminary results related to them. Then the proof of Theorem 5.1.1 is split in two parts. Section 5.4 is devoted to proving the upper bound, while Section 5.5 is devoted to proving the lower bound of Theorem 5.1.1.

5.2 Preliminaries

This section introduces the reader to the main notions of this chapter.

A *directed multigraph* is a quadruple $G = (V, E, \varphi^-, \varphi^+)$, where V and E are sets and their elements are called vertices and edges respectively. The function $\varphi^- \colon E \to V$ assigns a start vertex to each edge and the function $\varphi^+ \colon E \to V$ assigns an end vertex to $\varphi^+ \colon E \to V$ each edge.

In a graph G, an *oriented closed walk* is a finite sequence $v_0, e_0, \ldots, v_n, e_n$ of vertices v_i and edges e_i together with a family $(\sigma_i \mid i \in \mathbb{Z}_{n+1})$ of *traversals* $\sigma_i \in \{-, +\}$ such that $\varphi^{\sigma_i}(e_i) = v_{i+1} = \varphi^{-\sigma_{i+1}}(e_{i+1})$ for all i.

A 2-*complex* is a pair $C = (G, F)$, where G is a graph and F is a family of oriented closed walks in G. We call the elements of F the *faces* of C. A (2-dimensional) simplicial complex is a 2-complex whose underlying graph is simple and all its faces are of size three. We will call a 2-complex with only one vertex a *looped* 2-*complex*. Note that all of the edges of such a 2-complex are necessarily loops.

Definition 5.2.1. A *topological loop* in a topological space X is a continuous map $\ell : \mathbb{S}^1 \to X$.

The *geometric realisation* of a graph $G = (V, E, \varphi^-, \varphi^+)$ is the topological space obtained from $V \bigoplus([0, 1] \times E)$, where V and E carry the discrete topology, by identifying $(0, e)$ with $\varphi^-(e)$ and $(1, e)$ with $\varphi^+(e)$ for every edge $e \in E$.

The *geometric realisation* of a 2-complex $C = (G, F)$ is the topological space obtained from $G \bigoplus(\mathbb{D}^2 \times F)$, where F carries the discrete topology, G is the 1-complex of the graph G and $\mathbb{D}^2$ is the two-dimensional unit disc, by identifying $\partial\mathbb{D}^2 \times \{f\}$ with the closed walk f in G for every $f \in F$. The gluing map is defined as the composition $\partial\mathbb{D}^2 \times \{f\} \to \mathbb{S}^1 \to f$. The first map in the composition is a homeomorphism and the second one comes from the fact that the closed walk f in the geometric realisation of G is a topological loop.

Taking the geometric realisation is an invertible operation and gives a natural bijection between abstract and topological graphs and 2-complexes. In light of this duality, we will use the abstract and topological notions interchangeably when there is no ambiguity.

An *embedding* of a topological space X into a topological space Y is an injective continuous map from X to Y. We say that a 2-complex C is *embeddable* in a topological space T if there exists an embedding of the geometric realisation of C into T. If a 2-complex C is embeddable in the topological space T, we say that C is T-embeddable for short.

Example 5.2.2. A graph is $\mathbb{R}^2$-embeddable if and only if it is planar.

5.3 Link graphs

Consider an embedding in $\mathbb{R}^3$ of the 2-complex C and a 2-sphere S, centered at v, which is small enough so that it intersects C only in edges and faces that are incident to v. This sphere intersects each non-loop edge of C in a point and each loop in two points. For every occurrence of the vertex v in a face, there are two edges that meet at this occurrence of v. The sphere intersects C in one arc for every occurrence of v and this arc connects the two points corresponding to the two edges meeting at this occurrence of v. Thus, the resulting intersection is a graph, which we call the link graph of C at the vertex v.

This geometric definition gives us a recipe to define an abstract link graph by considering the incidence of edges and faces containing the vertex v.

Consider a 2-complex C with a vertex v. Its *link graph* $L(v)$ is defined as follows. For each edge e of C, incident with v we have a vertex e^+ in $L(v)$ if v is the start vertex of e and a vertex e^- in $L(v)$ if v is the end vertex of e. This way if an edge is not a loop it has only one vertex in $L(v)$ and if it is a loop it has two vertices because for loops the start vertex and the end vertex are the same. This agrees with the geometric picture where the sphere intersects a non-loop in one vertex and a loop in two vertices. A face f of C is an oriented closed walk of its underlying graph G. For every occurrence of v in this oriented closed walk there is a pair of edges e_i, e_{i+1} of f such that $\varphi^{\sigma_i}(e_i) = v = \varphi^{-\sigma_{i+1}}(e_{i+1})$. For every such occurrence we get one edge $e_i^{\sigma_i} e_{i+1}^{-\sigma_{i+1}}$ in the link graph $L(v)$.

Remark 5.3.1. We will denote the link graph of the 2-complex C at the vertex v by $L_C(v)$ when we want to stress the underlying 2-complex.

Observation 5.3.2. *The only way for a loop to occur in the link graph $L(v)$ is if a face traverses some edge incident to v twice in a row back and forth.*

Proof. Recall from the paragraph above that an edge of $L(v)$ is of the form $e_i^{\sigma_i} e_{i+1}^{-\sigma_{i+1}}$, for e_i and e_{i+1} – two consecutive edges of some face meeting at v. In order for this to be a loop, we need $e_i = e_{i+1}$ from which the result follows. $\square$

Definition 5.3.3. The *total link graph* of a 2-complex C is defined as the graph that is the disjoint union of all the link graphs of C. We will denote it by $L(C)$.

Definition 5.3.4. A graph with $2n$ vertices is called *paired* graph if its vertices are in n pairs. Two vertices that are in the same pair are called *paired*.

In this paper, we follow the convention that every paired graph on $2n$ vertices has vertex pairs $\{v_i^+, v_i^-\}$, $1 \leq i \leq n$.

Remark 5.3.5. The *total link graph* of a 2-complex C can be considered as a paired graph by letting each pair be the two vertices coming from the same edge of C. From here on when we talk about a paired graph that is the total link graph of some 2-complex we consider it with this pairing.

Remark 5.3.6. If a 2-complex is M-embeddable for any 3-manifold M, we know that its total link graph is planar as the disjoint union of planar graphs. Thus, we get that the total link graph of an M-embeddable 2-complex is a paired planar graph.

A rotation system of a graph G is a family $(\sigma_v | v \in V(G))$ of cyclic orientations σ_v of the edges incident with the vertices v. The orientations σ_v are called rotators. Any rotation system of a graph G induces an embedding of G in an oriented (2-dimensional) surface S. The rotation system is planar if that surface S is a disjoint union of 2-spheres.

A rotation system of a 2-complex C is a family $(\sigma_e | e \in E(C))$ of cyclic orientations σ_e of the faces incident with the edge e. A rotation system of a 2-complex C induces a rotation system on its total link graph $L(C)$ by taking the rotator of e^- to be the rotator of e and the rotator of e^+ to be its inverse for each edge e of C. A rotation system of a 2-complex is planar if the induced rotation system of the total link graph is planar.

Definition 5.3.7. We will call a paired planar graph a 2-pire map.

Remark 5.3.8. In a 2-pire map any two vertices that are in a pair have the same degree.

Definition 5.3.9. A *decoration* of a 2-pire map G is a set of rotators at each vertex of G together with a set of bijections indexed by the vertex pairs of G. For each pair of vertices, its corresponding bijection reverses the order between their rotators. A *decorated 2-pire map* is a 2-pire map together with a decoration.

For a picture of a 2-pire map, look at Figure 5.1. For a picture of its decoration, look at Figure 5.2.

Lemma 5.3.10. *Let M be any 3-manifold. Then the total link graph of an M-embeddable 2-complex C can be given the structure of a decorated 2-pire map.*

Proof. The fact that $L(C)$ is a 2-pire map follows from Remark 5.3.5 and Remark 5.3.6.

Since the 2-complex is M-embeddable, there exists a planar rotation system on the edges of C, induced by some embedding of C in M. This planar rotation system induces a planar rotation system on its total link graph. It is a decoration because the rotator at each edge of C induces reverse rotators to its two copies in the total link graph. $\square$

One of the important ideas in this paper is that there exists a converse to Lemma 5.3.10 which we present below.

Let G be a decorated 2-pire map. Let a^+ and a^- be two vertices in a pair and let a^+b and a^-c be two edges in a bijection dictated by the decoration. We define $\sigma_{a+}(b) = c$ and $\sigma_{a-}(c) = b$. We will call a set of edges of G of the form $v_1^{s_1}v_2^{-s_2}, v_2^{s_2}v_3^{-s_3}, \dots, v_n^{s_n}v_{n+1}^{-s_{n+1}}$, where $s_i \in \{+,-\}$ for $1 \le i \le n$, $v_{i+2}^{-s_{i+2}} = \sigma_{v_i^{s_i}}(v_{i+1}^{-s_{i+1}})$ for $1 \le i \le n$, and $v_{n+j}^{s_{n+j}} = v_j^{s_j}$ for $j \in \{1,2\}$ a *precycle*.

Observation 5.3.11. *Note that any precycle p has a decorated 2-pire map structure induced by G. Hence, for every decorated 2-pire map and a precycle p in G, the subgraph $G - p$ is also a decorated 2-pire map.*

Example 5.3.12. For a 2-complex C the face $v_1^{s_1}v_2^{s_2}\dots v_n^{s_n}$ yields the precycle $v_1^{s_1}v_2^{-s_2}$, $v_2^{s_2}v_3^{-s_3}, \dots, v_n^{s_n}v_1^{-s_1}$ in the total link graph $L(C)$.

Lemma 5.3.13. *For every decorated 2-pire map G, there exists a 2-complex C whose total link graph is G and has a planar rotation system that induces the decoration of G.*

Proof. Suppose that the connected components of G are $c_1, c_2, \dots, c_m$ and its vertices are $v_1^+, v_1^-, \dots, v_\ell^+, v_\ell^-$. We will build a 2-complex C with vertices $c_1, c_2, \dots, c_m$ and edges $v_1, v_2, \dots, v_\ell$, where $v_i = c_j c_k$ if the vertices v_i^+ and v_i^- are in components C_j and C_k.

The final step in building C is to define its set of faces. In order to do that, let us construct a set of edges of G in the following way. Start with an arbitrary edge $e_1 = u_1u_2$ of G. There is a bijection coming from the decoration of G that maps this edge to an edge incident to the vertex that is paired with u_2. Denote this edge by e_2. We can continue this process to obtain a sequence of edges $e_1, e_2, e_3, \dots$. We stop when we get an edge e_n that has the vertex that is paired with u_1 as an endvertex. It is easy to see that this always terminates and when it does, the sequence of edges $e_1, e_2, \dots, e_n$ consists of distinct elements. This set of edges is a precycle $p_1 = v_1^{s_1}v_2^{-s_2}, v_2^{s_2}v_3^{-s_3}, \dots, v_n^{s_n}v_1^{-s_1}$ and defines the walk $f_1 = v_1^{s_1}v_2^{s_2}\dots v_n^{s_n}$. We know from Observation 5.3.11 that $G - p_1$ is a decorated 2-pire map, so we can find a precycle p_2 in $G - p_1$ and a corresponding walk f_2. This process will terminate and will yield precycles $p_1, p_2, \dots, p_n$ and their corresponding walks $f_1, f_2, \dots, f_n$. The 2-complex C with a 1-skeleton as defined above and faces defined by the walks $f_1, f_2, \dots, f_n$ has a total link graph G by Example 5.3.12.

We can give a rotation system σ_C of C by defining the rotator at each edge e of the 2-complex to be the rotator of e^+ in $G = L(C)$ as a cyclic orientation of the edges around e^+ corresponds to a cyclic orientation of the faces around e. A rotator at the edge e of the 2-complex C induces a rotator at the vertex e^+ and the reverse rotator at the vertex e^- in $L(C)$. These induced rotators agree with the decoration of G by construction of σ_C and because in a decoration the rotators at e^+ and e^- are also reverse. Therefore, the rotation system σ_C that we defined induces the decoration of G. Since the decoration of G is a planar rotation system, it follows that σ_C also is, which completes the proof. $\square$

Remark 5.3.14. A crux in the strategy to prove the lower bound of Theorem 5.1.1 is to find a graph with high chromatic number and use Lemma 5.3.13 to build a 2-complex with a high chromatic number. It is easiest to look for the required graph when we restrict our attention to connected graphs. Since the 2-complex that is built in Lemma 5.3.13 has as many vertices as the graph has connected components, in this chapter we are mostly concerned with 2-complexes with only one vertex, or in other words – looped 2-complexes.

5.4 Upper bound of the chromatic number of embeddable 2-complexes

Definition 5.4.1. From a paired graph G we obtain its *paired quotient* graph $Q(G)$ by identifying vertices in the same pair. We denote the vertex obtained from identifying v_i^+ with v_i^- to be v_i.

Remark 5.4.2. Given a 2-complex C, the vertex set of $Q(L(C))$ is indexed by the set of edges of C.

