A study of automorphism groups of spheres via intersection graphs

by

Joshua Reynolds

under the supervision of Agelos Georgakopoulos

MA5P1 Dissertation

Submitted in partial fulfilment of the requirements for the degree of

Master of Science in Mathematics

Warwick Mathematics Institute

September 2025

THE UNIVERSITY OF WARWICK

Contents

Acknowledgements

I would like to thank Professor Agelos Georgakopoulos for his steadfast supervision throughout this project, for suggesting the main problem of this dissertation and for providing feedback on my many drafts. I would also like to thank my family for their support throughout the course of this year.

Declaration

This dissertation is submitted to the University of Warwick in support of my application for the degree of Master of Science in Mathematics. It has been composed by myself and has not been submitted in any previous application for any degree. The work presented was carried out by the author except where otherwise specified, and any featured illustrations were created for this dissertation by the author.

Abstract

By answering a question of A. Georgakopoulos, we generalise a result stating that the automorphism group of the intersection graph of the chords of the 1-sphere is isomorphic to the group of homeomorphisms of the 1-sphere to higher dimensions. We subsequently investigate the automorphism groups of intersection graphs of more general classes of curves on the 2-sphere and demonstrate that the class of Jordan curves is sufficient to recover the entire automorphism group of the 2-sphere. We conclude by providing examples of an important invariant, which we term the region graph, used throughout this dissertation. We also examine the properties of such graphs on the classes of curves investigated in earlier chapters.

Chapter 1

Introduction

In a recent paper, A. Georgakopoulos has provided a natural graph $\mathcal{C}$, such that $\mathrm{Aut}(\mathcal{C})$ is isomorphic to $\mathrm{Aut}(\mathbb{S}^1)$, the group of homeomorphisms of the 1-sphere [1]. The study of $\mathrm{Aut}(\mathbb{S}^1)$ is a classical topic, blending many areas of study; we refer the reader to several surveys and monographs of the area for further reading [2][3][4]. The graph $\mathcal{C}$ is termed the *circle graph*; the graph whose vertex set consists of the chords of $\mathbb{S}^1$ and where two chords share an edge if and only if they intersect on or within $\mathbb{S}^1$. Such a graph is an *intersection graph*, a graph whose vertices are objects, and edges exist only between those objects which have non-empty intersection. We again refer the reader to several summary texts which have been written on the topic of intersection graphs [5][6]. It is natural to consider higher-dimensional analogues of this problem. In particular, we can consider the *sphere graph*, $\mathcal{C}^d$, of the d-sphere, where $\mathcal{C}^1 = \mathcal{C}$. The vertices of such a graph are the chords of $\mathbb{S}^d$, where a chord is the non-empty intersection of a non-tangential d-plane with the closed ball $\mathbb{D}^{d+1}$ in $\mathbb{R}^{d+1}$. Georgakopoulos proposed the following question in the aforementioned paper, which this dissertation seeks to answer:

Question 1 *Is* $\mathrm{Aut}(\mathbb{S}^d)$ *isomorphic to* $\mathrm{Aut}(\mathcal{C}^d)$ *for every* $d > 1$*?*

We show that this is false for $d \geq 2$ and that $\mathrm{Aut}(\mathcal{C}^d)$ is in fact isomorphic to the group of Möbius transformations of the d-sphere, which we denote $\mathrm{M\ddot{o}b}(\mathbb{S}^d)$, for $d \geq 2$.

Many readers will be familiar with Möbius transformations of $\widehat{\mathbb{C}}$, but may be less familiar with their higher-dimensional relatives. We introduce the reader to such transformations of $\widehat{\mathbb{R}}^d$ in Chapter 2, blending material from several texts to provide a solid overview of these maps. We proceed by demonstrating how such maps may be transferred to the d-sphere via stereographic projection, showing that such projections are in fact Möbius transformations themselves. We encounter various properties of Möbius transformations that prove useful in answering Question 1 and end by showing a characterisation of Möbius transformations via *chordal spheres*. A chordal sphere of $\widehat{\mathbb{R}}^d$ is a $(d-1)$-sphere or d-plane; under stereographic projection, these are sent to chords of $\mathbb{S}^{d+1}$, which makes clear their name.

In all later chapters most material will constitute original work, except where otherwise credited.

In Chapter 3 we prove our stated answer to Question 1. We follow a method of proof inspired by that of the original result and introduce several useful invariants under automorphisms of $\mathcal{C}^d$, including what we term *region graphs*. Such invariants heavily restrict the possible symmetries of the sphere graphs and allow us to prove many unexpected properties of automorphisms of $\mathcal{C}^d$. For example, such automorphisms distinguish between types of intersection, despite this information not being explicitly encoded in $\mathcal{C}^d$. We further show that such invariants can be extended to more general classes of simple closed curve.

In order to construct an isomorphism between the two groups, much like Georgakopoulos, we build *boundary cliques* - the induced subgraphs on those vertices of $\mathcal{C}^d$ that contain a given point on $\mathbb{S}^d$. We demonstrate that the action of an automorphism $h \in \mathcal{C}^d$ on boundary cliques provides a well-defined map of points on $\mathbb{S}^d$ and thereby construct a homeomorphism. We end by demonstrating a homeomorphism constructed in this manner must satisfy a characterisation of Möbius transformations proved in Lemmas 2.5 and 2.6. We construct our method of proof to also cover the case of $d = 1$, and hence we are able to provide an alternative method of proof to that of Georgakopoulos. We consequently demonstrate that the methods used in proving this result provide a purely combinatorial way to extend the action of the Möbius group to the interior of $\mathbb{S}^d$. We conclude the chapter by making several observations on the structure of $\mathcal{C}^d$ and its automorphism group.

In Chapter 4 we look to find a natural graph whose automorphism group is isomorphic to $\mathrm{Aut}(\mathbb{S}^2)$, to recover a result analogous to that of Georgakopoulos in higher dimensions. We prove that the intersection graph $\mathcal{J}$ whose vertices are the Jordan curves on $\mathbb{S}^2$ is such a graph, by extending methods of proof developed through the prior chapter. A Jordan curve is the image of an injective and continuous map from $\mathbb{S}^1$ to $\mathbb{S}^2$. We discuss the pathological nature of many such curves and give a summary of several crucial results which allow us to tackle the problems these curves present.

In Chapter 5 we focus on the intersection graphs $\mathcal{S}^k$ of the C^k Jordan curves on $\mathbb{S}^2$, where $k \in \mathbb{N}^{\geq 0} \cup \{\infty\}$. We provide a brief summary of smooth manifolds to justify our analysis of such classes of simple closed curves. We show that the subgroup of $\mathrm{Aut}(\mathbb{S}^2)$ isomorphic to $\mathrm{Aut}(\mathcal{S}^k)$ is in fact larger than the subgroup of C^k homeomorphisms of the 2-sphere, by augmenting a recent proof of Le Roux and Wolff [7]. In doing so we demonstrate that there are non-C^k homeomorphisms which preserve all simple closed C^k curves.

Finally, in Chapter 6 we explore region graphs further, demonstrating various properties of these graphs in different dimensions. Highlighting some unexpected cases of region graphs that arise when considering curves in $\mathcal{J}$ or $\mathcal{S}^k$, we suggest that the case for circular curves on $\mathbb{S}^2$ is the most tractable. We classify all finite region graphs on two curves in $\mathcal{J}$ and show that such region graphs are perfect. By this classification we demonstrate quite how complex region graphs may become, even on small numbers of curves, when considering more general classes of curve. We demonstrate that for any finite number of Jordan curves on $\mathbb{S}^2$, there are planar region graphs, and provide examples of region graphs that are not perfect. We classify all region graphs on $0 \leq n \leq 3$ circles on $\mathbb{S}^2$, demonstrating the difficulty of such a task for larger collections. We then define a graph class $Circle_k$, which contains all region graphs on collections of circles with maximum degree $k \in \mathbb{N}^{\geq 0}$. We consider the degree of a circle to be its degree within the corresponding region graph. We end by characterising $Circle_k$ for $k = 0, 1, 2$ and demonstrate an iterative process for constructing any graph in $Circle_k$ from the single vertex graph.

Chapter 2

General Möbius transformations and their characterisations

In this chapter we focus on the d-dimensional sphere $\mathbb{S}^d \subset \mathbb{R}^{d+1}$ with the Euclidean metric. We remind the reader of the *extended reals*, $\widehat{\mathbb{R}}^d := \mathbb{R}^d \cup \{\infty\}$ and demonstrate that they are homeomorphic to $\mathbb{S}^d$ via *stereographic projection*. We introduce the notions of Möbius transformations, the maps $\widehat{\mathbb{R}}^d \to \widehat{\mathbb{R}}^d$ generated by reflections in chordal spheres. In particular we focus on certain properties and characterisations of these maps which we make use of in later chapters. The content of this chapter largely follows Chapter 3 of *The Geometry of Discrete Groups* by Beardon [8], Chapter 3 of *An Introduction to the Theory of Higher-Dimensional Quasiconformal Mappings* by Gehring, Martin and Palka [9] and *Inversive Geometry* by J.B. Wilker [10].

2.1 Reflections, inversions and chordal spheres

The core building blocks of Möbius transformations in higher dimensions are *reflections* in spheres and planes.

Definition 2.1 (Reflections) We denote by $\mathbb{S}^{d-1}(a,r) = \{y \in \mathbb{R}^d \colon |y-a| = r\}$ the sphere with centre a and radius r. We will denote the unit sphere in $\mathbb{R}^d$ dimensions by $\mathbb{S}^{d-1}$. We define the closed disc with centre a and radius r as $\mathbb{D}^d(a,r) = \{y \in \mathbb{R}^d \colon |y-a| \leq r\}$, and the analogous open disc as the interior of the closed disc. The *reflection* (or *inversion*) in $\mathbb{S}^{d-1}(a,r)$ of a point $x \in \mathbb{R}^d$ is defined as:

$$R(x) = a + \left(\frac{r}{|x-a|}\right)^2 (x-a) \tag{2.1}$$

In the special case of $\mathbb{S}^{d-1}$, $R(x) = x/|x|^2$. We let $x^* = x/|x|^2$. Ergo, we can reformulate equation 2.1 as:

$$R(x) = a + r^2(x-a)^*, \tag{2.2}$$

where we define $R(\infty) = a$ and $R(a) = \infty$ to define an inversion on the whole of $\widehat{\mathbb{R}}^d$.

The reflection of a point $x \in \mathbb{R}^d$ in the plane $P(a,t) := \{y \in \mathbb{R}^d \colon \langle y, a\rangle = t\}$ where $\langle \cdot, \cdot \rangle$ is the standard dot product is defined as:

$$R(x) = x - 2(\langle x, a\rangle - t)a^*, \tag{2.3}$$

and we define $R(\infty) = \infty$ to define a reflection on $\widehat{\mathbb{R}}^d$.