Definition 5.4.3. For a 2-complex C, its *incidence graph* $I(C)$ is the graph on the edges of C in which two vertices are joined by one edge for each occurrence of their corresponding edges being adjacent in a face.

Lemma 5.4.4. *For a 2-complex C, we have that $I(C) = Q(L(C))$.*

Proof. Firstly, let us note that the object $Q(L(C))$ is well defined by Remark 5.3.5. The two graphs have the same vertex set by definition. Every edge xy in $Q(L(C))$ corresponds to an edge $x^\pm y^\pm$ in $L(C)$. An edge $x^\pm y^\pm$ in $L(C)$ corresponds to an adjacency of the edges x and y in some face of C. An adjacency of the edges x and y in some face of C correspond to an edge xy in $I(C)$. This correspondence gives an isomorphism between the edges of $Q(L(C))$ and $I(C)$, which finishes the proof. □

Lemma 5.4.5. *A 2-complex C has a chromatic number equal to the chromatic number of $Q(L(C))$.*

Proof. By the definition of chromatic number of 2-complexes, we have that $\chi(C) = \chi(I(C))$, so the result follows from Lemma 5.4.4. □

Definition 5.4.6. The maximum average degree $\text{mad}(G)$ of a graph G is the maximum of the average degrees of all the subgraphs of G.

Lemma 5.4.7. *For any graph G, we have that $\chi(G) \leq mad(G) + 1$.*

Proof. We will prove this by induction on the number of vertices of G, which we will denote by n. If $n = 1$, we have $\text{mad}(G) = 0$ and $\chi(G) = 1$, which proves the base case. Suppose we have a graph G with n vertices and suppose that the claim holds for all graphs with $n-1$ vertices. There exists a vertex v whose degree is lower than the average degree of G, which in particular means that $\deg(v) \leq \text{mad}(G)$. Since the set of subgraphs of $G - v$ is a subset of the set of subgraphs of G, it follows that $\text{mad}(G - v) \leq \text{mad}(G)$. Hence, we can apply the inductive hypothesis to colour $G - v$ in $\lfloor \text{mad}(G) + 1 \rfloor$ colours. Since $\deg(v) \leq \text{mad}(G)$, there exists a colour of this colouring which is not used by any of the neighbours of v. Colour v in any such colour to obtain a colouring of G with $\lfloor \text{mad}(G) + 1 \rfloor$ colours. This completes the inductive step and hence the proof. □

Now we are ready to prove that for any 3-manifold M, the chromatic number of the class of M-embeddable 2-complexes is at most 12. To do that we start by showing that the chromatic number of the paired quotient of a 2-pire map is at most 12, which is equivalent to the map colour theorem of Heawood [41]. We add a proof for completeness.

Lemma 5.4.8. *For any 2-pire map G, its paired quotient $Q(G)$ can be coloured in 12 colours.*

Proof. Any subgraph of $Q(G)$ is the paired quotient of some 2-pire map that is a subgraph of G. The average degree of any 2-pire map is less than 6 as it is planar, so the average degree of any subgraph of $Q(G)$ is less than 12; that is, $\text{mad}(Q(G)) < 12$. Using Lemma 5.4.7, we can conclude that $\chi(Q(G)) \leq 12$. □

Lemma 5.4.9. *Let M be an arbitrary 3-manifold. Then the chromatic number of any M-embeddable 2-complex is at most 12.*

Proof. Consider an M-embeddable 2-complex C. Since C is M-embeddable, we have that its combined link graph is a 2-pire map. By Lemma 5.4.8 the graph $Q(L(C))$ is 12-colourable. By Lemma 5.4.5 we obtain that C has chromatic number at most 12. □

5.5 Lower bound of the chromatic number of embeddable 2-complexes

As vertices of link graphs correspond to edges of the 2-complex, and edges of the link graphs correspond to faces of the 2-complex, we can also say that faces in a planar embedding of a link graph correspond to chambers (of an embedding in some 3-manifold) of the 2-complex as follows. Let C be a 2-complex with an embedding in some 3-manifold, M, and v a vertex of C. For each face, f, in the plane embedding of $L_C(v)$ that is induced by the embedding of the 2-complex C in M, there exists a chamber in this embedding of C that is incident with all of the edges and faces in C that correspond to the vertices and edges that are incident to f in $L_C(v)$.

Lemma 5.5.1. *Let C be a 2-complex with a planar rotation system, such that all faces have size exactly 3 and for every face f there exists an edge e that is incident only with the face f. Then C is $\mathbb{S}^3$-embeddable.*

Proof. Given a $C = (G, F)$ satisfying the above assumptions, take an arbitrary face f and let e be the edge that is incident only with f. Denote the other two edges of f to be e_1 and e_2 with a common vertex v. Let C_f be the 2-complex obtained from C after deleting the face f and let σ_C be the planar rotation system of C. We will prove that C is embeddable if there exists an embedding of C_f in $\mathbb{S}^3$ induced by σ_C.

Consider the embeddings of $L_C(v)$ and $L_{C_f}(v)$ that are induced by σ_C. Clearly, this embedding of $L_{C_f}(v)$ is contained in the embedding of $L_C(v)$, and the only difference between the two is that the edge corresponding to f is in $L_C(v)$, but it is not in $L_{C_f}(v)$. Since both of these link graphs are planar, this means that there must be a face c that is incident with the vertices that correspond to e_1 and e_2 in $L_{C_f}(v)$. So there exists a chamber in C_f that corresponds to the face g of $L_{C_f}(v)$ that is incident with the edges e_1 and e_2. We can embed the edge e and the face f into c.

This embedding of e and f in c extends the embedding of C_f in $\mathbb{S}^3$ to an embedding of C in $\mathbb{S}^3$. The 2-complex C_f with the planar rotation system induced by σ_C satisfies the same initial assumptions as C, therefore we can inductively continue removing faces to obtain a chain $C = C_0, C_1, \ldots, C_{|F(C)|} = G$ of subcomplexes of C, such that there exists an embedding of C_i in $\mathbb{S}^3$ induced by σ_C if such an embedding exists for C_{i+1}. As G trivially has an embedding induced by σ_C, all the 2-complexes in the chain do. In particular, it follows that C is $\mathbb{S}^3$-embeddable. □

Definition 5.5.2. We say that a paired graph H is a *pruning* of a paired graph G if H can be obtained from G by iteratively removing pairs that consist of vertices of degree 1.

Lemma 5.5.3. *Let C be a 2-complex with a planar rotation system and no faces of size 1. Then there exists C', a 2-complex with a planar rotation system such that all faces have size exactly 3, for every face f there exists an edge e that is incident only with the face f, and $L(C)$ is a pruning of $L(C')$.*

Proof. To construct C', begin by taking the 1-skeleton G of C and apply one operation for each face of C, each of which iteratively adds edges and faces. For every face f in C, let the closed walk that bounds it be W_f and let the operation indexed by it be ι_f. To obtain C', we do a total of $|F|$ operations in an arbitrary order and each ι_f is one of the following, depending on the size of W_f.

If W_f has length 2, it consists of two parallel edges, e_1 and e_2. Attach two new loops, l_1 and l_2, to C', one at each endvertex of e_1 and e_2. Then add the two faces $e_1e_2l_1$ and $e_1e_2l_2$ to C'.

If W_f has length at least 3, for each pair of consecutive edges e_1 and e_2 in W_f, add a new edge e_3 to form the new face $e_1e_2e_3$ in C'.

Let the planar rotation system of the 2-complex C be σ and the planar rotation system that it induces on $L(C)$ be σ_1. For each edge e of C replace the face f in its σ-rotator by the two faces obtained from ι_f containing e (in any order) to obtain a rotator of e in C'. Each edge of C' that is not in C has a trivial rotator of one face. Combine these to get a rotation system σ' of C' that induces the rotation system σ_1' on $L(C')$.

Since the edge set of C is a subset of the edge set of C', we can consider the vertex set of $L(C)$ as a subset of the vertex set of $L(C')$. With this consideration in mind, we have that σ_1' restricted to $V(L(C))$ is isomorphic to σ_1. The elements of $V(L(C')) - V(L(C))$ correspond to the new edges added by the operations ι_f. All of these edges lie only on one face, hence the vertices of $L(C')$ that are not in $L(C)$ are leaves. Adding leaves preserves the planarity of a rotation system, therefore σ_1' is a planar rotation system, hence σ' also is. Note that in the above, we also collaterally proved that $L(C)$ is a pruning of $L(C')$ so we are done. □

Theorem 5.5.4. *Let C be a 2-complex with a planar rotation system and no faces of size 1. Then, there exists a 2-complex C' which is embeddable in $\mathbb{S}^3$ and is such that $L(C)$ is a pruning of $L(C')$. Moreover, if $\chi(C) \geq 3$, then $\chi(C)$ is equal to $\chi(C')$.*

Proof. Given such a 2-complex C, we can use Lemma 5.5.3 to construct a 2-complex C' that satisfies the conditions of Lemma 5.5.1. Then C' is embeddable in $\mathbb{S}^3$ and $L(C)$ is a pruning of $L(C')$.

By Lemma 5.4.5, we have $\chi(C) = \chi(Q(L(C)))$, and $\chi(C') = \chi(Q(L(C')))$. Now, to see that $\chi(Q(L(C))) = \chi(Q(L(C')))$, note that $L(C')$ is obtained from the paired graph $L(C)$ by adding pairs of vertices that each have degree 1. This does not change the chromatic number because $\chi(Q(L(C))) = \chi(C) \geq 3$ and since each added vertex will see at most 2 others, there will be a free colour available to colour it. Then $\chi(C) = \chi(C')$. □

Before we finish the proof of the main result of this section, we need the following.

Lemma 5.5.5. *The paired quotient of the graph G_{12} from Figure 5.1 is isomorphic to K_{12}.*

Proof. We will list the neighbours of each vertex of G_{12}. To see that it quotients to K_{12}, we need to check that the union of the neighbourhoods of each pair $\{a^+, a^-\}$, $\{b^+, b^-\}, \ldots, \{l^+, l^-\}$ has precisely one element from each of the other 11 pairs.

$a^+ : g^-, d^-, i^-, c^-, h^+$

$a^- : k^+, b^+, j^-, f^-, l^-, e^-$

$b^+ : a^-, j^-, l^+, h^-, e^-, k^+$

$b^- : c^+, d^-, i^-, f^-, g^-$

$c^+ : b^-, d^-, g^-$

$c^- : a^+, h^+, j^-, e^+, f^+, k^-, l^-, i^-$

$d^+ : j^-, e^+, f^+, k^-, h^-, l^+$

$d^- : g^-, a^+, i^-, b^-, c^+$

$e^+ : d^+, j^-, c^-, f^+$

$e^- : k^+, a^-, l^-, g^+, i^+, h^-, b^+$

$f^+ : d^+, e^+, c^-, k^-$

$f^- : a^-, j^-, h^+, g^-, b^-, i^-, l^-$

$g^+ : i^+, j^+, k^-, l^-, e^-$

$g^- : a^+, d^-, c^+, b^-, f^-, h^+$

$h^+ : j^-, f^-, g^-, a^+, c^-$

$h^- : b^+, e^-, i^+, k^-, d^+, l^+$

$i^+ : k^-, j^+, g^+, e^-, h^-$

$i^- : f^-, b^-, d^-, a^+, c^-, l^-$

$j^+ : i^+, k^-, g^+$

$j^- : h^+, f^-, a^-, b^+, l^+, d^+, e^+, c^-$

$k^+ : e^-, a^-, b^+$

$k^- : d^+, h^-, i^+, j^+, g^+, l^-, c^-, f^+$

$l^+ : h^-, b^+, j^-, d^+$

$l^- : a^-, f^-, i^-, c^-, k^-, g^+, e^-,$

□

Corollary 5.5.6. *There exists a looped* 2*-complex* C' *that is embeddable in* $\mathbb{S}^3$ *and has chromatic number* 12*.*

Proof. Consider the 12-chromatic 2-pire map G_{12} from Figure 5.1. By replacing each edge of G_{12} with a parallel class of size 2 and adding as many loops as necessary, we obtain a 2-pire map such that any two vertices in a pair have the same degree. Give this 2-pire map an arbitrary decoration to get the decorated 2-pire map G. Then, using Lemma 5.3.13, we can find a 2-complex C whose total link graph is G which has a planar rotation system coming from the decoration of G. The chromatic number of the paired quotient $Q(G)$ is 12 by Lemma 5.5.5, therefore C has a chromatic number 12 as well by Lemma 5.4.5. Since there are no two vertices in a pair connected by an edge in G_{12}, the same holds for G and therefore C has no face of size one. Moreover, C is looped since G_{12} as well as G have only one connected component. Now we can apply Theorem 5.5.4 to demonstrate the existence of the desired looped 2-complex C' that is embeddable in $\mathbb{S}^3$. Note that C' is looped because by construction, C' has the same number of vertices as C. □

Proof of Theorem 5.1.1. Follows directly from Corollary 5.5.6 and Lemma 5.4.9. □

Corollary 5.5.7. *The chromatic number of the class of* M*-embeddable 2-complexes is equal to* 12 *for any* 3*-manifold* M*.*

Proof. According to Corollary 5.5.6, there exists a 2-complex C that is $\mathbb{S}^3$-embeddable and is 12-chromatic. This 2-complex can also be embedded in $\mathbb{R}^3$ and any neighbourhood of M is homeomorphic to $\mathbb{R}^3$, hence C can be embedded in M. We found an M-embeddable 2-complex with chromatic number 12, thus the chromatic number of the class of M-embeddable 2-complexes is at least 12. The upper bound follows from Lemma 5.4.9 which completes the proof. □

5.6 Alternative proof of the lower bound of the chromatic number of embeddable 2-complexes

In Section 5.5 a recipe to construct an $\mathbb{S}^3$-embeddable 2-complex with a chromatic number 12 is given but without explicitly showing an example of such a 2-complex. In this section, at the cost of a bit more work, we will construct such an example which will serve as an alternative proof of the lower bound of Theorem 5.1.1.