In both cases a quick check shows that $R^2(x) = x$ and so R is 1-1 on $\widehat{\mathbb{R}}^d$, and that $R(x) = x$ if and only if $x \in \mathbb{S}^{d-1}(a,r)$ or $x \in P(a,t)$ respectively. Thus both of these maps satisfy the properties we would expect from a reflection. Geometrically, an inversion in a sphere can be thought of as mapping a point $x \in \widehat{\mathbb{R}}^d \setminus \{a\}$ to the point $R(x)$ lying on the Euclidean ray from a through x satisfying $|R(x) - a| \cdot |x - a| = r^2$. It interchanges the open disc $\mathrm{int}(\mathbb{D}^d(a,r))$ and $\mathbb{R}^d \setminus \mathbb{D}^d(a,r)$ while preserving $\mathbb{S}^{d-1}(a,r)$. A reflection in a plane $P(a,t)$ can be thought of as taking a point $x \in \widehat{\mathbb{R}}^d$ and mapping it to the point $R(x)$ such that $P(a,t)$ forms the perpendicular bisector of the line between the two points. Again, it interchanges the two partitions of $\mathbb{R}^d$. We will refer to reflections in spheres as inversions, and reserve the term "reflection" for those in planes.

These may seem like fundamentally different operations, however we can unify these two methods of reflection via *chordal spheres*.

Definition 2.2 (Chordal spheres) A subset Σ in $\widehat{\mathbb{R}}^d$ is a *chordal sphere* if it is a Euclidean sphere in $\mathbb{R}^d$ or if $\Sigma = P \cup \{\infty\}$ where P is a $(d-1)$-dimensional hyperplane.

In general, a chordal sphere Σ can be expressed as:

$$\Sigma \colon a|x|^2 - 2\langle b, x\rangle + c = 0, \tag{2.4}$$

where $a, c \in \mathbb{R}$, $b \in \mathbb{R}^d$ and $ac < |b|^2$. We also require that ∞ is counted as a solution if and only if $a = 0$.

One should think of the chordal spheres through ∞ as those infinitely large spheres. The name "chordal" is certainly suggestive, and we will see why this name is suitable in §2.2. Let us then denote the reflection in a particular chordal sphere as R_Σ. Having done so, we are now ready to define the Möbius transformations of $\widehat{\mathbb{R}}^d$.

Definition 2.3 (Möbius transformations) A Möbius transformation is a finite composition of inversions in spheres and reflections in planes. Equivalently, it is a finite composition of reflections R_Σ in chordal spheres.

The composition of two Möbius transformations is again a Möbius transformation. For $\phi = R_1 \circ R_2 \circ \cdots \circ R_m$, $\phi^{-1} = R_m \circ R_{m-1} \circ \cdots \circ R_1$. Finally $\mathrm{id} = R \circ R$ and so the identity is a Möbius transformation. This motivates the following:

Definition 2.4 (Möbius Group) The Möbius group Möb(d) is the group of all Möbius transformations $\widehat{\mathbb{R}}^d \to \widehat{\mathbb{R}}^d$.

We now make a quick comment on the structure of Möb(d).

Remark Let us define $O(1,d)$ as the collection of $(d+1) \times (d+1)$ invertible matrices that preserve the quadratic form $q(x,x) = x_0^2 - (x_1^2 + \cdots + x_d^2)$, where $x \in \mathbb{R}^{d+1}$. That is to say, $q(x,x) = q(xA, xA)$ for any $A \in O(1,d)$. These matrices form a group under matrix multiplication, and one may show that the collection of matrices $A \in O(1,d)$ with $a_{00} > 0$ forms a subgroup, denoted $O^+(1,d)$ - in fact $O(1,d)$ is a special example of a generalised orthogonal group [11]. We will omit proof of these facts, but proofs may be found in both [8] and [9]. Indeed, both books proceed to show that Möb(d) is isomorphic as a topological group to $O^+(1, d+1)$, proof of which we shall also omit.

We now provide several examples of Möbius transformations.

2.1.1 Examples of Möbius transformations

Firstly, all dilations of $\widehat{\mathbb{R}}^d$ are Möbius. For $f = \lambda x$, $f = R \circ R_0$ where R_0 is the inversion in $\mathbb{S}^{d-1}$, and R is the inversion in $\mathbb{S}^{d-1}(0, \sqrt{\lambda})$.

Secondly, all translations of $\mathbb{R}^d$ are Möbius. If we have $f(x) = x + a$, then we have $f = R_1 \circ R_2$, where R_1 is the reflection in the plane $P_1\colon \langle a, x\rangle = 0$ and R_2 is the reflection in $P_2\colon \langle a, x\rangle = |a|^2/2$.

Both of the above can be proven through simple calculations using the definitions of reflections and inversions.

Thirdly, and less obviously, all orthogonal linear transformations of $\mathbb{R}^d$ are Möbius, which we shall now prove. We denote the orthogonal group of dimension d as $O(d)$.

Lemma 2.1. *Any $U \in O(d)$ can be represented as a composition of d or fewer reflections in $(d-1)$-planes that pass through the origin.*

Proof. If U is the identity, then $U = R \circ R$ for any such reflection, and so we assume $U \neq I$. We now provide an iterative construction for transformations $V_1, V_2, \ldots, V_d$ in $O(d)$, where each V_i is either the identity I or a reflection as described in the statement of the lemma. The V_i are constructed so that for $k = 1, 2, \ldots, d$ the transformation $U_k = V_k V_{k-1} \ldots V_1 U$ fixes the vectors $e_1, e_2, \ldots, e_k$. Given this, we may then conclude that $U_d = I$ and so $U = V_1 V_2 \ldots V_d$, where we recall that $V_i^{-1} = V_i$. This is exactly a composition of d or fewer reflections in $(d-1)$-planes through the origin.

Let us now describe the construction of such reflections. Let $b_1 = U(e_1) - e_1$. If $b_1 = 0$ then take $V_1 = I$ otherwise let V_1 be the reflection in $P_1 : \langle b_1, x \rangle = 0$. Then in the second case, $U(e_1) + e_1 \in P_1$ as

$$\langle b_1, U(e_1) + e_1 \rangle = \langle U(e_1) - e_1, U(e_1) + e_1 \rangle = |U(e_1)|^2 - |e_1|^2 = 1 - 1 = 0.$$

Thus $U_1 = V_1 U$ fixes e_1 in both cases; if $b_1 = 0$ then $U_1(e_1) = U(e_1) = e_1$. If instead $b_1 \neq 0$, then

$$\begin{aligned} U_1(e_1) = V_1(U(e_1)) &= V_1\left(\frac{U(e_1) + e_1 + b_1}{2}\right) \\ &= \frac{1}{2}V_1(U(e_1) + e_1) + \frac{1}{2}V_1(b_1) = \frac{1}{2}(U(e_1) + e_1) - \frac{1}{2}b_1 = e_1. \end{aligned}$$

Now let us assume $k < d$ and assume we have constructed $V_1, V_2, \ldots, V_k$ such that U_k fixes $e_1, e_2, \ldots, e_k$. Let us now construct V_{k+1} in much the same way we constructed V_1. Let $b_{k+1} = U_k(e_{k+1}) - e_{k+1}$ and let $V_{k+1} = I$ if $b_{k+1} = 0$ otherwise let V_{k+1} be the reflection in $P_{k+1} : \langle b_{k+1}, x \rangle = 0$. If $b_{k+1} = 0$ then it is clear to see that $U_{k+1} = V_{k+1} U_k$ fixes e_{k+1} as before and also fixes e_i for $1 \leq i \leq k$ as $U_{k+1} = U_k$ in this case. If instead $b_{k+1} \neq 0$ then as U_m fixes e_i and that

$$\langle b_{k+1}, e_i \rangle = \langle U_k(e_{k+1}) - e_{k+1}, e_i \rangle = \langle U_k(e_{k+1}), e_i \rangle \quad = \langle U_k(e_{k+1}), U_k(e_i) \rangle = \langle e_{k+1}, e_i \rangle = 0,$$

we see that V_{k+1} fixes e_i. We can therefore continue until we reach the desired result. $\square$

Finally, a *similarity* is a Möbius transformation that does not involve inversion and so preserves distances up to scaling. That is, $\phi\colon \mathbb{R}^d \to \mathbb{R}^d$ is a similarity if and only if $|\phi(x) - \phi(y)| = k|x - y|$. In general, a similarity is of the form $\phi = kA + b$, where $k \in \mathbb{R}$, $A \in O(d)$ and $b \in \mathbb{R}^d$. These include all translations, dilations and linear orthogonal transformations.

In fact, we can see from these definitions that the group of Möbius transformations of $\mathbb{R}^d$ will be generated by the similarities and R_0, the inversion in $\mathbb{S}^{d-1}$. In particular, suppose we have an inversion R in $\mathbb{S}^{d-1}(a,r)$. Then $R = R_1 \circ R_2 \circ g \circ R_0 \circ R_2 \circ R_1$, where $g(x) = r^2 x$. So any inversion can be expressed as the inversion in the unit sphere composed with translations and dilations.

2.2 Stereographic Projections

We would like to consider maps from $\mathbb{S}^d \to \mathbb{S}^d$, while our Möbius transformations are from $\widehat{\mathbb{R}}^d \to \widehat{\mathbb{R}}^d$. In order to do so we must define a homeomorphism $\widehat{\mathbb{R}}^d \to \mathbb{S}^d$. We do this via a *stereographic projection*, π. We first embed $\widehat{\mathbb{R}}^d$ in $\widehat{\mathbb{R}}^{d+1}$ via $x \mapsto \tilde{x}$, the map such that:

$$x = (x_1, x_2, x_3, \ldots, x_d),\ \tilde{x} = (x_1, x_2, x_3, \ldots, x_d, 0),$$

and $\tilde{\infty} = \infty$. Thus we have a 1-1 map of $\widehat{\mathbb{R}}^d$ to the hyperplane $x_{d+1} = 0 \subset \widehat{\mathbb{R}}^{d+1}$. We can now map this hyperplane onto $\mathbb{S}^d$ by projecting $\tilde{x}$ towards e_{d+1} until it meets $\mathbb{S}^d$ in the unique point $\pi(\tilde{x})$.

It is now clear that $x \mapsto \pi(\tilde{x})$ is a 1-1 map of $\widehat{\mathbb{R}}^d$ onto $\mathbb{S}^d$. Explicitly:

Definition 2.5 (Stereographic projection) The stereographic projection $\pi\colon \widehat{\mathbb{R}}^d \to \mathbb{S}^d$ is the map defined by:

$$\pi(\tilde{x}) = \left(\frac{2x_1}{|x|^2+1}, \ldots, \frac{2x_d}{|x|^2+1}, \frac{|x|^2-1}{|x|^2+1} \right), \tag{2.5}$$

and $\pi(\infty) = e_{d+1}$.

While not immediately apparent, π is in fact a Möbius transformation of $\widehat{\mathbb{R}}^d$. To demonstrate this, consider the inversion R_4 in $\mathbb{S}^d(e_{d+1}, \sqrt{2})$ of $\widehat{\mathbb{R}}^d$. Then

$$R_4(\tilde{x}) = e_{d+1} + \frac{2(x_1, x_2, \ldots, x_d, -1)}{1+|x|^2} = \left(\frac{2x_1}{|x|^2+1}, \ldots, \frac{2x_d}{|x|^2+1}, \frac{|x|^2-1}{|x|^2+1} \right)$$

which is exactly Equation 2.5. Thus a stereographic projection is in fact an inversion - and thus a Möbius transformation - of $\widehat{\mathbb{R}}^{d+1}$. In fact, conjugating an inversion or reflection by a stereographic projection results in an inversion on $\mathbb{S}^d$. This now explains the name “chordal sphere”; every chord of $\mathbb{S}^d$ is sent to a $(d-1)$-sphere or $(d-1)$-plane by π. One can then define a stereographic projection from any point on $\mathbb{S}^d$ by composing with a suitable rotation, which is again Möbius.

Now, given a stereographic projection, we can transfer the Euclidean metric on $\mathbb{S}^d$ onto $\widehat{\mathbb{R}}^d$ using the following.

Definition 2.6 (Chordal metric) The *chordal metric* is the metric

$$d(x, y) = |\pi(\tilde{x}) - \pi(\tilde{y})| \tag{2.6}$$

Using Equation 2.5 we can compute an explicit expression for the chordal metric. From this expression, we may deduce that the chordal metric induces the same topology when restricted to $\mathbb{R}^d$ as the Euclidean metric. Hence, reflections and inversions as in Definition 2.1 are in fact homeomorphisms $\widehat{\mathbb{R}}^d \to \widehat{\mathbb{R}}^d$ under the chordal metric, and so a Möbius transformation is also such a homeomorphism.

The fact that stereographic projection is induced by an inversion implies that all invariants under Möbius transformations apply to both $\widehat{\mathbb{R}}^d$ and $\mathbb{S}^d$. Hence we can now define maps $\mathbb{S}^d \to \mathbb{S}^d$ with the same properties as our Möbius transformations.

Definition 2.7 The group of Möbius transformations of the d-sphere, denoted $\text{Möb}(\mathbb{S}^d)$ is the group of maps $\pi \circ \text{Möb}(d) \circ \pi^{-1}$, recalling that stereographic projections are in fact Möbius.

2.3 Properties of Möbius transformations

Having established the equivalence of Möbius transformations on $\widehat{\mathbb{R}}^d$ and $\mathbb{S}^d$, we now proceed to investigate their properties.

Definition 2.8 (Cross-ratio) Given four distinct points $x, y, u, v \in \widehat{\mathbb{R}}^d$, the *cross-ratio* of these points is

$$[x, y, u, v] = \frac{|x - u||y - v|}{|x - y||u - v|}, \tag{2.7}$$

where $|\cdot|$ is the Euclidean distance between points. If one of these points is ∞:

$$\begin{cases} [\infty, y, u, v] = \dfrac{|y - v|}{|u - v|}, \; [x, \infty, u, v] & = \dfrac{|x - u|}{|u - v|}, \\ [x, y, \infty, v] = \dfrac{|y - v|}{|x - y|}, \; [x, y, u, \infty] & = \dfrac{|x - u|}{|x - y|}. \end{cases} \tag{2.8}$$

The cross-ratio is the fundamental invariant under Möbius transformation.

Lemma 2.2. *A map $f : \widehat{\mathbb{R}}^d \to \widehat{\mathbb{R}}^d$ is a Möbius transformation if and only if it preserves cross-ratios.*

Proof. For the forward direction, observe that reflections and inversions as defined earlier preserve cross-ratio. Hence all Möbius transformations do so as the composition of reflections and inversions.

For the backwards implication, note that if we have a map ϕ that preserves cross-ratio but does not fix ∞, then we may compose it with an inversion such that the resulting map does fix ∞. Let ϕ be a map that preserves cross-ratio and has ∞ as a fixed point. Let $x_1, x_2, x, y \in \mathbb{R}^d$ be distinct points. Then

$$\frac{|x-x_2|}{|x_2-x_1|} = \frac{|x-x_2||x_1-\infty|}{|x_2-x_1||x-\infty|} = \frac{|\phi(x)-\phi(x_2)||\phi(x_1)-\infty|}{|\phi(x_2)-\phi(x_1)||\phi(x)-\infty|} = \frac{|\phi(x)-\phi(x_2)|}{|\phi(x_2)-\phi(x_1)|}$$

where we use the fact that ϕ preserves cross-ratio and fixes ∞ and Equalities 2.8. By similar logic, we obtain that

$$\frac{|x-y|}{|x-x_2|} = \frac{|\phi(x)-\phi(y)|}{|\phi(x)-\phi(x_2)|}.$$

Now, by multiplying these equal ratios,

$$\frac{|x-y|}{|x_2-x_1|} = \frac{|\phi(x)-\phi(y)|}{|\phi(x_2)-\phi(x_1)|} \implies |\phi(x)-\phi(y)| = \frac{|\phi(x_2)-\phi(x_1)|}{|x_2-x_1|}|x-y|.$$

If we consider x_1, x_2 to be fixed, then this implies that ϕ is a similarity with scale factor $k = |\phi(x_2)-\phi(x_1)|(|x_2-x_1|)^{-1}$ as in §2.1.1. □

The proof of this result also provides the following corollary.

Corollary 2.1. *If $\phi \in \text{Möb}(d)$ and $\phi(\infty) = \infty$ then ϕ is a similarity.*

Proof. The result follows immediately from the proof of Lemma 2.2. □

In fact, this provides the following result for Möbius transformations of $\mathbb{S}^d$.

Corollary 2.2. *If $\phi \in \text{Möb}(\mathbb{S}^d)$ has a fixed point, then ϕ is a similarity.*

Proof. If $\phi(x_0) = x_0$, then choose a stereographic projection sending x_0 to ∞. Then ϕ fixes ∞ as an element of $\text{Möb}(d)$ and so is a similarity by Corollary 2.1. □

Having investigated similarities and the cross ratio, let us briefly mention closed forms. We may express any Möbius transformation of $\widehat{\mathbb{R}}^d$ as

$$\phi(x) = b + \frac{\alpha A(x-a)}{|x-a|^\epsilon}, \tag{2.9}$$

where $a, b \in \mathbb{R}^d$, $\alpha \in \mathbb{R}$, $A \in O(d)$ and $\epsilon \in \{0, 2\}$. Note that $\epsilon = 0$ gives us the similarities, and $\epsilon = 2$ gives us those transformations involving inversion.

Also note that $\mathbb{S}^d$ is orientable, and so we may classify every Möbius transformation as either *orientation preserving* or *reversing.* A single inversion or reflection is orientation reversing. A composition of an odd number of reflections/inversions is also orientation reversing while an even number of reflections/inversions is orientation preserving. The orientation preserving transformations form a subgroup, $\text{Möb}^+(d) < \text{Möb}(d)$.

In the special case $d = 2$ any Möbius transformation can be expressed in its normalised form as

$$\phi(z) = \frac{Az + B}{Cz + D} \tag{2.10}$$

if ϕ is orientation preserving, or

$$\phi(z) = \frac{A\bar{z} + B}{C\bar{z} + D} \tag{2.11}$$

if ϕ is orientation reversing, where we use complex notation. Here $z \in \mathbb{R}^2 = \mathbb{C}$ and $A, B, C, D \in \mathbb{C}$ with $AD - BC = 1$. Möbius transformations of $\widehat{\mathbb{C}}$ are usually taken to be those that preserve orientation, $\text{Möb}^+(2)$, as those that reverse orientation are not analytic.

In fact, for the general case of $\widehat{\mathbb{R}}^d$, any Möbius transformation can be expressed as $(ax + b)(cx + d)^{-1}$ where $x \in \mathbb{R}$ and a, b, c, d are in a Clifford algebra and satisfy some specific constraints. This serves to demonstrate the form in Equation 2.10 is merely a special case. We will not go into more detail on this, but direct the reader to several sources on Clifford algebras and their applications to Möbius transformations [12][13].

We now return to demonstrating results for use in later chapters. It is well known that for $d = 2$, a Möbius transformation is determined by its action on three distinct points. We provide a proof of this before introducing a similar result for higher dimensions.

Lemma 2.3. *For two triples (z_1, z_2, z_3) and (z_1', z_2', z_3'), there is a unique Möbius transformation $\phi\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ such that $\phi(z_1, z_2, z_3) = (z_1', z_2', z_3')$.*

Proof. Suppose we have two Möbius transformations $\phi(z_1, z_2, z_3) = (z_1', z_2', z_3')$ and $\phi'(z_1, z_2, z_3) = (z_1', z_2', z_3')$. Take a third Möbius transformation $\psi(z_1', z_2', z_3') = (0, 1, \infty)$. To justify the existence of ψ, take

$$\psi_1(z) = \frac{1}{z - z_3}, \quad \psi_2(z) = z - \psi_1(z_1), \quad \psi_3(z) = \frac{z}{\psi_1(z_2) - \psi_1(z_1)}$$

then a quick check confirms that $\psi = \psi_3 \circ \psi_2 \circ \psi_1$ is such a map.

We may then define $\sigma = \psi \circ \phi$ and $\sigma' = \psi \circ \phi'$. Thus $\sigma^{-1} \circ \sigma'$ maps $(0, 1, \infty)$ to

$(0,1,\infty)$. We show that this implies $\sigma \circ \sigma' = \mathrm{id}$. Express $\sigma \circ \sigma' = (Az+B)/(Cz+D)$ as in Equation 2.10. Firstly, $\sigma \circ \sigma'(0) = 0$ implies $B = 0$. Now $\sigma \circ \sigma'(\infty) = \infty$ implies $C = 0$. Thus $\sigma \circ \sigma' = (A/D)z$. But $\sigma \circ \sigma'(1) = 1$ implies $A = D$. Because $AD - BC = 1$, $A = D \in \{1, -1\}$. In both cases $\sigma \circ \sigma'(z) = z$ and so we obtain the identity. But $\sigma^{-1} \circ \sigma$ is also the identity by similar logic, and thus $\sigma = \sigma' \implies \psi \circ \phi = \psi \circ \phi' \implies \phi = \phi'$. $\square$

The same is also clearly true of orientation reversing Möbius transformations of $\widehat{\mathbb{C}}$ by a very similar method of proof. Having demonstrated the 3-transitivity of Möbius transformations of $\widehat{\mathbb{C}}$, the next lemma gives us a generalisation of this result.

Lemma 2.4. *Given two subsets $S = \{x_i\}$, $S' = \{x'_i\}$, $i \in I$, lying in $\widehat{\mathbb{R}}^d$, there is a Möbius transformation ϕ with $\phi(S) = S'$ if and only if every cross ratio $[x_i, x_j, x_k, x_l] = [x'_i, x'_j, x'_k, x'_l]$ for $i, j, k, l \in I$. Given that ϕ exists, it is unique if and only if S contains $d+2$ points not all on the same chordal sphere.*

Proof. We follow the proof structure of Wilker in [10], fleshing out details he omits. We begin by proving existence.

For the forward implication, we have already seen that Möbius transformations preserve cross ratio by Lemma 2.2 and thus cross-ratios match in S and S'.