To this end, we are going to prove the following theorem, which has very similar statement to Corollary 5.5.6. But before that, we need a definition that we will use in the theorem's proof.

Definition 5.6.1. Given a multigraph G, the unique simple subgraph of G is called the *simplification* of G.

Theorem 5.6.2. *There exists a looped 2-complex C that is embeddable in $\mathbb{S}^3$ and has chromatic number* 12.

Proof. We are going to prove this in two steps. First, we will find a looped 2-complex C_{12} with a link graph that has simplification the graph G_{12} from Figure 5.1. The graph G_{12} is a 2-pire map, with pairs $\{a^+, a^-\}$, $\{b^+, b^-\}$, $\dots$, $\{l^+, l^-\}$ and has a paired quotient K_{12}. Since the chromatic number of K_{12} is 12, by Lemma 5.4.5 it follows that C_{12} has chromatic number 12. Second, we will prove that the so found 2-complex C_{12} is $\mathbb{S}^3$-embeddable by proving that it is simply connected and using [13, Theorem 1.1.].

Consider the decorated 2-pire map G defined by Figure 5.2. The unique 2-complex that has G as a link graph is our desired 12-chromatic 2-complex C_{12}. In Lemma 5.6.3 we describe in detail how to interpret Figure 5.2, we define fully both G and C_{12} and explain their connection to the graph G_{12}.

The next step is to prove that C_{12} is simply connected. To do that, we will prove that its fundamental group $\pi_1(C_{12}, v_0)$ is trivial, where v_0 is the unique vertex of C_{12}. The 2-complex C_{12} is a graph with 12 loops sharing a single vertex and 37 faces glued along words on these loops. Each loop is a generator of $\pi_1(C_{12}, v_0)$ and each face is a relation of $\pi_1(C_{12}, v_0)$. This yields that $\pi_1(C_{12}, v_0)$ is a group with 12 generators and 25 relations since 12 of the 37 relations are repeated. Our strategy is given these relations to prove that all of the generators are equal to the identity element $\mathbb{1}$. As shown in Lemma 5.6.3, the 2-complex C_{12} has the faces dkb^{-1}, fkb^{-1}, gkb^{-1}, kai, kah^{-1}, $l^{-1}ah^{-1}$, cba, cbh, hfe^{-1}, $ef^{-1}g$, cdl^{-1}, lji^{-1}. We have the following equations:

1. dkb^{-1} and fkb^{-1} are faces $\implies dkb^{-1} = fkb^{-1} = \mathbb{1} \implies d = bk^{-1} = f \implies f = d$.
2. fkb^{-1} and gkb^{-1} are faces $\implies fkb^{-1} = gkb^{-1} = \mathbb{1} \implies f = bk^{-1} = g \implies f = g$.
3. kah^{-1} and $l^{-1}ah^{-1}$ are faces $\implies kah^{-1} = l^{-1}ah^{-1} \implies k = ha^{-1} = l^{-1} \implies k = l^{-1}$.
4. $ef^{-1}g$ is a face and $g = f \implies e = ef^{-1}g = \mathbb{1}$, so $e = \mathbb{1}$.
5. hfe^{-1} is a face, $e = \mathbb{1} \implies hfe^{-1} = hf = \mathbb{1} \implies h = f^{-1}$.
6. cba and cbh are faces $\implies cba = cbh = \mathbb{1} \implies a = b^{-1}c^{-1} = h \implies a = h$.
7. kah^{-1} is a face and $a = h \implies k = kah^{-1} = \mathbb{1}$, so $k = \mathbb{1} = l$.
8. cdl^{-1} is a face, $l = \mathbb{1} \implies cd = cdl^{-1} = \mathbb{1} \implies d = c^{-1}$.
9. lji^{-1} is a face, $l = \mathbb{1} \implies ji^{-1} = lji^{-1} = \mathbb{1} \implies i = j$.
10. fkb^{-1} is a face $\implies fkb^{-1} = \mathbb{1} \implies fb^{-1} = \mathbb{1}$, since $k = \mathbb{1} \implies f = b$.
11. kai is a face $\implies ai = \mathbb{1}$, since $k = \mathbb{1}$, $\implies i = a^{-1}$.
12. cba is a face, $a = f^{-1} = b^{-1} \implies c = cba = \mathbb{1}$, so $c = \mathbb{1}$.

From these 12 equations we get that $a^{-1} = b = c^{-1} = d = f = g = h^{-1} = i = j = \mathbb{1}$ and that $e = k = l = \mathbb{1}$. From this, we arrive at the conclusion that $a = b = c = d = e = f = g = h = i = j = k = l = \mathbb{1}$. Therefore, all of the 12 generators of $\pi_1(C_{12}, V)$ are equal to the identity element which means that $\pi_1(C_{12}, v_0)$ is the trivial group. Thus, we proved that the 2-complex C_{12} is simply connected.

Now we give a triangulation of C_{12}. We start by subdividing each loop x into a triangle vx_1x_2. Next we give each face a triangulation as follows. Since all faces of C are made solely out of loops and every loop has length 3, a face of size n has a boundary of length $3n$. We triangulate the inside by adding one more $3n$-gon, one vertex inside the new $3n$-gon connected to all its vertices and connecting the vertices of the boundary $3n$-gon with the corresponding and the adjacent to the corresponding vertices. We give an example for a face of size 3 in Figure 5.3, the construction is similar for a face of larger size. Since triangulation preserves simply connectedness and planarity of rotation systems, we obtain a simply connected simplicial complex T_{12} with a planar rotation system that is homeomorphic to C_{12}. By [13, Theorem 1.1], it follows that T_{12} is $\mathbb{S}^3$-embeddable. Since T_{12} and C_{12} are homeomorhpic it follows that the 2-complex C_{12} is also $\mathbb{S}^3$-embeddable.

We have found the desired simply connected looped 2-complex C_{12} with a link graph that has G_{12} as its simplification so we are done. □

Lemma 5.6.3. *The graph G from the proof of Theorem 5.6.2 is a decorated 2-pire map with simplification G_{12}. Furthermore, we can build a 2-complex C_{12} from G.*

Proof. Figure 5.2 shows a coloured table. The rows of the table are labeled by the 24 vertices $\{a^+, a^-\}$, $\{b^+, b^-\}$, ..., $\{l^+, l^-\}$ of G. Each row of the table shows the rotator of the corresponding vertex with the addition that each entry is doubled in G. So for example the rotator of a^+ is two times g^-, eight times d^-, 12 times i^-, and so on. Each row labeled by a vertex with a '+' superscript shows the neighbours of the corresponding vertex in the positive direction, while each row labeled by a vertex with a '−' superscript shows the neighbours of the corresponding vertex in the negative direction. It is easy to confirm that the simplification of G is G_{12} by comparing the neighbours of each vertex. Since G_{12} is a 2-pire map, it follows that G also is. The rotators that we defined above give a decoration to G, hence G is a decorated 2-pire map.

To build C_{12} from G, we will find all the precycles of G which in turn will tell us all the faces of C_{12}. Start from an arbitrary entry of the table, it corresponds to an edge of G. Its paired edge in the decoration corresponds to the entry above it if the entry is in a row with a negative superscript and the entry below it otherwise, this is the next edge in the precycle. Every edge has two copies in this table, one for each of its endpoints. Take the other entry corresponding to the new edge and repeat the process. At some point we have to repeat an edge (or equivalently an entry), which will clearly be the edge (entry) we started from. The ordered edge set arising from this process is a precycle of G. We can mark the entries corresponding to this precycle and pick a non-marked entry to obtain a new precycle by repeating the same algorithm. Continuing this procedure we obtain all precycles of G and hence all the faces of C_{12}. In Figure 5.2 each unique precycle is coloured in a corresponding unique colour (25 in total). Each of the 12 precycles marked on the table that do not contain any loops have two copies in G. On the other hand, the 13 precycles containing a loop have a single copy in G, because each loop takes both copies of the 'doubled' entry of the table. This makes a total of 37 precycles in G. Let us give an example derivation of two precycles – one without a loop and one with loops.

Observation 5.6.4. *Consider the red entry of a^+ in the row c^-, this gives that the first edge in the precycle is a^+c^-. Above it we have the entry b^- in the row c^+, so the next edge is c^+b^-. The other entry corresponding to this edge is the c^+ entry in the b^- row. Above it is the a^- entry in the b^+ row, so the next edge is b^+a^-. The other entry corresponding to this edge is the b^+ entry of the a^- row. Above it is the c^- entry of the a^+ row. The next step completes the precycle by getting back to the entry of a^+ in the row c^- corresponding to the edge a^+c^-. This gives us the precycle $\{a^+c^-, c^+b^-, b^+a^-\}$ and in turn gives us the face acb.*

Observation 5.6.5. *Consider the purple entry of b^+ in the row b^+, the first edge in this precycle is b^+b^+ or in other words the loop based at b^+. Below it we have the entry g^+ in the row e^-, so the next edge is e^-g^+. Continuing in this manner we get the edges g^-h^+, h^-k^- and k^+k^+. The second entry corresponding to the edge k^+k^+ is the cell right next to it and below it we have the entry h^- in the row k^- corresponding to the edge h^-k^-. The edges in the rest of the precycle are h^+g^- and g^+e^- after which we go back to e^+e^+. Overall, the precycle is $\{e^+e^+, e^-g^+, g^-h^+, h^-k^-, k^+k^+, k^-h^-, h^+g^-, g^+e^-\}$. Note that the* *edges between the two loops are mirrored. This holds for all thirteen precycles with two loops. The resulting face is* $e^{-1}g^{-1}h^{-1}kk^{-1}hge$.

In a similar way as we did in Observation 5.6.4 and Observation 5.6.5, we can find all precycles and thus all faces of C_{12}. Hence, we found a 2-complex C_{12} which has a link graph with a simplification G_{12} as required. □

5.7 Concluding remarks

This chapter was devoted to analyzing the chromatic number of a looped 2-complex, which as we showed is equivalent to finding the chromatic number of $Q(L(C))$, where $L(C)$ is the total link graph of C. In this case $L(C)$ has one connected component and we can give it a structure of a 2-pire map. With this, we showed that the chromatic number of a general $\mathbb{S}^3$-embeddable 2-complex is 12. The final goal of the project in this chapter is to find the chromatic number of the entire set of simplicial complexes. Below, we will present this question in terms of the chromatic number of a graph as well as give the current best known bounds. We will also do that for an intermediate class of 2-complexes – the class of loopless 2-complexes.

The upper bound of 12 holds for all 2-complexes and is the best known. So this is the current upper bound both for loopless 2-complexes and simplicial complexes. We are also going to explore one more type of 2-complexes.

Definition 5.7.1. A 2-complex is called *parallel* if it has two vertices and has no loops.

Definition 5.7.2. Given a graph G, its *thickness* is the smallest integer t such that G can be represented as the union of t planar graphs and is denoted by $\Theta(G)$.

If we consider a parallel complex C, we have that its total link graph $L(C)$ has two connected components. Then $Q(L(C))$ can be seen as a graph with thickness two. Hence, finding the chromatic number of parallel 2-complexes is equivalent to finding the chromatic number of graphs with thickness two. There exist 9-chromatic thickness two graphs, the first example due to T. Sulanke is shown in Figure 5.4. A family of infinitely many examples later given in [8]. In a similar way that we constructed an $\mathbb{S}^3$-embeddable 2-complex C_{12} from the 2-pire map in Figure 5.1, we can construct an $\mathbb{S}^3$-embeddable 2-complex C_9 from the example in Figure 5.4. This shows us that the chromatic number of loopless 2-complexes embeddable in $\mathbb{S}^3$ is between 9 and 12.

If we have a simplicial complex C, its total link graph $L(C)$ has as many connected components as the number of vertices of C. For every edge e of C, its two copies in $L(C)$ are in different connected components because simplicial complexes do not have loops. Moreover, for every pair of edges e and f, it is not possible for their four corresponding vertices to be in 2 connected components, because simplicial complexes do not have parallel edges. Hence, finding the chromatic number of simplicial complexes reduces to the following question.

Open Question 5.7.3. *What is the chromatic number of a* 2*-pire map with the following restrictions.*

1. *Two vertices in a pair are in different connected components.*
2. *Four vertices from two pairs are in at least three components.*

The best known lower bound to the author comes from the cone over a tetrahedron. This is a simplicial complex C which has a 1-skeleton K_5 and has all 10 possible faces. It is easy to check that the graph $Q(L(C))$ is 5-chromatic.