For the backwards direction, we begin by inverting x_1 and x'_1 to ∞. Call these inversions ι and ι'. Then S becomes $T = \{y_i\}$ and S' goes to $T' = \{y'_i\}$. Clearly ι and ι' are Möbius and so

$$[x_i, x_j, x_k, x_l] = [y_i, y_j, y_k, y_l]$$

and

$$[x'_i, x'_j, x'_k, x'_l] = [y'_i, y'_j, y'_k, y'_l],$$

for $i, j, k, l \in I$. Now suppose there is ψ with $\psi(T) = T'$. Then $\psi(\infty) = \infty$, and so by the Corollary 2.1, ψ is a similarity. Observe that $\phi := \iota'^{-1} \circ \psi \circ \iota$ is a Möbius transformation with $\phi(S) = S'$. As a result, ϕ exists if and only if ψ exists. We show the existence of ψ. By assumption, we have

$$[x_i, x_j, x_k, x_l] = [x'_i, x'_j, x'_k, x'_l]$$

which implies

$$[y_i, y_j, y_k, y_l] = [y'_i, y'_j, y'_k, y'_l],$$

and so $|y'_i - y'_j| = k|y_i - y_j|$ holds for all $i, j \in I$ by rearranging the cross ratio. This demonstrates such a similarity ψ exists, and so such a ϕ exists. This concludes our proof of existence.

We now proceed to prove the condition on uniqueness. To do so, we show that the following are equivalent:

(i) ϕ is unique.

(ii) ψ is unique.

(iii) T contains a d-simplex (A collection of $d+1$ points that do not all lie on a $(d-1)$-plane).

(iv) S contains $d+2$ points not on the same chordal sphere.

(i) $\implies$ (iv): Suppose S did lie entirely on a chordal sphere. Then we can find multiple ϕ which preserve S but differ elsewhere, for example the rotations about the chordal sphere containing S. Thus assuming a unique ϕ exists means (iv) holds.

(iv) $\implies$ (iii): If S contains $d+2$ points not all on the same chordal sphere, then under ι, all $d+2$ points again do not lie on a single chordal sphere. Take $d+1$ of these points in T, which define a unique chordal sphere, which may be a $(d-1)$-plane. The remaining point then does not lie on this chordal sphere, and so we may easily form a d-simplex.

(iii) $\implies$ (ii): If T contains a d-simplex, then so does T'. Call the simplices σ and σ', and note that $\psi(\sigma) = \sigma'$ must hold. Suppose there exists a similarity $\psi_1 \neq \psi$ with the same property, which implies $\psi_1^{-1}(\psi(\sigma)) = \sigma$. But then $\psi_1^{-1} \circ \psi$ is a similarity fixing a d-simplex, and so is an isometry fixing a d-simplex, and is thus the identity. This demonstrates that ψ must therefore be unique.

(ii) $\implies$ (i): If $\psi = \iota' \circ \phi \circ \iota^{-1}$ is unique, then $\phi = \iota'^{-1} \circ \psi \circ \iota$ is unique.

Thus we have shown (i) holds if and only if (iv) holds, and so the required uniqueness condition. $\square$

Note that the conditions of Lemma 2.4 are vacuously true if S and S' are triples, and so Möb(d) is always at least sharply 3-transitive. Furthermore, the lemma tells us that a Möbius transformation is determined by its action on at most $d+2$ points in general. This means we only ever have to check a finite number of points to entirely determine the map.

2.4 Characterising Möbius transformations by chordal sphere preservation

We now proceed to show a crucial characterisation of Möbius transformations via their preservation of chordal spheres.

Lemma 2.5. *If $d \geq 2$ and $f : \widehat{\mathbb{R}}^d \to \widehat{\mathbb{R}}^d$ is a bijective function with the property that $f(\Sigma)$ is a chordal sphere for every chordal sphere Σ, then f is a Möbius transformation of $\widehat{\mathbb{R}}^d$.*

Proof. Let us first recall Lemma 2.4, which demonstrated that the group of Möbius transformations is certainly sharply 3-transitive on $\widehat{\mathbb{R}}^d$. Given this, we may compose f with a Möbius transformation, and thereby assume that f fixes 0, ∞ and a point x on $\mathbb{S}^{d-1}$. As f is assumed to be chordal sphere preserving, we see in particular that f is $(d-1)$-plane preserving. We may use this fact and an induction to show that f preserves lines. Then by the fundamental theorem of affine geometry, f is an affine transformation, and is in fact linear as it fixes 0.

Consider a $(d-2)$-sphere centred on 0 and passing through x. Then as f is linear, it must preserve all such spheres. The union of all such spheres is exactly $\mathbb{S}^{d-1}$, and so f fixes the unit sphere. But the only linear transformations that fix the unit sphere are the orthogonal transformations, which we earlier demonstrated are themselves Möbius. Thus the proof is complete. □

Let us now make a short remark concerning the history of this lemma. It has been an objective to weaken the requirements on f in Lemma 2.5 for many years. Problems of this sort belong to the geometrical discipline referred to as "characterisations of geometrical mappings under mild hypotheses" [14]. There are many papers on problems in the discipline within the literature.

It may already be surprising that we need not assume f is continuous, yet it turns out that neither surjectivity nor injectivity are required. In fact Li and Yao showed in [15] that the following assumptions are sufficient:

(i) f preserves r-dimensional chordal spheres for some $1 \leq r \leq d$.

(ii) f is non-degenerate, that is to say, $f(\widehat{\mathbb{R}}^d)$ is not an r-dimensional chordal sphere.

then given (i), f is a Möbius transformation of $\widehat{\mathbb{R}}^d$ if and only if it satisfies (ii).

All that remains to complete the characterisation is to show that Möbius transformations preserve chordal spheres.

Lemma 2.6. *If ϕ is a Möbius transformation, then it preserves chordal spheres.*

Proof. It is apparent that similarities preserve chordal spheres; they preserve distance up to scaling. As $\text{Möb}(d)$ is generated by the similarities and R_0, we only have to check that R_0 preserves chordal spheres. Take a chordal sphere Σ with equation $a|x|^2+2\langle b, x\rangle+c = 0$. Then for $x \neq 0, \infty$, $y = R_0(x) = x/|x|^2$ satisfies

$$c|y|^2 + 2\langle b, y\rangle + a = \frac{c}{|x|^2} + 2\left\langle b, \frac{x}{|x|^2}\right\rangle + a = \frac{c - 2\langle b, x\rangle + a|x|^2}{|x|^2}.$$

Hence x lies on Σ if and only if y lies on $\Sigma'\colon c|y|^2 + 2\langle b, y\rangle + a = 0$. We then see that 0 lies on Σ if and only if $R_0(\infty)$ lies on Σ', and vice versa. Thus $R_0(\Sigma) = \Sigma'$. □

Combining Lemmas 2.5 and 2.6, we see that chordal sphere preservation characterises the Möbius transformations. We see then that maps in $\text{Möb}(\mathbb{S}^d)$ are characterised by their preservation of chords of $\mathbb{S}^d$.

Finally, we provide a quick result on how the Möbius group of the sphere acts on chordal spheres.

Lemma 2.7. *The action of* $\text{Möb}(d)$ *on the set of chordal spheres of* $\widehat{\mathbb{R}}^d$ *is sharply transitive. That is, for Σ and Σ' chordal spheres, there is a unique Möbius transformation ϕ such that $\phi(\Sigma) = \Sigma'$.*

Proof. Let S be a set of $d+1$ distinct points, and S' another set of $d+1$ distinct points in $\widehat{\mathbb{R}}^d$. Then S defines a unique chordal sphere Σ and similarly S' defines a unique chordal sphere Σ'. Then as S does not contain $d+2$ points on a single chordal sphere, there is a unique Möbius transformation ϕ such that $\phi(S) = S'$ by Lemma 2.4. This implies $\phi(\Sigma) = \Sigma'$ as claimed. □

As stated earlier, results proved in the previous two sections also apply to $\text{Möb}(\mathbb{S}^d)$ and chords of $\mathbb{S}^d$. We will make use of this in the next chapter.

Chapter 3

Extending results of A. Georgakopoulos

In [1], A. Georgakopoulos has proved certain results on sphere graphs and the d-sphere. In particular, he has proven that there is an isomorphism between $\mathrm{Aut}(\mathbb{S}^d)$ and $\mathrm{Aut}(\mathcal{C}^d)$ for $d = 1$, where $\mathrm{Aut}(\mathbb{S}^d)$ is the group of homeomorphisms from $\mathbb{S}^d$ to $\mathbb{S}^d$. We aim to prove the following generalisation for $d \geq 2$ in order to answer Question 1:

Theorem (Theorem 3.2) The map π from $\text{Möb}(\mathbb{S}^d)$ to $\mathrm{Aut}(\mathcal{C}^d)$ is an isomorphism for $d \geq 2$.

Where $\pi(g)$ is the graph automorphism canonically induced by a homeomorphism g of $\mathbb{S}^d$, that is, for a chord C, $\pi(g)(C) = g(C)$. Also note that each Möbius transformation h induces a graph automorphism $\pi(h)$ canonically.

3.1 Graph Notation

A large portion of the remaining chapters will concern graph theory, and so we introduce the relevant terminology.

A graph X consists of $V(X)$, its vertex set and $E(X)$ the edge set. An edge consists of an unordered pair xy where $x, y \in V(X)$. We then say x is adjacent to y, and write $x \sim y$. We consider only simple graphs; those graphs where no edge has identical end-vertices and no two vertices may be connected by multiple edges. We consider two graphs X, Y *isomorphic* if there is a bijection ϕ from $V(X)$ to $V(Y)$ such that $x \sim y$ if and only if $\phi(x) \sim \phi(y)$. A graph isomorphism $X \to X$ is a *graph automorphism*, a permutation of the vertices that maps edges to edges, and non-edges to non-edges. For

$S \subseteq V(G)$, the subgraph $G[S]$ of G *induced* by S has vertex set S and edge set consisting of all edges in $E(G)$ with both end-vertices in S.

We assume knowledge of the most basic terminology in graph theory and will endeavour to explain any less well-known notions that we use in later sections. There are many texts the reader may refer to for further reading; *Algebraic Graph Theory* by Godsil and Royle [16] is perhaps the closest in spirit to the focus of this dissertation.

3.2 Preliminary definitions and results

In order to tackle this problem, we must introduce some key definitions and results. Definitions 3.1, 3.3, 3.4 and 3.5 are analogous to definitions in [1]. Recall our definitions of $\mathbb{S}^d$ and $\mathbb{D}^d$ from Chapter 2 and that $\mathbb{S}^d = \partial\mathbb{D}^{d+1}$.

Definition 3.1 (Hyperchord) A hyperchord of $\mathbb{S}^d$ is the non-empty intersection of a d-dimensional, non-tangential hyperplane with $\mathbb{D}^{d+1}$. We refer to hyperchords as chords for the remainder of this dissertation.

Definition 3.2 (Sheaf of Planes) A sheaf of d-dimensional hyperplanes is the collection of all hyperplanes that contain V where V is a $(d-1)$-dimensional subspace of $\mathbb{R}^{d+1}$.

Definition 3.3 (Incident Chords) We refer to two chords P, Q as *incident* if $P \cap Q$ is a singleton. We often say "incident at x" when referring to a collection of pairwise incident chords whose intersection is x.

Definition 3.4 (Sphere Graph) The sphere graph $\mathcal{C}^d$ is the graph whose vertices are the chords of $\mathbb{S}^d$. Two chords form an edge whenever they intersect in $\mathbb{D}^{d+1}$.

Definition 3.5 (Boundary Clique) We let a boundary clique be the set of all chords containing $x \in \mathbb{S}^d$. This clearly induces a clique K_x in $\mathcal{C}^d$. Furthermore, $K_x \neq K_y$ for $x \neq y$.

Definition 3.6 (Incident Clique) We let an incident clique be a maximal set of pairwise incident chords at $x \in \mathbb{S}^d$. This also induces a clique I_x in $\mathcal{C}^d$, as all such chords contain the point x and so are adjacent in the sphere graph. We may see that $I_x \subset K_x$ certainly holds.

It is important to stress that these cliques are subsets of the vertex set of $\mathcal{C}^d$, and so we can view them as induced subgraphs. It is also simple to view them as subsets of $\mathbb{D}^{d+1}$, and in fact we will often go back and forth between the sphere graph and the d-sphere to construct proofs.

Furthermore, a notable difference between $d = 1$ and $d \geq 2$ arises here. All boundary cliques are also incident cliques for $d = 1$ however, this is not the case for higher dimensions. In fact we will shortly see that a boundary clique is just a union of incident cliques.

We now introduce *spherical caps* to construct our proof; their use ensures we also cover the case $d = 1$, allowing us to provide an alternative method of proof for the original result of Georgakopoulos. It is possible to instead use the unique embedded circle that is the intersection of a given chord with $\mathbb{S}^d$, and to prove analogous results. Such an avenue is perhaps slightly more in line with the material of later chapters, but does not cover all dimensions.

Definition 3.7 (Spherical caps) A chord P partitions $\mathbb{S}^d$ into two distinct regions. A spherical cap C_P is the union of P with exactly one of the aforementioned partitions. Each cap clearly uniquely defines a corresponding chord. For $d = 1$ a cap is a circular segment. Note that a given chord defines two spherical caps, giving us choice.

Having introduced the foundations, we now begin proving lemmas that will allow us to extend Georgakopoulos' results on sphere graphs.

Lemma 3.1. *If P and Q are a pair of intersecting chords of $\mathbb{S}^d$, then they partition $\mathbb{S}^d$ into four disjoint sets. Namely the following:*

- $A = \{x \in \mathbb{S}^d | x \notin C_P, x \notin C_Q\}$
- $B = \{x \in \mathbb{S}^d | x \notin C_P, x \in C_Q\}$
- $C = \{x \in \mathbb{S}^d | x \in C_P, x \notin C_Q\}$
- $D = \{x \in \mathbb{S}^d | x \in C_P, x \in C_Q\}$

If P and Q are incident, then either $|D| = 1$ or one of A, B or C is empty.

Proof. As our planes are distinct, intersecting d-dimensional subspaces inside a $(d+1)$-dimensional space, they must partition the d-sphere into four regions. We now analyse the four possible cases when P and Q are incident, see Figure 3.1.

In Case 1, $|D| = 1$ as D contains only the incident point, and A, B, C are all non-empty. In Case 2, $B = \varnothing$. In Case 3, $C = \varnothing$. In Case 4, $A = \varnothing$.

□

The above proof demonstrates that a pair of incident chords partition the d-sphere into three.

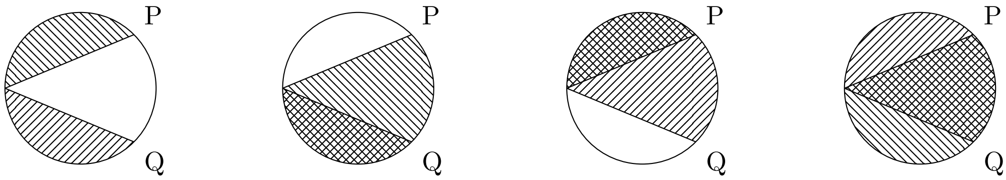

(a) Case 1 (b) Case 2 (c) Case 3 (d) Case 4

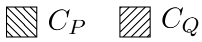

Figure 3.1: The four possible choices of caps for an incident pair.

Several of the following lemmas will rely on basic knowledge of vector spaces and smooth manifolds. There are many texts dealing with such material, but we would refer the reader to *An Introduction to Smooth Manifolds* by J. M. Lee for further reading on the topic of smooth manifolds. Given this, we now analyse the structure of incident cliques.

Lemma 3.2. *A collection of chords* $\{P_i\}_{i\in I}$ *with* $x \in P_i, \forall i \in I$ *is an incident clique if and only if it is the intersection of a sheaf containing* $V \subset T_x\mathbb{S}^d$*, excluding the tangent hyperplane, with* $\mathbb{D}^{d+1}$*, where* $\dim(V) = d - 1$*.*

Proof. For each chord P, recall that $P = H_P \cap \mathbb{D}^{d+1}$ where H_P is a hyperplane.

For the forward direction, suppose we have an initial chord $P \ni x$, with its associated hyperplane H_P. We may then define $T := T_x\mathbb{S}^d$, the tangent space at x, which is d-dimensional. We also define $\ell_P = H_P \cap T$. Now suppose we pick a chord Q incident to P at x. This necessarily implies $H_P \cap H_Q \cap \mathbb{S}^d = \{x\}$. Thus $H_P \cap H_Q \subset T$. We also define $\ell_Q = H_Q \cap T$. Thus

$$\ell_P \cap \ell_Q = (H_P \cap T) \cap (H_Q \cap T) = H_P \cap H_Q \cap T = H_P \cap H_Q. \tag{3.1}$$

Analysing dimensions, $\dim(H_P \cap H_Q) = d - 1$ as the intersection of two distinct d-dimensional spaces. Thus $\dim(\ell_P \cap \ell_Q) = d - 1$. But ℓ_p and ℓ_q are $(d-1)$-dimensional spaces themselves, and thus $\ell_P = \ell_Q$. Hence all chords incident to P at x must lie in the sheaf around $\ell_P \subset T_x\mathbb{S}^d$, and thus the incident clique containing P must be the sheaf around ℓ_P minus T, as it is tangent to $\mathbb{S}^d$.

To show the reverse implication, suppose we have a sheaf S around $V \subset T_x\mathbb{S}^d$ minus $T_x\mathbb{S}^d$. Then for $H_i, H_j \in S$,

$$\{x\} = H_i \cap H_j \cap \mathbb{S}^d = (H_i \cap \mathbb{S}^d) \cap (H_j \cap \mathbb{S}^d) = P_i \cap P_j \tag{3.2}$$

and so pairwise incidence is satisfied. Maximality is also satisfied, any plane not in S cannot be pairwise incident to all members of S. $\square$

We now show that incident cliques are the fundamental building blocks of boundary cliques, which form the basis of our proof structure.

Lemma 3.3. *Each boundary clique K_x is the disjoint union of the incident cliques at x.*

Proof. Firstly, given a chord P that contains x, it is clearly the intersection of a hyperplane, say H_P, contained in a sheaf around a tangent to x with $\mathbb{D}^d$. Thus P is in some incident clique I_x. This must be a unique incident clique, otherwise H_P would contain two $(d-1)$-dimensional tangents to $\mathbb{S}^d$. Thus every chord in K_x is contained in a single incident clique, and so the disjoint union of incident cliques at x must be K_x. $\square$

Intuitively, a chord containing a fixed point x can be determined by two parameters:

1. Its "angle" (i.e. orientation within a fixed sheaf)

2. Its "rotation" (i.e. the direction of the tangent at x)

An incident clique at x is the collection of chords with a fixed rotation containing x. In particular, this tells us that there are infinitely many disjoint incident cliques at a point x. A boundary clique is the collection of chords containing x with any angle and rotation, and so it is the union of incident cliques at x.

Having determined the structure of incident and boundary cliques, we introduce two useful invariants under automorphisms of intersection graphs of curves.

3.3 Region counting

We will often refer to a *region counting* argument throughout this dissertation. Given an intersection graph $\mathcal{Z}$ of some class of curves embedded in $\mathbb{S}^d$ we may take the induced subgraph of $\mathcal{Z}$ consisting of all vertices not adjacent to a given collection of curves. Removing such curves corresponds to removing the *closed neighbourhood* of every curve in the collection. Recall that the closed neighbourhood of a vertex, v, in a graph is the set of vertices adjacent to v as well as v itself. We denote such a neighbourhood $\overline{N}(v)$. In this chapter, $\mathcal{Z} = \mathcal{C}^d$ - as mentioned earlier every chord defines a unique embedded circle in $\mathbb{S}^d$; in later chapters we will consider more general classes of curves.

Given a collection of curves, say $\{C_i\} := \{C_i\}_{i\in I} \subset V(\mathcal{Z})$, we define:

$$\mathrm{Reg}(\{C_i\}) = \#\text{ of connected components of ISG on } V(\mathcal{Z})\backslash\left(\bigcup_{i\in I}\overline{N}(C_i)\right), \tag{3.3}$$

the region function, where "ISG" stands for "induced subgraph". The vertices not adjacent to any C_i are exactly those that do not intersect any C_i, and are not in the collection themselves. Then the region function gives us a direct correspondence between the graph and the geometry of the sphere. The connected components of this induced subgraph correspond to the topological components of $\mathbb{S}^d\backslash\{C_i\}$. It is clear that a graph automorphism $h \in \mathrm{Aut}(\mathcal{Z})$ must preserve the number of connected components, and so the region function satisfies:

- $\mathrm{Reg}(\{C_i\}) = \mathrm{Reg}(h(\{C_i\}))$
- $\mathrm{Reg}(H) = \mathrm{Reg}(h(H))$

where $H \subset \{C_i\}$. This fact will prove crucial in understanding the behaviour of graph automorphisms on various classes of curves.

In general, if we claim that a collection of curves gives us n regions, we are implicitly using the region function.

3.4 Region graphs

In addition to our region function, we may define a *region graph*, $R_{\{C_i\}}$ on a set of curves in $\mathbb{S}^d$. Each vertex corresponds to a unique connected component in the aforementioned induced subgraph on

$$V(\mathcal{Z})\backslash\left(\bigcup_{i\in I}\overline{N}(C_i)\right) \tag{3.4}$$

For two connected components X, Y, we have an edge $X \sim Y$ if and only if there is a $j \in I$ with X, Y induced subgraphs of Z, where Z is a connected component in the induced subgraph on

$$V(\mathcal{Z})\backslash\left(\bigcup_{i\neq j\in I}\overline{N}(C_i)\right) \tag{3.5}$$

That is to say, there is a curve, which when removed from our picture, results in the two regions becoming one. In subsequent chapters we consider more general curves on $\mathbb{S}^2$, but in these cases region graphs may be defined for higher dimensions equally well.

We now prove that region graphs are indeed an invariant under any $h \in \mathrm{Aut}(\mathcal{Z})$.

Lemma 3.4. *A region graph $R_{\{C_i\}}$ is preserved under $h \in \mathrm{Aut}(\mathcal{Z})$. That is, $R_{\{C_i\}} \cong R_{h(\{C_i\})}$, where $h(\{C_i\}) = \{h(C_i) | i \in I\}$.*

Proof. Assume we are given $\{C_i\}_{i \in I}$ and the corresponding region graph. We show $X \sim Y$ if and only if $h(X) \sim h(Y)$ where X and Y are vertices of the region graph.