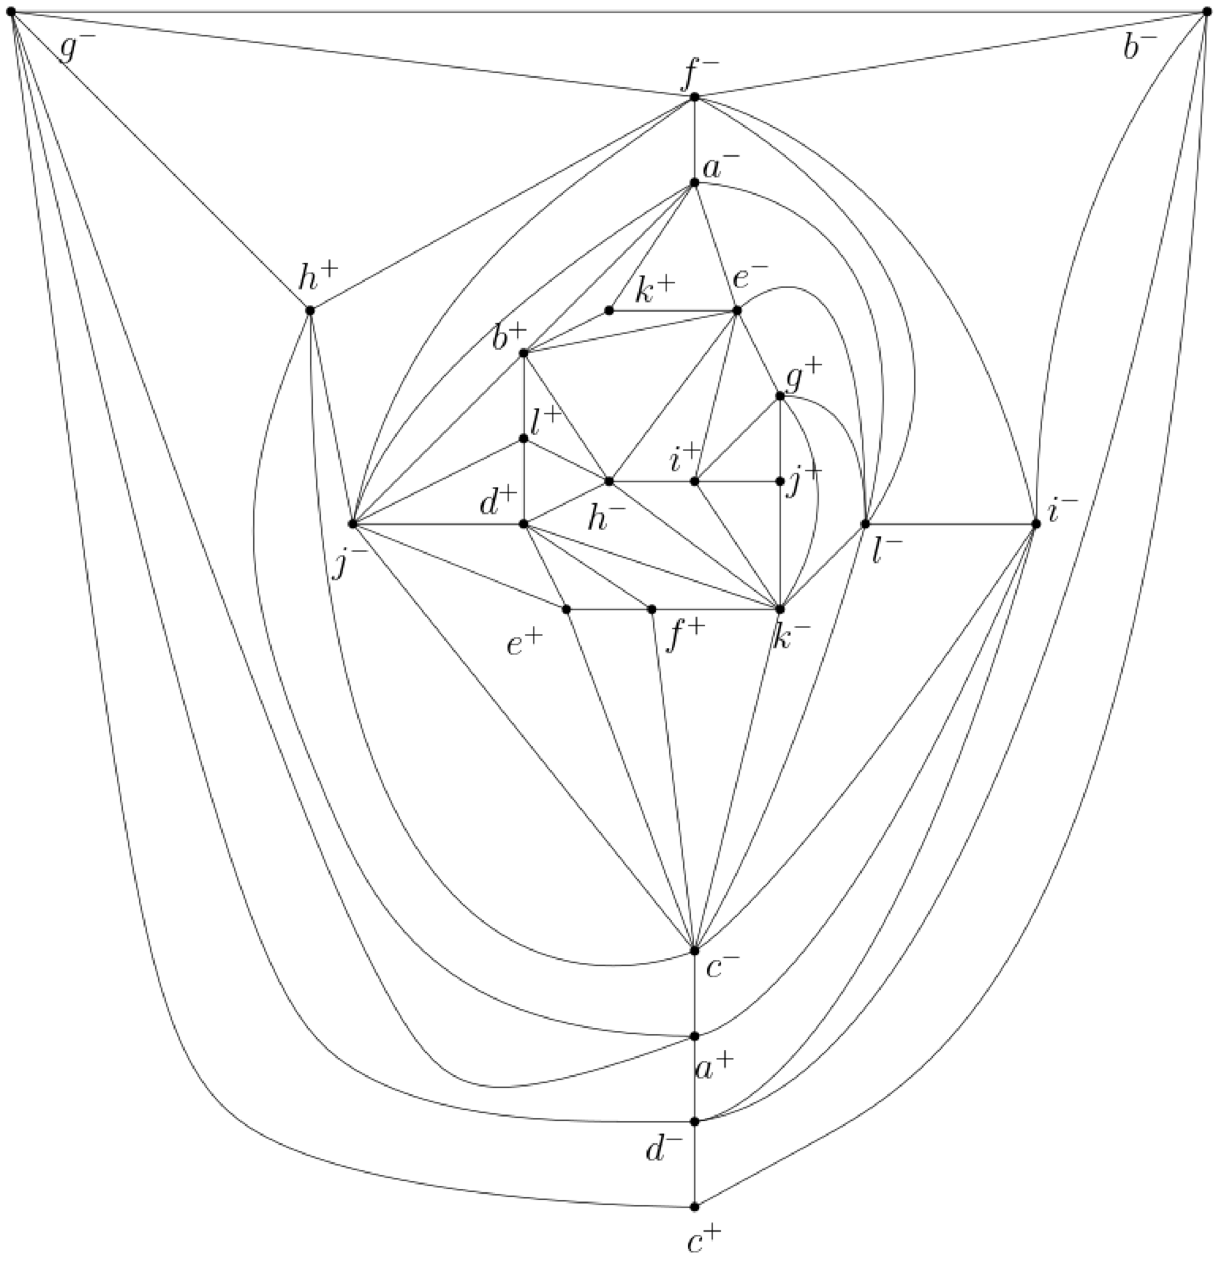

Figure 5.1: The planar paired graph G_{12} which quotients to K_{12}. This example was found in [33], where it is attributed to Heawood.

a+	g-	d-	d-	d-	d-	i-	i-	i-	i-	i-	i-	c-	c-	c-	c-	h+	h+	h+		
a-	k+	k+	k+	k+	k+	k+	k+	k+	b+	b+	b+	b+	j-	f-	l-	l-	e-	k+		
b+	a-	a-	a-	a-	j-	j-	l+	l+	h-	h-	h-	e-	k+	k+	k+	k+	k+	k+		
b-	c+	c+	c+	c+	c+	c+	c+	c+	c+	d-	d-	d-	d-	i-	f-	f-	g-	g-		
c+	b-	b-	b-	b-	b-	b-	c+	c+	c+	c+	c+	c+	d-	d-	g-	g-	g-	b-	b-	b-
c-	a+	a+	a+	a+	h+	h+	j--	e+	e+	f+	f+	k-	l-	i-	c-	c-	c-	c-	c-	c-
d+	j-	e+	e+	e+	f+	f+	f+	k-	k-	h-	l+	l+	l+							
d-	g-	a+	a+	a+	a+	i-	i-	b-	b-	b-	b-	c+	c+							
e+	d+	j-	c-	c-	f+	e+	e+	e+	f+	f+	f+	d+	d+							
e-	k+	k+	a-	l-	g+	g+	i+	i+	h-	b+	e-	e-	e-							
f+	d+	e+	e+	e+	f+	e+	c-	c-	k-	k-	d+	d+								
f-	a-	j-	h+	h+	h+	g-	b-	f-	b-	i-	i-	l-								
g+	i+	j+	j+	j+	j+	k-	l-	e-	e-											
g-	a+	d-	c+	c+	c+	b-	b-	f-	h+											
h+	j-	f-	f-	f-	g-	a+	a+	a+	c-	c-										
h-	b+	b+	e-	i+	k-	k-	d+	l+	l+	b+										
i+	k-	k-	k-	k-	k-	k-	j+	j+	j+	j+	j+	j+	j+	g+	e-	e-	h-			
i-	f-	f-	b-	d-	d-	a+	a+	a+	a+	a+	a+	c-	l-	l-	l-	l-	l-			
j+	i+	i+	k-	k-	g+	g+	g+	g+	i+	i+	i+	i+	i+							
j-	h+	f-	a-	b+	b+	l+	l+	l+	l+	l+	d+	e+	c-							
k+	e-	e-	a-	k+	a-	a-	a-	a-	a-	a-	a-	a-	b+	b+	b+	b+	b+	b+		
k-	d+	k-	h-	h-	i+	i+	i+	i+	i+	i+	j+	j+	g+	l-	c-	f+	f+	d+		
l+	h-	b+	b+	j-	j-	j-	j-	j-	d+	d+	l+	l+	d+	h-						
l-	a-	f-	i-	i-	i-	i-	i-	l-	l-	c-	k-	g+	e-	a-						

Legend:	
kdb^{-1}	
$fb^{-1}k$	
$gb^{-1}k$	
kia	
$h^{-1}ak$	
$h^{-1}al$	
abc	
chb	
fhe^{-1}	
egf^{-1}	
$dc^{-1}l$	
$li^{-1}j$	

Figure 5.2: G is a decorated 2-pire map with simplification G_{12}. Each row in the figure shows the rotators of one of the vertices of G in terms of its neighbours. The rotator goes in the positive direction for vertices with a '+' superscript and in the negative direction for the vertices with a '−' superscript. Each different colour represents one of the 25 faces of C_{12}. The 12 faces which are used to prove simply connectedness are the faces of size three shown to the right. The other 13 colours corresponding to the other 13 faces are also added for convenience of the reader. Loops are coloured in white.

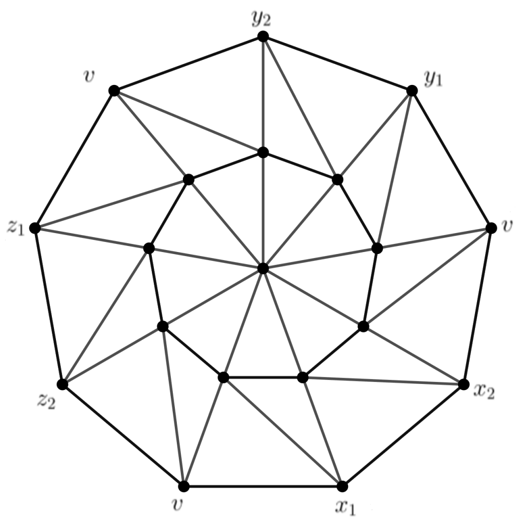

Figure 5.3: Triangulation of a face of size 3

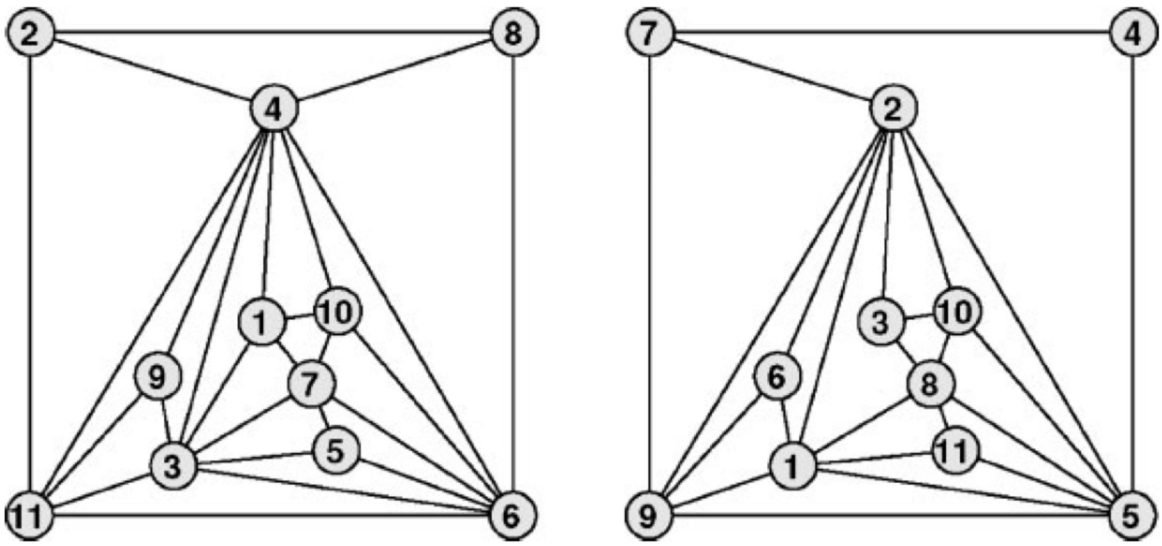

Figure 5.4: A 9-chromatic graph of thickness two. Two vertices are in a pair if they are labeled with the same number.

Chapter 6

Mac Lane's embeddability criterion in three dimensional space

6.1 Introduction

This chapter is part of a project aiming to extend theorems from planar graph theory to three dimensions. The starting point of this project is the series of papers [12, 13, 14, 15, 17]. In there a three-dimensional analogue of Kuratowski's theorem was proved: embeddability of simply connected 2-complexes in the 3-sphere was characterised by excluded 'space minors'. Another papers from this project are [18] where the authors extend the forbidden minor characterisation of outerplanar graphs [20] to three dimensions and [19] where the authors find an extension of the three colour theorem for even degree planar graphs [40] in three dimensions.

The result that this chapter aims to extend in three dimensions is Mac Lane's theorem, which goes as follows.