For the forward implication, suppose $X \sim Y$. Then X, Y are induced subgraphs of Z a connected component in the induced subgraph on Expression 3.5. Then $h(Z)$ is a connected component of the induced subgraph on

$$V(h(\mathcal{Z}))\backslash \left(\bigcup_{i \neq j \in I} \overline{N}(h(C_i)) \right) \tag{3.6}$$

for some $j \in I$. If not, then we contradict that h preserves adjacency. In particular, the number of connected components must be preserved by h. Furthermore, $h(X)$ and $h(Y)$ must be induced subgraphs of $h(Z)$. Now consider the induced subgraph on

$$V(h(\mathcal{Z}))\backslash \left(\bigcup_{i \in I} \overline{N}(h(C_i)) \right) \tag{3.7}$$

Clearly, $h(X)$ and $h(Y)$ must be distinct connected components of this induced subgraph, otherwise we again contradict that h preserves adjacency. This implies that removing $h(C_j)$ from the collection unites $h(X)$ and $h(Y)$ into a single connected component. This in turn implies that $h(X) \sim h(Y)$.

It only remains to show that $h(X)$ and $h(Y)$ are indeed vertices of $R_{\{h(C_i)\}}$. We must show the diagram in Figure 3.2 commutes.

$$\begin{tikzcd} \{C_i\} \arrow[r, "h"] \arrow[d, "\text{Region}"'] & h(\{C_i\})) \arrow[d, "\text{Region}"] \\ \text{components} \arrow[r, "h"'] & h(\text{components}) \end{tikzcd}$$

Figure 3.2: Commutative diagram

We let Region: $\mathcal{P}(V(\mathcal{Z})) \to \mathbb{N}$ be a function that takes a collection of curves $\{C_i\}$, a subset of the vertex set of some intersection graph of curves, $\mathcal{Z}$, and gives us the connected components of the resulting induced subgraph on Expression 3.4. Note that $\operatorname{Reg}(C) = |\operatorname{Region}(C)|$. Clearly, $\overline{N}(C_i) \cong \overline{N}(h(C_i))$ for all $i \in I$. Thus $\bigcup_{i\in I} \overline{N}(C_i) \cong \bigcup_{i\in I} \overline{N}(h(C_i))$ and hence Expression 3.4 is isomorphic to Expression 3.7 as a graph. Hence the connected components of each are isomorphic, and so $h(X)$ and $h(Y)$ are vertices of $R_{\{h(C_i)\}}$ as expected.

The backwards direction follows by exactly the same arguments and thus the region graph is preserved by any $h \in \operatorname{Aut}(\mathcal{Z})$. □

We will see that region count and region graphs are useful tools in answering Question 1.

3.5 Constructing a well-defined map of the d-sphere via boundary cliques

The aim of this section is to show that boundary cliques are mapped bijectively under the action of a given graph automorphism of $\mathcal{C}^d$, and so provide a well-defined map of the d-sphere. We begin by recalling Lemma 3.1 from [1], which forms the basis of this section.

Lemma 3.5 (Lemma 3.1)**.** *Every $h \in \operatorname{Aut}(\mathcal{C}^d)$ maps each boundary clique onto a boundary clique for $d = 1$.*

We aim to extend this result to $d \geq 2$. To do this, we first prove results on incident cliques and then combine these to obtain a result for boundary cliques.

Lemma 3.6. *Every $h \in \operatorname{Aut}(\mathcal{C}^d)$ maps an incident pair onto an incident pair.*

Proof. Let P, Q be an incident pair of chords which are sent to a non-incident pair of chords, $h(P)$, $h(Q)$ by our automorphism h. There are two scenarios to consider in the image. Either $h(P)$ and $h(Q)$ do not intersect or they intersect but are not incident. The first is impossible as h is an automorphism of $\mathcal{C}^d$ and so preserves adjacency. Thus we may assume that $h(P)$ and $h(Q)$ intersect non-incidentally. Ergo, they divide the sphere into four disjoint sets by Lemma 3.1.

We now use a region counting argument, explicitly stating all steps to make the process clear. Consider the collection of all chords $R \in V(\mathcal{C}^d)$ which are not adjacent to P or Q. This is the collection of all chords in of $\mathbb{S}^d$ which are disjoint to both P and Q. Let us call the collection of all such chords $\mathcal{R}$. Considering $\mathcal{R}$ as an induced subgraph of $\mathcal{C}^d$, we show that it is composed of three connected components. Firstly, only three topological components as demonstrated in Figure 3.1(a) may contain chords, and so $\mathcal{R}$ must have exactly three connected components. To show connectedness of each component, let us consider non-adjacent chords M, N in a particular partition. We can always choose C_M and C_N so that they do not intersect P or Q. By the connectedness of $\mathbb{S}^d$, there is a shortest path connecting C_M to C_N which lies entirely within the selected topological component. We may then specify a sequence of caps not intersecting P or Q - giving us chords - but intersecting the previous cap in the sequence. Thus there is a path from M to N in our component of $\mathcal{C}^d$ and so each component of $\mathcal{R}$ is indeed connected. Hence under h each of these components remains connected, as edges are preserved. However, each component must be mapped by h to a single partition of $\mathbb{S}^d$ in the image, as we cannot split or add $h(P)$ or $h(Q)$ to a component. But we have only three components of $\mathcal{R}$, and four partitions in the image, so we have a partition with no component mapped to it. But this provides a contradiction. There is clearly a chord in this partition which intersects neither $h(P)$ nor $h(Q)$ and hence must be the image of some element of $\mathcal{R}$. But it is not adjacent to any of the images of components of $\mathcal{R}$, which has no isolated vertices. Thus h cannot map P and Q to a pair of non incident chords. □

In future proofs we will omit the explicit construction of a contradiction if an automorphism fails to preserve region count. We now demonstrate that incident cliques are preserved by automorphisms of $\mathcal{C}^d$.

Lemma 3.7. *Every $h \in \mathrm{Aut}(\mathcal{C}^d)$ maps an incident clique onto an incident clique.*

Proof. For an n-tuple of incident chords, $\mathbb{S}^d$ is partitioned into $n+1$ sets, as a consequence of Lemma 3.2. If our tuple is not mapped to another n-tuple of incident chords, then there are two cases to consider. Either we have at least one pair of chords in the image that are disjoint, or we have a pair that intersect but are not incident. In the first case, we contradict that h preserves adjacency. In the second, the image chords will partition $\mathbb{S}^d$ into at least $n+2$ regions which contradicts that h preserves region count. Hence incident n-tuples are preserved.

Thus incident cliques are sent to incident cliques as all finite subsets of an incident clique are also preserved. Explicitly, $h(I_x) \subseteq I_y$ for some $y \in \mathbb{S}^d$ and I_x, I_y incident cliques at x and y. If $h(I_x)$ is a proper subset of I_y, then $h^{-1}(I_y)$ intersects I_x but does not contain it. This then contradicts that h^{-1} preserves incident chords. Thus $h(I_x) = I_y$ □

Now we may show that distinct incident cliques at a point are preserved injectively.

Lemma 3.8. *Let $I_x^1 \neq I_x^2$ be distinct incident cliques at $x \in \mathbb{S}^d$. Then $\forall h \in \mathrm{Aut}(\mathcal{C}^d)$, $h(I_x^1) = I_y^1$ and $h(I_x^2) = I_y^2$ with $I_y^1 \neq I_y^2$.*

Proof. Observe that $I_x^1 \cup I_x^2$ forms a clique in $\mathcal{C}^d$, as all of these chords contain x. Suppose now that $h(I_x^1) = I_y$ and $h(I_x^2) = I_z$ with $y \neq z$, by Lemma 3.7. Then we may find $P \in I_y$ and $Q \in I_z$ with $P \cap Q = \varnothing$. Specifically, as $y \neq z$, there is a diameter between the two, so that they are in distinct hemispheres. We can then pick P and Q so that they are contained solely in one hemisphere. As $h^{-1}(P),\ h^{-1}(Q) \in I_x^1 \cup I_x^2$ this contradicts that h preserves edges, and thus $y = z$. We call these I_y^1 and I_y^2.
Now suppose $I_y^1 = I_y^2$. Take a chord $P \in I_x^1 / I_x^2$. Then $h(P) \in I_y^1 = I_y^2$. Thus $h^{-1}(h(P)) \in I_x^2$, a contradiction. Thus $I_y^1 \neq I_y^2$, and distinct incident cliques at a point are mapped to distinct incident cliques at a point. □

We can also show that distinct incident cliques at different points are mapped injectively by graph automorphisms.

Lemma 3.9. *Let I_x and I_y be incident cliques at distinct points. Then $h(I_x)$ and $h(I_y)$ are also incident cliques at distinct points, for $h \in \mathrm{Aut}(\mathcal{C}^d)$.*

Proof. We may find a chord $P \in I_x$ and a chord $Q \in I_y$ such that $P \cap Q = \varnothing$. If $h(I_x) = h(I_y)$ then $h(P) \cap h(Q) \neq \varnothing$, contradicting that h preserves adjacency. Thus incident cliques at different points are also mapped injectively by graph automorphisms. □

We finally show that graph automorphisms map incident cliques surjectively.

Lemma 3.10. *For an incident clique I_y, there exists I_x such that $h(I_x) = I_y$. That is, $h \in \mathrm{Aut}(\mathcal{C}^d)$ is surjective on incident cliques.*

Proof. We know from Lemma 3.7 that $h^{-1}(I_y)$ is an incident clique I_x satisfying the statement of the lemma. Clearly this incident clique must be unique, otherwise we contradict Lemma 3.9 when applied to h^{-1}. □

Thus we see that any automorphism of $\mathcal{C}^d$ is bijective in its action on incident cliques at a point, by combining Lemmas 3.8 and 3.10. Furthermore, automorphisms of $\mathcal{C}^d$ act bijectively on incident cliques at distinct points, as a consequence of Lemmas 3.9 and 3.10. Given that incident cliques form the building blocks of boundary cliques, we suspect that a similar result should hold for boundary cliques. We now prove so.

Lemma 3.11. *For a boundary clique K_x, $h(K_x) = K_y$ for all $h \in \mathrm{Aut}(\mathcal{C}^d)$ and for some unique $y \in \mathbb{S}^d$. That is, boundary cliques are sent to boundary cliques by graph automorphisms.*

Proof. By Lemma 3.3, a boundary clique K_x is simply the union of all incident cliques at x. So for $I_x \subset K_x$ we know that $h(I_x) = I_y \subset K_y$ for a unique $y \in \mathbb{S}^d$ by Lemma 3.9. Thus

$$h(K_x) = \bigcup_{I_x \subset K_x} h(I_x) = \bigcup_{I_y \subset K_y} I_y = K_y \tag{3.8}$$

by the bijectivity of h, and where we take the union to be the induced subgraph on the union of the vertex sets. By the same logic $h^{-1}(K_y) = K_x$. Thus a boundary clique is mapped to a unique boundary clique. □

Lemma 3.12. *Boundary cliques are mapped bijectively under automorphisms of $\mathcal{C}^d$.*

Proof. For injectivity, take two distinct boundary cliques K_x, K_y. Suppose $h(K_x) = h(K_y) = K_z$, but then we may find a small chord P in K_x and a small chord Q in K_y such that $P \cap Q = \varnothing$. But then $h(P) \cap h(Q) \neq \varnothing$ as both contain z. Hence we have a contradiction, and so h maps boundary cliques injectively. For surjectivity, take a boundary clique K_y and a graph automorphism h. Then by Lemma 3.11, $h^{-1}(K_x)$ is a boundary clique. □

Note that we do not have to consider distinct boundary cliques at a point x, as there is a unique such clique for a given point.