Theorem 6.1.1. *[49] A graph is planar if and only if there exists a sparse spanning set of its cycle space,*

In the two sections below we are going to prove Mac Lane's embeddability criterion in three dimensions using two different planarity criteria as a basis. In Section 6.2 we will directly use Mac Lane's planarity criterion to prove Theorem 6.2.2. However, using the original author's argument does not suffice for our three dimensional needs because planarity of the link graphs is not enough for embeddability of the underlying simplicial complex. That is why we are going to use a more sophisticated approach coming from [47], which gives not only planarity but also a planar rotation system on the link graphs, which we can then extend to a planar rotation system of the whole simplicial complex. The other planarity criterion is in terms of dual matroids – Whitney's theorem, which says that a graph is planar if and only if it dual matroid is graphic. In [15] a three dimensional analogue of Whitney's theorem was proved. In Section 6.3 we are going to give a proof of Theorem 6.3.2 using this three-dimensional analogue of Whitney's theorem. Let us remark that Theorem 6.2.2 and Theorem 6.3.2 have the same statement.

6.2 Mac Lane's embeddability criterion in three dimensions using Mac Lane's planarity criterion

6.2.1 Preliminaries

The task in this section is to prove an extension of Mac Lane's theorem to three dimensions. We have our desired three dimensional extension of a graph, namely the simplicial complex. However, we still need an extension of the notion of a cycle space one dimension higher. We do this as follows.

We will use signs to signify orientations, so for example $ABC = -ACB$ and $AB = -BA$. Given a face ABC we define the boundary operation ∂_2 on it by $\partial_2(ABC) \coloneqq AB + BC - AC$. Given an edge AB, we define the boundary operation ∂_1 on it by $\partial_1(AB) = B - A$. We will use just ∂ for both boundary operations when there is no ambiguity.

Consider a simplicial complex C with a vertex set V, an edge set E and a face set F. Let $\mathcal{F}_C = \mathbb{F}_p^{|F|}$ be the vector space generated by basis vectors labeled by the set F over the field $\mathbb{F}_p$ for some prime p. We will call $\mathcal{F}_C$ the face space of the 2-complex C. Let $\mathcal{C}(C)$ be the kernel of the map ∂_2, or in other words, the vector space consisting of vectors corresponding to sets of faces whose boundaries sum to 0. We refer to $\mathcal{C}(C)$ as the *p-ary cycle space* of C. Its elements are called the 2-*chains* of C. We will call a minimal (not containing others), nonempty 2-chain a 2-cycle. We can similarly define the edge space $\mathcal{E}_G$ of a graph G. Using ∂_1 we can also define the cycles space of G and we will denote it by $\mathcal{C}(G)$. Its elements are called the chains of G and its minimal elements are called cycles.

Recall the definition of a sparse set of vectors.

Definition 6.2.1. Let k be a finite field and E be a set. A set $\mathcal{S}$ of vectors in k^E is *sparse* if for each coordinate $e \in E$ either for all vectors we have $v(e) = 0$ or we have precisely one vector $v \in \mathcal{S}$ with $v(e) = 1$ and precisely one vector $v \in \mathcal{S}$ with $v(e) = -1$.

Now, we are ready to state the main result of this section.

Theorem 6.2.2. *Let C be a local, simply connected 2-dimensional simplicial complex. Then C embeds in $\mathbb{R}^3$ if and only if there exists a sparse spanning set of its cycle space.*

We will prove the locally 2-connected case.

Theorem 6.2.3. *Let C be a local, locally 2-connected and simply connected 2-dimensional simplicial complex. Then C embeds in $\mathbb{R}^3$ if and only if there exists a sparse spanning set of its cycle space.*

And show that it implies the general case.

Lemma 6.2.4. *Theorem 6.2.3 implies Theorem 6.2.2.*

Proof. Let say that a simplicial complex is of type A if it is local and simply connected. If it is additionally locally 2-connected, say that it is of type B. From the environment 'Proof of Theorem 1.2' in [15] we can extract the following.

Sublemma 6.2.5. *Any simplicial complex of type A can be built inductively by starting from a simplicial complex of type B and at each step gluing a new simplicial complex of type B at a vertex or an edge of the simplicial complex obtained so far. Moreover, the simplicial complex obtained at each step is of type A.*

We will also prove the following.

Sublemma 6.2.6. *Let Z be the simplicial complex obtained by gluing the simplicial complexes X and Y at a vertex or an edge. Then Z has a sparse spanning set of its cycle space if and only if X and Y do.*

Proof. It is easy to see that $\mathcal{C}(Z)$ is the direct sum of $\mathcal{C}(X)$ and $\mathcal{C}(Y)$. Furthermore, if $\mathcal{X}$ is a sparse spanning set of $\mathcal{C}(X)$ and $\mathcal{Y}$ is a sparse spanning set of $\mathcal{C}(Y)$, then $\mathcal{X} \cup \mathcal{Y}$ is a sparse spanning set of $\mathcal{C}(Z)$. Combining these two, the result follows. □

Assume that Theorem 6.2.3 holds. We will prove that all simplicial complexes of type A are embeddable if and only if they have a sparse spanning set by induction on the number of steps it takes to build it via Sublemma 6.2.5. The base case is Theorem 6.2.2 and the inductive step follows from Sublemma 6.2.6. The result that we proved is equivalent to Theorem 6.2.2, so we are done. □

As our result is extension of Mac Lane's planarity criterion, we will use the latter as a main ingredient in our proof. We will use the following algebraic topology version of Mac Lane's theorem due to Lefschetz. In particular, we will utilise a technique used there which not only gives planarity of the graph but builds a planar rotation system.

Theorem 6.2.7. *(Eq. (5) from [47]) A 2-connected graph G is planar if and only if it has a set of cycles $\lambda_0, \lambda_1, \ldots, \lambda_R$ satisfying the following conditions:*

(a) R is equal to the dimension of $\mathcal{C}(G)$.

(b) Every edge of G belongs to exactly two of the λ_i and $\sum \lambda_i = 0$.

(c) The only dependency relation between the cycles is $\sum \lambda_i = 0$.

We show below the embeddability criterion that we will use for the final argument for the convenience of the reader.

Theorem 6.2.8. *([13, Theorem 1.1]) A local simply connected simplicial complex is embeddable in $\mathbb{R}^3$ if and only if it has a planar rotation system.*

In this section, we will work over the field $\mathbb{F}_3$ as it is the simplest field which covers all intricacies of the proofs. However, Theorem 6.2.2 holds for every finite field as all the underlying results are true over the integers. This agrees with the insights of [16, Corollary 5.2].

The reminder of this section is structured as follows. The 'if' implication of Theorem 6.2.3 is proved in Subsection 6.2.2.The 'only if' implication of Theorem 6.2.3 is proved in Subsection 6.2.3. At the end of Subsection 6.2.3, we complete the proof of Theorem 6.2.2.

6.2.2 Proof of the 'if' implication of Theorem 6.2.3

Let C be a locally 2-connected and simply connected simplicial complex. The strategy of this subsection is given a sparse spanning set of the cycle space of C to obtain a sparse spanning set of the cycle space of each link graph and show that these sparse spanning sets induce a planar rotation system on their respective link graphs. Then we show that these planar rotation systems can be extended to a planar rotation system of C. As C is simply connected, we can finish with an application of Theorem 6.2.8.

Consider a face $f = uvw$ of the 2-complex C incident to some vertex v. We define a map $\mathcal{L}_v : \mathcal{F}_C \to \mathcal{E}_{L(v)}$ by the following $\mathcal{L}_v(uvw) := UW$, where U and W are the vertices of $L(v)$ corresponding to the edges uv and uw respectively. Note that $\mathcal{L}_v$ is induced by the map sending a face of C containing the vertex v to an edge of $L(v)$. Additionally, define $\partial_v(f) = vw - vu$ and note that $\partial\mathcal{L}_v(f) = \partial_v f$.

Definition 6.2.9. Define Q_v to be the map $\mathcal{C} \to \mathcal{C}(L(v))$ induced by $\mathcal{L}_v$. More specifically, given a 2-chain S of C, we define

$$Q_v(S) = \sum_{s \in S} \sum_{v \in s} \mathcal{L}_v(s)$$

Proposition 6.2.10. *Q_v is well-defined.*

Proof. This is equivalent to proving that the image of a 2-chain of C under the map Q_v is a chain of $L(v)$. Let S be a 2-chain of the 2-complex C, which is to say $\sum\limits_{s \in S} \partial s = 0$. Since edges incident to v are contained only in faces incident to v this implies that $\sum\limits_{s \in S} \sum\limits_{v \in s} \partial_v s = 0$. Therefore, we have $\partial \left(\sum\limits_{s \in S} \sum\limits_{v \in s} \mathcal{L}_v(s) \right) = \sum\limits_{s \in S} \sum\limits_{v \in s} \partial \mathcal{L}_v(s) = \sum\limits_{s \in S} \sum\limits_{v \in s} \partial_v s = 0$, which means that $Q_v(S)$ is a chain of $L(v)$. $\square$

Lemma 6.2.11. *A set of faces X in $F(v)$ which form a cycle in $L(v)$, lie in some 2-cycle of C.*

Proof. Since the dual matroid of C is local, we have $M^*(C) \restriction F(v) = M^*(L(v))$. By deleting the elements $F(v) - X$ from both matroids we get $M^*(C) \restriction X = M^*(L(v) \restriction X)$. Since $M^*(L(v) \restriction X)$ is a set of parallel edges, it follows that $M^*(C) \restriction X$ also is. This means that the set of faces X lie on some 2-cycle of C. $\square$

Corollary 6.2.12. *The map Q_v is surjective.*

Proof. From Lemma 6.2.11 we get that for any cycle z of $L(v)$, there exists a 2-cycle Z for which $Q_v(Z) = z$ or $Q_v(-Z) = z$. $\square$

Lemma 6.2.13. *If a local 2-complex C has a sparse spanning set $\mathcal{S}$ of its cycle space $\mathcal{C}$, then for any vertex v we have that $\mathcal{S}_v = \{Q_v(S) | S \in \mathcal{S}\}$ is a sparse spanning set of the cycle space of $L(v)$.*

Proof. Let us set $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ and let $Q_v(S_i) = s_i$, so $\mathcal{S}_v = \{s_1, s_2, \ldots, s_n\}$.

Consider any chain s of $L(v)$. By Corollary 6.2.12, there exist a 2-chain S of C such that $Q_v(S) = s$. Since $\mathcal{S}$ is spanning, there exists an $I \subseteq [n]$ such that $\sum\limits_{i \in I} S_i = S$. This means that $\sum\limits_{i \in I} s_i = s$, hence $\mathcal{S}_v$ is spanning.

Let $\mathcal{A}$ be the matrix whose columns are labeled by the faces of C and whose rows are the characteristic vectors of the elements of $\mathcal{S}$. Since $\mathcal{S}$ is sparse, every column of this matrix has precisely one 1, precisely one -1 and the rest of the entries are 0.

Given the characteristic vector of S_i, for $1 \leq i \leq n$, we obtain the characteristic vector of s_i by relabeling each entry f with $\mathcal{L}_v(f)$ and by deleting all entries corresponding to faces not incident to v. Thus, the matrix $\mathcal{A}'$ whose columns are labeled by the edges of $L(v)$ and whose rows are the characteristic vectors of the elements of $\mathcal{S}_v$ is obtained from $\mathcal{A}$ by deleting the columns corresponding to faces not containing v. Since every column of $\mathcal{A}$ has precisely one 1, precisely one -1 and the rest of its entries are 0, the same holds for $\mathcal{A}'$. Therefore, $\mathcal{S}_v$ is sparse.

Note that the fact that $\mathcal{S}_v$ is sparse and spanning implies that all the s_i are different, thus $\mathcal{S}_v$ is also a set. □

Corollary 6.2.14. *The map Q_v is a bijection when restricted to a sparse spanning set $\mathcal{S}$.*

Lemma 6.2.15. *Let G be a 2-connected graph with a sparse spanning set $\mathcal{S}$ of its cycle space $\mathcal{C}(G)$. Then the following three statements hold.*

(i) There does not exist a proper subset of $\mathcal{S}$ that is linearly dependent.

(ii) $|\mathcal{S}| = \dim(\mathcal{C}(G)) + 1$

(iii) Each element of $\mathcal{S}$ is a cycle.

Proof. (i) Suppose for contradiction that such a subset $\mathcal{S}'$ exists and let $\mathcal{S}'' = \mathcal{S} - \mathcal{S}'$. The set $\mathcal{S}'$ is linearly dependent and a subset of the sparse set $\mathcal{S}$, hence $\mathcal{S}'$ is also sparse. Therefore, all elements of $\mathcal{S}'$ are edge disjoint from all elements of $\mathcal{S}''$. Since $\mathcal{S}'$ is proper, there exists one edge covered by $\mathcal{S}'$ and another edge covered by $\mathcal{S}''$. We now get a contradiction with the fact that any two edges of a 2-connected graph lie on some cycle.

(ii) From (i) it follows that for any element s of $\mathcal{S}$, the set $\mathcal{S} - s$ is a basis of $\mathcal{C}(G)$. The result is an immediate consequence of this.

(iii) Suppose for a contradiction that there exists an element $s \in \mathcal{S}$ such that $s = s' + s''$, where s' and s'' are proper disjoint subchains of s. By (i), either s' or s'' is not in $\mathcal{S}$, we can assume that it is s'. From (ii) it follows that $s' \in Sp(\mathcal{S} - s)$. Furthermore, note that $\mathcal{S}' = (\mathcal{S} - s) \cup s' \cup s''$ is also a sparse spanning set of $\mathcal{C}(G)$. But now, we obtain that $(\mathcal{S} - s) \cup s'$ is a linearly dependent proper subset of $\mathcal{S}'$, which is a contradiction with (i). □

Consider a 2-connected graph G with a sparse spanning set $\mathcal{S}$ of its cycle space. Let v be an arbitrary vertex of G and let e_1 be an arbitrary edge pointing towards it. Since G is 2-connected, there exists a cycle containing e_1, thus it is covered by the cycle space of G. As $\mathcal{S}$ is sparse, there exists a unique element z_1 of $\mathcal{S}$ that has e_1 with a positive sign. Let $-e_2$ be the next edge in z_1. We can continue this process with e_2 to obtain a cycle z_2 and an edge $-e_3$ and so on until we repeat an edge (with opposite sign). It is clear that this edge will be $-e_1$. At this point we have the sequence $e_1, z_1, e_2, z_2, \ldots, e_n, z_n$ such that z_i contains e_i and $-e_{i+1}$ ($e_{n+1} = e_1$). The z_i are distinct by Lemma 6.2.15(iii) and since a cycle can contain at most two edges incident to a given vertex. Such a sequence is called an *umbrella* centered at v. Since each edge is in a unique cycle of $\mathcal{S}$ with a positive sign and in a unique cycle of $\mathcal{S}$ with a negative sign, we can partition the edges and cycles incident with v with umbrellas centered at v. This partition is unique up to cyclic reordering of the umbrellas, because once the first edge is fixed, the umbrella sequence is determined. We can do this at each vertex to obtain its corresponding set of umbrellas partitioning the edges and cycles incident to it.

Lemma 6.2.16. *If a 2-connected graph G has a sparse spanning set $\mathcal{S}$ of its cycle space, then it has a set of cycles $\lambda_0, \lambda_1, \ldots, \lambda_R$ satisfying conditions $(a) - (c)$ of Theorem 6.2.7.*

Proof. The set $\mathcal{S}$ consists only of cycles due to Lemma 6.2.15(iii). It satisfies condition (a) by Lemma 6.2.15(ii), it satisfies condition (b) because it is sparse and it satisfies condition (c) due to Lemma 6.2.15(i). Hence, $\mathcal{S}$ is the desired set of cycles λ_i. □

Lemma 6.2.17. *Given a 2-connected graph G with a sparse spanning set $\mathcal{S}$ of its cycle space, each vertex has a unique umbrella. Moreover, these umbrellas induce a planar rotation system of G.*

Proof. By Lemma 6.2.16, G satisfies the conditions $(a)-(c)$ of Theorem 6.2.7. In the proof of Theorem 6.2.7 is shown that each vertex has a unique umbrella. Moreover, if the umbrella at a vertex v is $u_v = e_1^v, z_2^v, e_2^v, z_2^v, \dots, e_n^v, z_n^v$, the rotation system with the rotator of v being $e_1^v, e_2^v, \dots, e_n^v$ is planar. $\square$

Proof of the 'if' implication of Theorem 6.2.3. Let C be a local, locally 2-connected and simply connected 2-complex with a sparse spanning set $\mathcal{S}$ of its cycle space and let v be an arbitrary vertex of C. The sparse spanning set $\mathcal{S}$ induces a sparse spanning set $\mathcal{S}_v$ of $\mathcal{C}(L(v))$ by Lemma 6.2.13. Since $L(v)$ is 2-connected, we can partition the edges and faces incident with v in umbrellas. By Lemma 6.2.17, all vertices of $L(v)$ have a unique umbrella and these umbrellas induce a planar rotation system σ_v. We can do this for each vertex u of C to obtain planar rotation systems σ_u for $L(u)$ induced by their umbrellas.

Given an edge $e = uv$, suppose its umbrella at $L(u)$ is $f_1, z_1, f_2, z_2, \dots, f_n, z_n$. Let $Z_i = Q_v^{-1}(Z_i)$ and $z_i' = Q_u(Z_i)$ – this is well defined due to Corollary 6.2.14. We can see that Z_i contains f_i and $-f_{i+1}$ in $\mathcal{C}(C)$ because z_i do so in $\mathcal{C}(L(v))$, therefore z_i' contains f_i and $-f_{i+1}$, for $1 \leq \ i \leq n$ and $f_{n+1} = f_1$. Now it follows that the umbrella of e at $L(v)$ is $-f_n, z_{n-1}', \dots, -f_1, z_n'$. Therefore, if we set the rotator at e of the 2-complex C to be $f_1, f_2, \dots, f_n$, it induces the same rotators on e in $L(u)$ and $L(v)$ that are induced by its respective umbrellas. Doing this for every edge we get a rotation system σ of C. This rotation system induces the planar rotation system σ_v on each link graph $L(v)$. Since each σ_v is planar, σ also is. As C is simply connected, the result follows from Theorem 6.2.8. $\square$

6.2.3 Proof of the 'only if' implication of Theorem 6.2.3 and completing the proof of Theorem 6.2.2

Let C be a locally 2-connected, simply connected and local simplicial complex. The strategy of this subsection is given an embedding of C in $\mathbb{R}^3$ to describe the chambers of this embedding and prove that their boundaries are the required sparse spanning set. This proof is based on the proof of Mac Lane's planarity criterion in [23, Theorem 4.5.1].

Lemma 6.2.18. *Consider a local 2-complex C embedded in $\mathbb{R}^3$ and an arbitrary vertex v of C. Then, there is a bijection between the faces of the (planar) link graph $L(v)$ and the chambers of C, containing v. Similarly, there is a bijection between the face boundaries of $L(v)$ and the boundaries of the chambers of C, containing v.*

Proof. The faces of $L(v)$ are in a bijection with the cycles of $M^*(L(v))$ and the chambers of C are in a bijection with the cycles of $M^*(C)$. Since, by locality $M^*(L(v)) \cong M^*(C) \upharpoonright F(v)$, it follows that the faces of $L(v)$ are in bijection with the chambers of C containing v, which proves the first part.

Similarly, the face boundaries of $L(v)$ are in a bijection with the cycles of $M(L(v))$ and the chamber boundaries of C are in a bijection with the cycles of $M(C)$. Again, by locality $M(L(v)) \cong M(C) \upharpoonright F(v)$ and thus the face boundaries of $L(v)$ are in bijection with the chamber boundaries of C containing v, proving the second part. □

Lemma 6.2.19. *Consider a locally 2-connected 2-complex embedded in $\mathbb{R}^3$ which is local. Then every chamber of the embedding is bounded by a 2-cycle.*

Proof. Take any vertex v of C and consider the plane embedding of $L(v)$ induced by the embedding of C in $\mathbb{R}^3$. By assumption, $L(v)$ is 2-connected. Therefore, by [23, Proposition 4.2.6], the faces of $L(v)$ are all bounded by cycles. Furthermore, by Lemma 6.2.3, it follows that all chambers containing v are bounded by a 2-cycle. Since v is arbitrary, we have that for each vertex, its adjacent chambers are bounded by 2-cycles. Since each chamber is adjacent to at least one vertex, it follows that every chamber is bounded by a 2-cycle. □

Observation 6.2.20. *Let C be a simplicial 2-complex embedded in $\mathbb{R}^3$, R be a chamber and $D \subseteq C$ be a subcomplex. Then, the induced embedding of D has a chamber R' containing R.*

Proof. Let R' be the connected component of $\mathbb{R}^3 - D$ that contains R. This selection is possible because $\mathbb{R}^3 - C \subseteq \mathbb{R}^3 - D$. □

Lemma 6.2.21. *Let C be a simplicial 2-complex embedded in $\mathbb{R}^3$ and let f be a face of C. If f lies on a 2-cycle $S \subseteq C$, then f lies on the boundary of exactly two chambers of C, and these are contained in distinct chambers bounded by S.*

Proof. Consider a point in the interior of f. We will prove first that this point lies on the boundary of two chambers. We will consequently show that the whole interior of the face f lies on the boundary of the same two chambers. Finally, the edges that bound f will also be on the boundary of these chambers because a neighbourhood of a boundary point is also a neighbourhood of some interior point of f.

Around every point $x \in \mathring{f}$, we can find an open ball B_x, with centre x, which intersects C only in $\mathring{f}$.

Let us pick a point $x_0 \in \mathring{f}$. Since $D_x \cap C = D_x \cap f$, which is a great disc of D_x, it follows that $D_x - C$ is the union of two open half-balls. Since these two half-balls do not intersect C, they each lie in a chamber bounded by C. Denote these two chambers by R_1 and R_2; they are the only two chambers of C that have x_0 on their boundary and a priori they may coincide.

By Theorem 3.2.22, the 2-cycle S, when considered as a subcomplex of C separates $\mathbb{R}^3$ into two chambers. By Observation 6.2.20, the chambers R_1 and R_2 are subchambers of these two chambers. Hence, R_1 and R_2 are different.

Now consider any other point $x_1 \in \mathring{f}$. Let P be an arc from x_0 to x_1 within $\mathring{f}$. Since P is compact there exist finitely many balls B_x with $x \in P$ covering P. By adding B_{x_0} or B_{x_1} if necessary, we can assume that $B_0 = B_{x_0}$ and $B_n = B_{x_1}$. By induction on n, one can prove that $\cup_{i=0}^{n}(B_{x_i} - S)$ is homeomorphic to the disjoint union of two balls, each ball containing one of the half-balls of $B_0 - S$. Hence, the two half-balls of $B_n - S$ are in the same chamber as the corresponding two half-balls of $B_0 - S$. Now, we can apply the same argument as for x_0 to conclude that x_1 is on the boundary of precisely R_1 and R_2. $\square$

Proof of the 'only if' implication of Theorem 6.2.3. Suppose that C is local, simply connected, locally 2-connected and consider an embedding of C into $\mathbb{R}^3$. By Lemma 6.2.19 the chamber boundaries of C are 2-cycles, so they are elements of $\mathcal{C}$. We shall show that the chamber boundaries generate all the 2-cycles of C and thus they form a sparse spanning set of $\mathcal{C}$ by Lemma 6.2.21. Let S be a 2-cycle of C and let R be its inner chamber. By Lemma 6.2.21, every face f with $\mathring{f} \subseteq R$ lies on the boundaries of exactly two chambers $R_1, R_2 \subseteq R$, and every face of S lies on exactly one such boundary. Hence, the sum in $\mathcal{C}$ of the boundaries of chambers within R is exactly S. Since S is arbitrary, this completes the proof. $\square$

Proof of Theorem 6.2.2. Combine the two implications of Theorem 6.2.3 and then apply Lemma 6.2.4. $\square$

6.3 Mac Lane's embeddability criterion in three dimensions using a Whitney type theorem

6.3.1 Preliminaries

We will use the same definitions and conventions regarding 2-cycles and cycle spaces of simplicial complexes as in Section 6.2. The only difference is that we used $\mathbb{F}_3$ as the underlying field in the previous section, while we will use $\mathbb{F}_2$ here. We use this field, because historically it was the hardest field to prove the result for. With the current machinery though, the proof works in $\mathbb{F}_2$ the same way as in any other finite field. We make this change to advertise this improvement.

Recall the series of definitions Definition 2.5.3 – Definition 2.5.5. They show how we can define a matroid over any finite field using a matrix. Subsequently, given a graph or a 2-complex, they give a construction of a matrix that defines the (dual) matroid of a graph or a 2-complex.

Let us restate Definition 2.5.3 in particular, as it defines cocircuits, which is the main tool used in this section.

Definition 6.3.1. Take an arbitrary prime number p. A p-ary matroid is defined using a matrix by its set of edges and its set of circuits in the following way. Start with the p-ary matrix A. The edges of the matroid are the columns of A. The circuits of the matroid are the sets of edges corresponding to minimal non-empty linearly dependent sets of column vectors. We will denote a matroid obtained from the matrix A in such a way by $M(A)$. The dual matroid $M^*(A)$ of the matroid $M(A)$ can be defined by the same matrix in the following way. The edges of $M^*(A)$ will be again the columns, while the circuits of $M^*(A)$ are the non-empty minimal supports of the row space of A. The circuits of $M^*(A)$ are also called *cocircuits* of $M(A)$.

In Definition 6.3.1, we define the cocircuits of a matroid M as a set of vectors of the matrix that defines M. The vector spaces spanned by this set will be called the *cocircuit space* of the matroid. We can similarly define the *circuit space* of a mattroid.

The main result of this section is as follows.

Theorem 6.3.2. *Let C be a simply connected local 2-dimensional simplicial complex. Then C embeds in $\mathbb{R}^3$ if and only if it has a sparse generating set of its cycle space.*

One of the vital elements in the proof of Theorem 6.3.2 is going to be the following Whitney type theorem in three dimensions.

Theorem 6.3.3. *[15, Theorem 1.2] Let C be a simply connected 2-dimensional simplicial complex whose dual matroid M^* is local. Then C is embeddable in 3-space if and only if M is graphic.*

Tutte's characterisation [66] of graphic matroids yields the following corollary.