At this stage, one might question whether incident and boundary cliques are necessary to construct a bijective map of $\mathbb{S}^d$. We address this concern in the following section.

3.6 Motivating incident cliques

The motivation behind defining incident cliques and boundary cliques is that they tell us exactly how a point in $\mathbb{S}^d$ is mapped under our graph automorphism, and thus allow us to define a homeomorphism of the sphere in §3.7. In fact, we do not strictly need infinite collections; representing a point x by an incident pair would suffice. As in Theorem 3.2, we pick a pair incident at x, and map x to y, where their image under a graph automorphism is incident at y. One concern around finite arrangements is demonstrated in the following example, in $\mathbb{S}^3$, where chords are circles.

Example 1 We consider a collection of four pairwise incident chords, P_1, P_2, P_3, P_4, of $\mathbb{S}^3$ in the the arrangement demonstrated in Figure 3.3. It is clear that pairwise incidence is satisfied, but there is not a single point of incidence. In general such a configuration is possible for up to $d+2$ pairwise incident chords of $\mathbb{S}^d$; a pyramid with an extra chord in the middle. In these cases, a region counting argument does not distinguish this arrangement from an arrangement that defines a single point of incidence, except for $d = 2$. This is because such an arrangement divides the sphere up into $d+1$ regions, the same as for a sheaf arrangement of pairwise incident chords. However, an argument via region graphs does demonstrate a sheaf arrangement cannot be split into a pyramid by any $h \in \mathrm{Aut}(\mathcal{C}^d)$. Namely, the region graph of a sheaf arrangement on n chords is P_{n+1}, while a pyramidal arrangement has at least one vertex of degree at least three in any dimension.

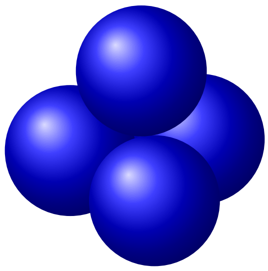

Figure 3.3: A pyramidal arrangement of spheres

However, it is exactly our results on incident and boundary cliques that demonstrate a map constructed using representative incident pairs is well-defined:that any choice of representative pair for a given point will be mapped to the same point under a given graph automorphism. While we will not provide a proof, this follows from our lemmas on incident cliques, and is not obvious a priori. Furthermore, the use of incident and boundary cliques demonstrates the surprising geometric properties that can be deduced entirely combinatorially about graph automorphisms of $\mathcal{C}^d$. Finally, boundary cliques do not require a choice of representative for each point which aids in succinctness when proving our main result.

We now have two ways to construct a map of $\mathbb{S}^d$. We can use incident cliques, choosing a representative for each $x \in \mathbb{S}^d$, or we can use boundary cliques, in which case no choice of representative need be made. In the following section we employ boundary cliques, but replacing each K_x with a unique I_x suffices to show the same result.

3.7 Constructing homeomorphisms of the d-sphere via boundary cliques

We start by recalling Theorem 1.1 from [1], which states Georgakopoulos' result mentioned at the start of this chapter.

Theorem 3.1 (Theorem 1.1)**.** *The map π is an isomorphism from* $\mathrm{Aut}(\mathbb{S}^d)$ *onto* $\mathrm{Aut}(\mathcal{C}^d)$ *for $d = 1$.*

Where $\pi(g)$ is the graph automorphism canonically induced by a homeomorphism g of $\mathbb{S}^1$, that is, for a chord xy, $\pi(g)(\{x, y\}) = \{g(x), g(y)\}$. Also note that each Möbius transformation h induces a graph automorphism $\pi(h)$ canonically.

We prove the following statement for higher dimensions:

Theorem 3.2. *The map π from* $\mathrm{Möb}(\mathbb{S}^d)$ *to* $\mathrm{Aut}(\mathcal{C}^d)$ *is an isomorphism for $d \geq 2$.*

where π is defined accordingly.

Proof. We follow a proof structure similar to that of Theorem 1.1 in [1]. Specifically, we use the action of h on boundary cliques to construct a corresponding homeomorphism.

Firstly, π is clearly an injective homomorphism, as $\ker(\pi) = \mathrm{id}$. For surjectivity, we want to find $h' \in \mathrm{Möb}(\mathbb{S}^d)$ with $\pi(h') = h$ for a given $h \in \mathrm{Aut}(\mathcal{C}^d)$. By Lemma 3.12, for every $x \in \mathbb{S}^d$, $h(K_x) = K_y$ for a unique $y \in \mathbb{S}^d$. Define h' by $x \mapsto y$. We must show this is indeed a Möbius transformation. To show h' is injective, suppose $x \neq y \in \mathbb{S}^d$ and $h'(x) = h'(y) = z$. Then $h(K_x) = h(K_y) = K_z$, contradicting Lemma 3.12. For surjectivity, pick $y \in \mathbb{S}^d$ and note that $h^{-1}(K_y) = K_x$ for some $x \in \mathbb{S}^d$ by Lemma 3.11, and so $h'(x) = y$. Furthermore, for $P, Q \in V(\mathcal{C}^d)$, we have that:

$$h(P) = Q \text{ if and only if } h'(P \cap \mathbb{S}^d) = Q \cap \mathbb{S}^d \tag{3.9}$$

as the unique boundary cliques containing P and Q lie along the $(d-1)$-spheres that are their intersections with $\mathbb{S}^d$. (The boundary cliques containing P must be sent to the boundary cliques containing Q for $h(P) = Q$ to be satisfied.) From this we see that $\pi(h') = h$, and so π is bijective. Furthermore, $\pi(g' \circ h')(P) = g'(h'(P)) = g(h(P)) = (g \circ h)(P)$ for g' and h' constructed as above from two graph automorphisms g, h. This demonstrates that π respects composition and so is indeed an isomorphism.

Note that Condition 3.9 implies that h' must preserve $(d-1)$-spheres, and is bijective. Thus we satisfy the characterisation of Lemma 2.5 and so h' is a Möbius transformation of $\mathbb{S}^d$. $\square$

Furthermore, by using hyperspherical caps, our method of proof extends to the case $d = 1$. In this case we find that homeomorphisms constructed from graph automorphisms preserve nestedness of pairs of points, which suffices to demonstrate that all elements of $\mathrm{Aut}(\mathbb{S}^1)$ correspond to a unique graph automorphism. In this way we have also provided an alternative proof of Georgakopoulos' result.

We now demonstrate explicitly that $\mathrm{Möb}(\mathbb{S}^d)$ is not isomorphic to $\mathrm{Aut}(\mathbb{S}^d)$. Firstly, for $d \geq 2$, there are homeomorphisms of $\mathbb{S}^d$ which do not preserve chords. Consider a local dilation, for example (a dilation of the sphere composed with a bump function [17]). A $(d-1)$-sphere which intersects but is not contained in the scaled region will no longer be a $(d-1)$-sphere after the homeomorphism is applied. This certainly demonstrates that the natural candidate π cannot be an isomorphism between $\mathrm{Aut}(\mathcal{C}^d)$ and $\mathrm{Aut}(\mathbb{S}^d)$, but there may still exist some other isomorphism.

Let us now show such an isomorphism cannot exist, via Thompson's T group. Firstly, we note that $T \leq \mathrm{Aut}(\mathbb{S}^1)$ - indeed, $T \leq \mathrm{Aut}^+(\mathbb{S}^1)$, the group of orientation preserving homeomorphisms of $\mathbb{S}^1$ [18]. Hence, we may embed T in $\mathrm{Aut}(\mathbb{S}^d)$ as follows. Let us create the *suspension* of $\mathbb{S}^d$, $S\mathbb{S}^d := (\mathbb{S}^d \times [-1,1])/(\mathbb{S}^d \times \{-1\})/(\mathbb{S}^d \times \{1\})$, which is equal to $\mathbb{S}^{d+1}$ [19]. Namely, we quotient by the end faces of $\mathbb{S}^d \times [-1,1]$, reducing each to a point. Now given $h \in \mathrm{Aut}^+(\mathbb{S}^d)$, define

$$\widehat{h}[x,r] = [h(x), r].$$

as a map $\mathbb{S}^{d+1} \to \mathbb{S}^{d+1}$. Then $\widehat{h \circ g}[x,r] = [(h \circ g)(x), r] = [h(g(x)), r] = \widehat{h} \circ \widehat{g}[x,r]$, and so $\widehat{\cdot}$ respects the group operation. Furthermore, if $\widehat{h} = \mathrm{id}$, then for $r \in (-1,1)$, $[h(x), r] = [x, r]$ which implies h is the identity, and so the map is injective. One may also check that $\widehat{h}$ is continuous under the compact open topology. Hence, we have a continuous embedding of $\mathrm{Aut}^+(\mathbb{S}^d)$ in $\mathrm{Aut}^+(\mathbb{S}^{d+1})$. By iterating, we may conclude $T \leq \mathrm{Aut}^+(\mathbb{S}^d)$ for all d. We now mention several well-known facts about T. In particular, T is infinite, finitely generated, and simple - it has no proper normal subgroups. There are many treatments of the group in the literature, see [20] for example. We say a group G is *residually finite* if for any $\mathrm{id}_G \neq g \in G$, there is a homomorphism $h : G \to H$ where H is a finite group and $h(g) \neq \mathrm{id}_H$ [21]. Then T cannot be residually finite; any homomorphism from a simple group must be either trivial or injective. So suppose T is residually finite; this implies we can find injective homomorphisms from an infinite group into a finite group - an impossibility. Let us then make use of Mal'cev's theorem, which states that any finitely generated, linear group is residually finite [21]. Then T

cannot be linear over any field. But recall from Chapter 2 that Möb($\mathbb{S}^d$) is isomorphic to $O^+(1, d+1)$, which is a subgroup of $GL_{d+2}(\mathbb{R})$. Hence the finitely generated subgroups of Möb($\mathbb{S}^d$) are linear and residually finite. Hence, any homomorphism between T and such a subgroup must be trivial, otherwise T is a linear group and is therefore residually finite by being finitely generated; a contradiction. This therefore demonstrates that Möb($\mathbb{S}^d$) cannot be isomorphic to Aut($\mathbb{S}^d$) for any $d \geq 2$, even as abstract groups.

Hence, we have proved our stated answer to Question 1.

3.8 Extending the action of $\mathrm{Aut}(\mathcal{C}^d)$ via interior cliques

We now extend the action of $\mathrm{Aut}(\mathcal{C}^d)$ to the interior of $\mathbb{S}^d$ via *interior cliques.*

Definition 3.8 (Interior clique) Let $x \in \mathbb{D}^{d+1} \backslash \mathbb{S}^d =: \mathrm{int}(\mathbb{S}^d)$. We define the *interior clique* J_x at x as:

$$J_x = \{P \in V(\mathcal{C}^d) \colon x \in P\}$$

So an interior clique is the maximal set of pairwise-intersecting chords whose intersection contains x.