Theorem 6.3.4. *[15, Corollary 1.3] Let C be a simply connected simplicial complex whose dual matroid is local.*

Then C is embeddable in 3-space if and only if M has no minor isomorphic to U_2^4, the fano plane, the dual of the fano plane or the dual of either $M(K_5)$ or $M(K_{3,3})$.

A key component of the characterisation in Theorem 6.3.4 is the exclusion of the matroid U_2^4, which is not a binary matroid. This means that the matroid used in Theorem 6.3.3 cannot be defined over $\mathbb{F}_2$. However, Carmesin has proved in [16] that the same result holds for binary matroids as well. Thus, from here on we will assume that Theorem 6.3.3 works over any field.

6.3.2 Proof of Theorem 6.3.2

The biggest hurdle in this section is proving Lemma 6.3.5. Having this lemma together with Theorem 6.3.3 we then complete the proof of Theorem 6.3.2 in a straightforward fashion.

Lemma 6.3.5. *Let M be a binary matroid. Then M is graphic if and only if its cocircuit space has a sparse generating set.*

Proof. For the 'only if' implication, suppose that M is graphic and let $M = M(G) = M(A_G)$, where A_G is a binary matrix. If there is a cutvertex, we can split it without changing the cycle space and thus without changing M, so we can assume that G has no cutvertices. Denote the set of row vectors of A_G by $\mathcal{S}$. The elements of $\mathcal{S}$ are cocircuits of M and correspond to atomic cuts of G. Furthermore, the cuts of G are in bijection with the cocircuits of M. Recall also that the atomic cuts generate all cuts of G over $\mathbb{Z}$ and thus in particular over $\mathbb{F}_2$. Piecing all of these together we obtain that the set of cocircuits $\mathcal{S}$ generates all cocircuits of M and therefore it generates its cocircuit space. Each column of A_G has precisely two non-zero elements since every edge has precisely two endpoints. As $\mathcal{S}$ is the set of row vectors of A_G, it is a sparse set of vectors. Hence, $\mathcal{S}$ is the desired sparse generating set of the cocircuit space of M.

For the 'if' implication suppose that the cocircuit space of M has a sparse generating set $\mathcal{S}$. Construct a graph G with $V(G) = \mathcal{S}$. Connect two distinct vertices in G with one edge for each element in M that the generating cocircuits share and label this edge with the respective element. For each loop in M we can freely add a loop in G, so we can assume for simplicity that there are no loops. Since $\mathcal{S}$ is a sparse generating set, each element of M is shared by exactly two or exactly zero sets, so this construction is well defined. Note that the set of edges incident to the vertex $s \in \mathcal{S}$ are precisely the elements of the cocircuit s in the matroid M. In other words, the atomic cuts of the graph G are the cocircuits of the matroid M. We will denote the atomic cut/cocircuit at the vertex v by c_v.

Sublemma 6.3.6. *The graph G has no cutvertices.*

Proof. Suppose for a contradiction that G has a cutvertex $v^* \in \mathcal{S}$. Let C be the connected component of G, containing v^*. Then, as v^* is a cutvertex, $C - v^*$ is disconnected and has components $C^1, C^2, \ldots, C^k$. The atomic cuts of the vertices of C^1 sum up to the set of edges between v^* and the its neighbours in C^1, which we will denote by $c^1_{v^*}$. Since the atomic cuts are cocircuits it follows that $c^1_{v^*}$ contains a cocircuit. Therefore, the atomic cut at v^* contains a smaller cocircuit, which is contradiction with the fact that c_{v^*} is a cocircuit itself. This contradcition proves the sublemma. □

Sublemma 6.3.7. *The cycle matroid $M(G)$ of G is equal to the matroid M.*

Proof. The matroid $M(G)$ can be defined through a matrix A_G, where the rows of this matrix correspond to characteristic vectors of the atomic cuts of G. As we noted before, the atomic cuts of G correspond to the cocircuits of M. Therefore, A_G is a matrix whose rows are the characteristic vectors of the cocircuits of M. Hence, the two matroids can be defined from the same matrix, or in other words, $M = M(A_G) = M(G)$ as claimed. □

The fact that M is graphic is a direct consequence of Sublemma 6.3.7. This completes the 'if' implication and finishes the proof. □

Proof that Lemma 6.3.5 implies Theorem 6.3.2. Assume that C is a simply connected, local simplicial complex and consider the following properties:

(1) The cycle space of C has a sparse generating set.

(2) C embeds in $\mathbb{R}^3$.

(3) The dual matroid of C is graphic.

(4) The cocircuit space of the dual matroid of C has a sparse generating set.

Theorem 6.3.3 states that (2) $\iff$ (3). Lemma 6.3.5 states that (3) $\iff$ (4). Proving Theorem 6.3.2 is equivalent to proving (1) $\iff$ (2). We have for a given that (2) $\iff$ (3) and (3) $\iff$ (4). Therefore, to prove Theorem 6.3.2, it is enough to show that (1) $\iff$ (4). To confirm the veracity of the last statement, we note that by definition, the cocircuit space of the dual matroid of C is isomorphic to the circuit space of the cycle matroid of C, which is in turn isomorphic to the cycle space of C. Both isomorphisms are isomorphisms between binary vector spaces. From this, the statement (1) $\iff$ (4) and thus the theorem follows. □

Chapter 7

Two additional results regarding outerspatial 2-complexes

In this chapter we are going to recall the definitions of outerspatial and weakly outerspatial 2-complexes and explore their properties. We have two main results about these two types of complexes.

7.1 Introduction

Let us first revisit the most essential notions for this subsection:

Definition 7.1.1. A 2-complex C is *weakly outerspatial* if there is an embedding of C such that the outer chamber is incident to all edges.

Definition 7.1.2. A 2-complex is *outerspatial* if the cone over it is embeddable in $\mathbb{R}^3$.

Recall that in Chapter 3, we considered two equivalent definitions of outerplanarity as candidates for the basis of outerspatiality – one concerning an apex vertex and the other concerning an apex face. There, we saw that attempting to extend these two definitions, we obtain two different properties of 2-complexes. One of them is weakly outerspatial and the other is outerspatial. As the name suggests and Lemma 3.2.21 shows, the latter is the more specific definition and, in this chapter, it is the one we chose to work with.

In Chapter 3 we characterise outerspatial 2-complexes by a set of forbidden minors. Two of the main stepping stones in getting this result are finding the proper definition of outerspatial and finding a closed form of the set of forbidden minor. The purpose of this chapter is to show the progress in these two aspects before coming to the complete result. This is done by showing two results similar to Theorem 3.1.1 but lacking one of these aspects.

The first result is a forbidden minor characterisation of outerspatial 2-complexes, which is derived as a direct consequence of the the 3-dimensional Kuratowski's theorem (Theorem 7.2.1). As such, the proof is simple but lacks the first aspect that we desire – presenting the set of forbidden minors in a closed form.

The second result is a characterisation of weakly outerspatial 2-complexes via properties of their dual graphs, which is stated below as Theorem 7.3.3. In this theorem we prove that a 2-complex is weakly outerspatial if and only if its dual graph exhibits certain tree-like properties. We present two proofs of this result. The first one uses directly properties of weakly outerspatial 2-complexes. The second one first demonstrates what further conditions are necessary for a weakly outerspatial 2-complex to be outerspatial, which is an interesting result in itself. The proofs and ideas are very insightful and relevant to the results in Chapter 3. However, Theorem 7.3.3 itself lacks the second aspect that we desire – using the best definition of outerspatial 2-complexes. This is presented in Section 7.3.

7.2 Forbidden minor characterisation of outerspatial 2-complexes

Let us introduce the following characterisation of embeddability, using a set of forbidden minors, the so-called three dimensional Kuratowski's theorem:

Theorem 7.2.1. [12, *Theorem 1.1.*] *Let C be a simply connected and locally 3-connected 2-dimensional simplicial complex. The following are equivalent.*

- *C embeds in 3-space;*

- *C has no space minor from the finite list $\mathcal{Z}$.*

We will also recall a lemma from Chapter 3 – Lemma 3.5.14.

Lemma 7.2.2. *Consider a 2-complex C that is locally 2-connected. Then the cone over it is locally 3-connected.*

With this we can prove the main result of this subsection.

Theorem 7.2.3. *A simply connected locally 2-connected 2-complex C is outerspatial if and only if it does not have a minor in the class $\mathcal{Z}'$, where*

$$\mathcal{Z}' = \{Z - u \mid Z \in \mathcal{Z} \text{ and } u \text{ is vertex of } Z\}$$

Proof. Firstly, we will prove the following.

Sublemma 7.2.4. *The 2-complex C has a minor in $\mathcal{Z}'$ if and only if its cone $\widehat{C}$ has a minor in $\mathcal{Z}$*

Proof. For the 'only if' implication suppose that C has a minor C' that is in $\mathcal{Z}'$. By the definition of $\mathcal{Z}'$ there exists a 2-complex C'_u in $\mathcal{Z}$ such that $C' = C'_u - u$. The cone $\widehat{C'}$ is a minor of $\widehat{C}$ since C' is a minor of C. Moreover, $\widehat{C'}$ contains C'_u as a subcomplex since $C'_u - u = \widehat{C'} - t = C'$ and $C'_u \restriction u \subseteq \widehat{C'} \restriction t$, where t is the top of the cone $\widehat{C'}$. Joining the results from the last two sentences, we obtain that C'_u is a minor of $\widehat{C}$. Thus, given a 2-complex C with a minor in $\mathcal{Z}$, we obtained a 2-complex C'_u that is a minor of $\widehat{C}$ and is in the set $\mathcal{Z}$.

For the 'if' implication, let t be the top of the cone $\widehat{C}$ and suppose that $\widehat{C}$ has a minor C^* in $\mathcal{Z}$. In the text below, when we refer to t in C^*, we refer to all copies of t as well as all equivalence classes of vertices containing t and its copies that might appear due to the space minor operations applied to obtain C^* from $\widehat{C}$.

If $t \notin C^*$, then C^* is also a minor of C. Hence, for an arbitrary $u \in V(C^*)$, we have that $C^* - u$ is a minor of C and is in $\mathcal{Z}'$.

If $t \in C^*$, let C^*_t be the 2-complex obtained from C^* by removing t. Now, note that all space minor operations done to obtain C^*_t from $\widehat{C}$ either disappear with the deletion of t or can be viewed as operations restricted to C. This is easy to see if t is not involved in any edge contractions. If t is involved in edge contractions, let the first such contraction be on the edge tx, where x is a vertex of C. We can consider all further operations on t as operations on x, so restricted to C. Hence, C^*_t is a minor of C and is an element of $\mathcal{Z}'$ because $t \in C^* \in \mathcal{Z}$.

Thus, in both cases we obtained a 2-complex that is a minor of C and is in $\mathcal{Z}'$, which gives us the 'if' implication. $\square$

Consider the following statements.

(1) C is outerspatial

(2) $\hat{C}$ is embeddable in $\mathbb{R}^3$

(3) C has no minor in the class $\mathcal{Z}'$.

(4) $\hat{C}$ has no minor in the class $\mathcal{Z}$.

The statement (1) $\iff$ (2) holds by definition. The statement (2) $\iff$ (4) is true according to Theorem 7.2.1. The contrapositive of Lemma 7.2.2 is (3) $\iff$ (4). Therefore, statement (1) $\iff$ (3) also holds true. But this is equivalent to our result, so we are done. $\square$

7.3 Characterisation of weakly outerspatial 2- complexes via their dual graphs

Consider a 2-complex C together with an embedding of it in $\mathbb{R}^3$. Let G_C be the *dual graph* of this embedding. Here, we will use a similar definition for the dual graph of an embedding as in Chapter 3 which is as follows.

Definition 7.3.1. The *dual graph* of the embedding of a 2-complex C is constructed in the following way. We have a vertex for each chamber apart from the outer chamber and two vertices are connected by one edge for each face their respective chambers share.

Let us start with a definition essential for this section.

Definition 7.3.2. Call a 2-complex *treelike*, if it is built inductively by starting from a sphere and at each step gluing a new sphere in the outer chamber at exactly one face from the complex obtained so far.

Now we are ready to state the main result of this subsection.

Theorem 7.3.3. *Consider a* 2*-complex* C *which is locally* 2*-connected, simply connected and local. Then* C *is weakly outerspatial if and only if there is an embedding of* C *in* $\mathbb{R}^3$ *for which* G_C *is a tree.*

Proof. Since C is local and locally 2-connected, we can use Lemma 6.2.19 for both implications to obtain that every chamber of C is bounded by a 2-cycle. In particular, since C is simply connected, we know that every chamber is bounded by a 2-sphere.

Suppose that we have an embedding for which G_C is a tree. Since every chamber of C is bounded by a 2-sphere and has a dual graph a tree, it easily follows by induction that it is treelike. Hence, the following sublemma is sufficient to prove the 'if' implication.

Sublemma 7.3.4. *A treelike* 2*-complex is weakly outerspatial.*

Proof. We prove this by induction on the number of gluing steps. For zero steps we have a sphere, which is clearly weakly outerspatial. Consider a treelike 2-complex that is weakly outerspatial and glue one more sphere to it. Since the gluing is in the outer chamber all new edges are on it. All the old edges which were not glued on are unchanged and are still on the outer chamber by the inductive hypothesis. Consider an arbitrary edge e that was glued on. Its copy on the new sphere lies on precisely two faces – one of them is the gluing face, call the other face f. After the gluing, e lies on f. Since f is on the outer chamber, it follows that e also is. Since e was an arbitrary gluing edge, we have covered all edges and proved that they are on the outer chamber. □

For the ‘only if’ implication consider a weakly outerspatial 2-complex C. Suppose for a contradiction that there exists an embedding of C such that G_C is not a tree. Then it should have a cycle. Since all chambers of C are bounded by a 2-sphere, a cycle in G_C corresponds to a chain of spheres each glued at a face to the previous one and the last sphere glued to the first. Let T be the subcomplex formed by restricting C to the spheres in this chain. Its boundary is homeomorphic to a sphere glued to itself at two discs with disjoint interiors. If the two discs are disjoint, we get a torus. If the two discs intersect at their boundary, we get a sphere. In the case where T is a sphere, the shared boundary of the two gluing discs correspond to a path in the 1-skeleton of T which is on the inside of the sphere bounding T. Any edge of this path will not be in the outer chamber of C, therefore C is not weakly outerspatial in this case. In the case where T is a torus, all edges of T need to be on the boundary of the outer chamber of C, so the whole 1-skeleton of T needs to be on it. However, the boundary of the outer chamber of C is a 2-sphere and the 1-skeleton of a torus cannot be embedded on a 2-sphere. This contradiction finishes the proof. □

An example of a 2-complex that is weakly outerspatial but does not have a dual graph which is a tree – $3 \times 3 \times 1$ cubic grid with the center top face removed. This example shows that locality is required for the ’if’ implication of Theorem 7.3.3. If we remove also the bottom center face we obtain an example that shows that simply connectedness is required as well.

Remark 7.3.5. Note that first proof of the ‘if’ implication of Theorem 7.3.3 has quite similar ideas to the proof of Lemma 3.5.7, while the first proof of the ‘only if’ implication has quite similar ideas to the proof of Proposition 3.2.34. Moreover, those two results are used in the second proof as an ingredient to the characterisation of outerspatial 2-complexes in terms of weakly outerspatial 2-complexes. Hence, the ideas presented in this section are not only first steps towards Chapter 3, but also have had some relevant use there.

Bibliography

[1] Wolfram alpha auxilary calculation. `https://shorturl.at/vLQV4`.

[2] N. Alon and J. H. Spencer. *The probabilistic method*. John Wiley & Sons, 2016.

[3] P. Angelini, G. Di Battista, M. Kaufmann, T. Mchedlidze, V. Roselli, and C. & Squarcella. Small point sets for simply-nested planar graphs. In *International Symposium on Graph Drawing*, pages 75–85. Springer, 2011.

[4] K. Appel and W. Haken. *Every planar map is four colorable*, volume 98. American Mathematical Soc., 1989.

[5] A. Asadi, L. Postle, and R. Thomas. Sub-exponentially many 3-colorings of triangle-free planar graphs. *Electronic Notes in Discrete Mathematics*, 34:81–87, 2009.

[6] E. Asplund and B. Grünbaum. On a Coloring problem. *Mathematica Scandinavica*, 8:181–188, 1960.

[7] K. J. Böröczky and A. Sagmeister. The isodiametric problem on the sphere and in the hyperbolic space. *Acta Mathematica Hungarica*, 160(1):13–32, 2020.

[8] D. L. Boutin, E. Gethner, and T. Sulanke. Thickness-two graphs part one: New nine-critical graphs, permuted layer graphs, and catlin's graphs. *Journal of Graph Theory*, 57(3):198–214, 2008.

[9] H. Breu, H. B. Hunt, M. V. Marathe, S. S. Ravi, and D. J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25(2):59–68, 1995.

[10] N. Buduma and N. Locascio. *Fundamentals of deep learning: Designing next-generation machine intelligence algorithms*. O'Reilly Media, Inc., 2017.

[11] J. P. Burling. On coloring problems of families of prototypes (PhD thesis). *University of Colorado, Boulder*, 1965.

[12] J. Carmesin. Embedding simply connected 2-complexes in 3-space – I. A Kuratowski-type characterisation. *arXiv preprint arXiv:1709.04642*, 2017.

[13] J. Carmesin. Embedding simply connected 2-complexes in 3-space – II. Rotation systems. *arXiv preprint arXiv:1709.04643*, 2017.

[14] J. Carmesin. Embedding simply connected 2-complexes in 3-space – III. Constraint minors. *arXiv preprint arXiv:1709.04645*, 2017.

[15] J. Carmesin. Embedding simply connected 2-complexes in 3-space – IV. Dual matroids. *arXiv preprint arXiv:1709.04652*, 2017.

[16] J. Carmesin. Dual matroids of 2-complexes–revisited. *arXiv preprint arXiv:2210.15594*, 2022.

[17] J. Carmesin. Embedding simply connected 2-complexes in 3-space. *arXiv preprint arXiv:2309.15504*, 2023.

[18] J. Carmesin and T. Mihaylov. Outerspatial 2-complexes: Extending the class of outerplanar graphs to three dimensions. *arXiv preprint arXiv:2103.15404*, 2021.

[19] J. Carmesin, E. Nevinson, and B. Saunders. A characterisation of 3-colourable 3-dimensional triangulations. *arXiv preprint arXiv:2204.02858*, 2022.

[20] G. Chartrand and F. Harary. Planar permutation graphs. In *Annales de l'IHP Probabilités et statistiques*, volume 3, pages 433–438, 1967.

[21] A. D. N. J. De Grey. The chromatic number of the plane is at least 5. *arXiv preprint arXiv:1804.02385*, 2018.

[22] C. P. Dettmann and O. Georgiou. Random geometric graphs with general connection functions. *Physical Review E*, 93(3):032313, 2016.

[23] R. Diestel. Graph theory. *Graduate texts in mathematics*, 173, 2000.

[24] D. Eppstein and B. Reed. Finding maximal sets of laminar 3-separators in planar graphs in linear time. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 589–605. SIAM, 2019.

[25] P. Erdős. On sets of distances of n points. *The American Mathematical Monthly*, 53(5):248–250, 1946.

[26] G. Exoo and D. Ismailescu. The chromatic number of the plane is at least 5: a new proof. *Discrete & Computational Geometry*, 64(1):216–226, 2020.

[27] K. J. Falconer and J. M. Marstrand. Plane sets with positive density at infinity contain all large distances. *Bulletin of the London Mathematical Society*, 18(5):471–474, 1986.

[28] S. Fiorini, N. Hardy, B. Reed, and A. Vetta. Approximate min-max relations for odd cycles in planar graphs. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 35–50. Springer, 2005.

[29] H. Fleischner, D. Geller, and F. Harary. Outerplanar graphs and weak duals. *The Journal of the Indian Mathematical Society*, 38(1-4):215–219, 1974.

[30] P. Frankl and R. M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1:357–368, 1981.

[31] F. Frati and M. Patrignani. A note on minimum-area straight-line drawings of planar graphs. In *International Symposium on Graph Drawing*, pages 339–344. Springer, 2007.

[32] S. Galhotra, A. Mazumdar, S. Pal, and B. Saha. Connectivity in random annulus graphs and the geometric block model. *arXiv preprint arXiv:1804.05013*, 2018.

[33] M. Gardner. Mathematical games. *Scientific american*, 222(6):132–140, 1970.

[34] E. N. Gilbert. Random plane networks. *Journal of the society for industrial and applied mathematics*, 9(4):533–543, 1961.

[35] M. Goemans and D. Williamson. Primal-dual approximation algorithms for feedback problems in planar graphs. *Combinatorica*, 18(1):37–59, 1998.

[36] P. Golovach, M. Kamiński, D. Paulusma, and D. Thilikos. Induced packing of odd cycles in a planar graph. In *International Symposium on Algorithms and Computation*, pages 514–523. Springer, 2009.

[37] A. Gräf. Coloring and recognizing special graph classes. *Bulletin of the European Association for Theoretical Computer Science*, 56:303–303, 1995.

[38] H. Grötzch. Zur theorie der diskreten gebilde VII: ein dreifarbensatz für dreikreisfreie netze auf der kugel. *Wiss. Z. Martin Luther Univ. Halle Wittenberg, Math-Nat.(1958/1959)*, pages 109–120, 1958.

[39] A. Hatcher. *Algebraic topology.* Cambridge University Press, 2005.

[40] P. J. Heawood. On the four-colour map theorem. *Quarterly Journal of Pure and Applied Mathematics*, 29:270–285, 1898.

[41] P. J. Heawood. Map-colour theorem. *Proceedings of the London Mathematical Society*, 2(1):161–175, 1949.

[42] G. A. Kabatyanskii and V. Levenshtein. On bounds for packings on a sphere and in space. *Problemy Peredachi Informatsii*, 14(1):3–25, 1978.

[43] C. Kidd, S. T. Piantadosi, and R. N. Aslin. The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one*, 7(5):e36399, 2012.

[44] M. Krug and D. Wagner. Minimizing the area for planar straight-line grid drawings. In *International Symposium on Graph Drawing*, pages 207–212. Springer, 2007.

[45] C. Kuratowski. Sur le probleme des courbes gauches en topologie. *Fundamenta mathematicae*, 15(1):271–283, 1930.

[46] M. Lackenby. Topology and groups. University of Oxford lecture notes, 2017.

[47] S. Lefschetz. Planar graphs and related topics. *Proceedings of the National Academy of Sciences*, 54(6):1763–1765, 1965.

[48] S. Li. Concise formulas for the area and volume of a hyperspherical cap. *Asian Journal of Mathematics and Statistics*, 4(1):66–70, 2011.

[49] S. Mac Lane. *A combinatorial condition for planar graphs.* Seminarium Matemathicum, 1936.

[50] H. Maehara. Dispersed points and geometric embedding of complete bipartite graphs. *Discrete & computational geometry*, 6(1):57–67, 1991.

[51] E. Malesińska, S. Piskorz, and G. Weißenfels. *On the approximability of coloring problems for disk graphs.* Bereich Musikinformatik, Musikwiss. Inst., 1996.

[52] B. Mohar and C. Thomassen. *Graphs on Surfaces.* John Hopkins, 2001.

[53] M. J. P. Peeters. *On coloring j–unit sphere graphs*, volume FEW 512 of *Research memorandum.* Department of Economics, Tilburg University, 1991.

[54] M. Penrose. *Random geometric graphs*, volume 5. Oxford University Press, 2003.

[55] M. D. Penrose. Connectivity of soft random geometric graphs. *The Annals of Applied Probability*, 26(2), Apr 2016.

[56] J. Reiterman, V. Rödl, and E. Šinajová. Geometrical embeddings of graphs. *Discrete Mathematics*, 74(3):291–319, 1989.

[57] Neil Robertson and Paul D Seymour. Graph minors. xx. wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.

[58] C. A. Rogers. Covering a sphere with spheres. *Mathematika*, 10(2):157–164, 1963.

[59] W. Schmaltz. The Jordan-Brouwer separation theorem. `http://www.math.uchicago.edu/~may/VIGRE/VIGRE2009/REUPapers/Schmaltz.pdf`, 2009.

[60] E. Schmidt. Die Brunn-Minkowskische Ungleichung und ihr Spiegelbild sowie die isoperimetrische Eigenschaft der Kugel in der euklidischen und nichteuklidischen Geometrie. I. *Mathematische Nachrichten*, 1(2-3):81–157, 1948.

[61] E. Schmidt. Die Brunn-Minkowskische Ungleichung und ihr Spiegelbild sowie die isoperimetrische Eigenschaft der Kugel in der euklidischen und nichteuklidischen Geometrie. II. *Mathematische Nachrichten*, 2(3-4):171–244, 1949.

[62] A. Scott and P. Seymour. A survey of χ-boundedness. *Journal of Graph Theory*, 95(3):473–504, 2020.

[63] W. Song, T. Di Matteo, and T. Aste. Nested hierarchies in planar graphs. *Discrete Applied Mathematics*, 159(17):2135–2146, 2011.

[64] J. Spencer, E. Szemerédi, and W. T. Trotter. Unit distances in the Euclidean plane. In *Graph theory and combinatorics*, pages 294–304. Academic Press, 1984.

[65] G. F. Töth and W. Kuperberg. Packing and covering with convex sets. In *Handbook of Convex Geometry*, pages 799–860. Elsevier, 1993.

[66] W.T. Tutte. Lectures on matroids. *Journal of Research of the National Bureau of Standards B*, 69:1–2, 1964.

[67] J-L. Verger-Gaugry. Covering a ball with smaller equal balls in n–dimensional space. *Discrete & Computational Geometry*, 33(1):143–155, 2005.

[68] B. Waxman. Routing of multipoint Connections. *IEEE Journal on selected areas in communications*, 6(9):1617Y1622, 1988.

[69] H. Whitney. Non-separable and planar graphs. In *Hassler Whitney Collected Papers*, pages 37–59. Springer, 1992.