Note that there is a unique interior clique for each point x.

We begin by showing that interior cliques are mapped bijectively by automorphisms of $\mathcal{C}^d$.

Lemma 3.13. *A pair of intersecting, non-incident chords cannot be mapped to a pair of incident chords by $h \in \mathrm{Aut}(\mathcal{C}^d)$.*

Proof. Let P and Q be chords in $V(\mathcal{C}^d)$ that intersect but are not incident. Then $\mathrm{Reg}(\{P, Q\}) = 4$. Clearly h cannot map P and Q to a disjoint pair, as h preserves adjacency. If h maps the pair P, Q to a pair of incident chords, then we get only three regions in the image, contradicting that h preserves region count. Hence the image pair must intersect non-incidentally. $\square$

Lemma 3.14. *Every interior clique is mapped to an interior clique by $h \in \mathrm{Aut}(\mathcal{C}^d)$.*

Proof. As a consequence of Lemma 3.13, we see that n-tuples of pairwise intersecting chords are preserved. An interior clique J_x is just such a tuple, and so J_x must be sent to a set of pairwise intersecting chords by h. In fact, the intersection of these chords must also be a single point. If this is not the case, then we contradict that h preserves adjacency. Thus $h(J_x) \subseteq J_y$ for some $y \in \mathbb{D}^{d+1}$. Firstly we show that such a y must be in $\text{int}(\mathbb{S}^d)$. If $y \notin \text{int}(\mathbb{S}^d)$, then we must have that $y \in \mathbb{S}^d$. But this makes J_y a boundary clique, contradicting Lemma 3.11 when applied to h^{-1} and J_y. Now if $h(J_x)$ is a proper subset of J_y, then $h^{-1}(J_y)$ is not contained in J_x but intersects it, which contradicts that h^{-1} preserves pairwise intersecting chords. Thus $h(J_x) = J_y$. $\square$

Lemma 3.15. *Interior cliques are mapped bijectively by automorphisms of $\mathcal{C}^d$.*

Proof. Lemma 3.14 tells us that $h(J_x) = J_y$ for all interior cliques J_x. Now suppose there are distinct interior cliques J_x and J_y both mapped to J_z by h. But then we may find a chord in J_x and a chord in J_y that are disjoint. But under h these chords must intersect at z, contradicting that h preserves adjacency. Hence h maps interior cliques injectively.

For surjectivity, suppose we have an interior clique J_y. Take a graph automorphism h. Then Lemma 3.14 tells us that $h^{-1}(J_y)$ is an interior clique. $\square$

We now use interior cliques to naturally extend the action of an automorphism of $\mathcal{C}^d$ to the interior of the d-sphere.

Theorem 3.3. *Every Möbius transformation $\phi\colon \mathbb{S}^d \to \mathbb{S}^d$ as in §3.7 extends to a Möbius transformation $\phi'\colon \mathbb{D}^{d+1} \to \mathbb{D}^{d+1}$ via automorphisms of $\mathcal{C}^d$.*

Proof. We have seen that for every $x \in \text{int}(\mathbb{S}^d)$, $h(J_x) = J_y$ for a graph automorphism $h \in \text{Aut}(\mathcal{C}^d)$. We thus extend ϕ via ϕ' by $x \mapsto y$ if and only if $h(J_x) = J_y$. That is to say, ϕ' is defined by the action of h on boundary cliques and interior cliques. Then ϕ' is bijective as h acts bijectively on both boundary and interior cliques, by Lemmas 3.12 and 3.15. Finally, $h(P) = Q$ if and only if $\phi'(P) = Q$, where we recall that P and Q are defined as the intersection of a hyperplane with $\mathbb{D}^{d+1}$. So in particular, $h(P) = Q$ implies that the boundary cliques and interior cliques containing P are sent to those containing Q. This implies that ϕ' is a bijective map that preserves $\mathbb{D}^{d+1}$ and also preserves $(d-1)$-dimensional spheres. If we now define a map $\widehat{\phi}' = \phi'$ on $\mathbb{D}^{d+1}$ and $\widehat{\phi}'(\infty) = \infty$, we find that $\widehat{\phi}'$ is a bijective map of $\mathbb{S}^{d+1}$ that preserves d-spheres. Hence we see that $\widehat{\phi}'$ is a Möbius transformation of $\mathbb{S}^{d+1}$ by Lemma 2.5, and in particular that ϕ' is a Möbius transformation of $\mathbb{D}^{d+1}$ as a restriction of $\widehat{\phi}'$. $\square$

This provides a novel way to extend Möbius transformations to the interior of the d-sphere via a purely combinatoric construction.

3.9 Analysing the structure of $\mathcal{C}^d$

Having concluded that $\mathrm{Aut}(\mathcal{C}^d) \cong \mathrm{Möb}(\mathbb{S}^d)$, we can make some observations on the structure of the sphere graph.

Observation 3.1 Lemma 3.6 implies that $\mathcal{C}^d$ is not edge-transitive. No $h \in \mathrm{Aut}(\mathcal{C}^d)$ can map an edge of $\mathcal{C}^d$ comprising a pair of incident chords to an edge comprising a pair of intersecting non-incident chords. We see then that the sphere graph is not arc-transitive. However, $\mathcal{C}^d$ is vertex transitive, as $\mathrm{Möb}(\mathbb{S}^\mathrm{d})$ is transitive on the chords of $\mathbb{S}^d$, by Lemma 2.7.

Observation 3.2 Let us consider $h \in \mathrm{Aut}(\mathcal{C}^d)$. We already know that $\pi^{-1}(h)$ is a Möbius transformation on $\mathbb{S}^d$. Following Lemma 2.3, we may pick $d+1$ arbitrary points in $\mathbb{S}^d$. These necessarily will lie on a chord, as a $(d-1)$-sphere is defined by $d+1$ points. As long as we pick our $(d+2)^{th}$ point such that it is *not* on this chord, $\pi^{-1}(h)$ is *completely determined* by its action on these points. Thus h is completely determined by its action on the boundary cliques determining these points, and furthermore, is determined by its action on $d+2$ pairs incident at these points. So we only need to check finitely many pairs, and we will determine the entire structure of h. Considering how complex a graph $\mathcal{C}^d$ is - a graph with uncountably many vertices and uncountably many edges, containing uncountably many cliques - this is a surprising property. It tells us that the symmetries imposed by intersections and incidence restrict the possible automorphisms enormously.

Observation 3.3 Further to Observation 3.1, we see that automorphisms h distinguish intersecting, non-incident pairs of chords, and incident pairs of chords, by Lemmas 3.6 and 3.13 Every edge of $\mathcal{C}^d$ can therefore be labelled as such. Let us colour incident edges blue, and non-incident edges red. Then the subgraphs of $\mathcal{C}^d$ containing exactly the blue or red edges have the same automorphism group as $\mathcal{C}^d$, as every $h \in \mathrm{Aut}(\mathcal{C}^d)$ is also an automorphism of these subgraphs. It is important to stress $\mathcal{C}^d$ is not defined so that its edge set carries such information about the sort of intersection between two chords, so it is interesting to find that such information is indeed encoded in its edges. Note also that all prior observations apply to both of these subgraphs.

Chapter 4

Jordan curves on the 2-sphere

We now consider the intersection graph $\mathcal{J}$ of the simple closed curves on $\mathbb{S}^2$. Throughout this chapter, "curve" or "Jordan curve" will refer to the image of a continuous and injective map $\gamma\colon \mathbb{S}^1 \to \mathbb{S}^2$. We use "arc" or "Jordan arc" to refer to the image of a continuous, injective map $A\colon [0,1] \to \mathbb{S}^2$. We often abuse notation and refer to γ or A itself as the curve/arc, rather than its image. If we have two Jordan arcs A, A' and $A' \subset A$, we refer to A' as a *subarc* of A. The aim of this chapter is to prove:

Theorem (Theorem 4.2) $\pi\colon\ \mathrm{Aut}(\mathbb{S}^2) \to \mathrm{Aut}(\mathcal{J})$ is an isomorphism.

Again, $\pi(g)$ is the graph automorphism canonically induced by a homeomorphism g of $\mathbb{S}^2$; for a curve γ, $\pi(g)(\gamma) = g(\gamma)$. That is to say, the group of automorphisms of the intersection graph of Jordan curves on $\mathbb{S}^2$ is isomorphic to the group of all homeomorphisms of $\mathbb{S}^2$.

We begin with a short overview of Jordan curves.

4.1 The Jordan curve theorem and generalisations

The Jordan curve theorem states that a simple closed curve in the plane divides it into two disjoint, connected regions, the interior and exterior of the curve. The statement of this theorem intuitively seems obvious, but the general case requires very careful proof. A generalisation of the Jordan curve theorem to immersions of manifolds is known as the Jordan-Brouwer separation theorem [22]. In our case, we are only interested in the case of spheres, and so we provide a statement of the theorem below.

Theorem 4.1 (Jordan-Brouwer Separation theorem for d-spheres)**.** *The image X of an injective, continuous map $\gamma\colon \mathbb{S}^{d-1} \to \mathbb{S}^d$ separates $\mathbb{S}^d$ into exactly two connected* *components, each with X as their boundary.*

We will generally refer to these connected components as *regions* of the 2-sphere to align with our terminology from previous chapters.

The generalised Schönflies Theorem [23] states that the closure of each component as in Theorem 4.1 is homeomorphic to $\mathbb{D}^d$, as long as the closure is indeed a manifold. For $d = 2$, we run into no problems with the closure not being a manifold. However, for $d \geq 3$, this is not the case, the canonical counterexample in three dimensions being the Alexander horned sphere [24] - in which one region is not even simply connected. Furthermore, Bing showed in [25] that one can join together two horned spheres in such a manner that the resulting boundary between the two, K, is still homeomorphic to $\mathbb{S}^2$, but neither component of $\mathbb{S}^3 \backslash K$ is simply connected. For the remainder of this chapter we work only in the case $d = 2$.

A result equivalent to the generalised Schönflies theorem in two dimensions is what we will call the Jordan-Schönflies theorem. On $\mathbb{R}^2$, this states that any homeomorphism between two Jordan curves γ and λ extends to a homeomorphism of the whole of $\mathbb{R}^2$. In particular, given a collection of Jordan curves, there is a homeomorphism of the plane that maps one curve to the unit circle. The same is true on the 2-sphere; given any curve γ on $\mathbb{S}^2$, there is a global homeomorphism sending γ to an equator. There are several proofs of the theorem in the literature, C. Thomassen gives an interesting graph theory based proof in [26].

Furthermore, it is a consequence of the generalised Schönflies Theorem that any point of a Jordan curve is *curve-accessible* [26]. That is, given a point x on a Jordan curve γ and a point y not on γ we may find a Jordan arc with endpoints x, y which intersects γ only at x. Furthermore, due to a result of Alexander [22] the set of points which are *finitely accessible* is dense on a given Jordan curve. By finitely accessible, we mean that there is a Jordan arc of finite length satisfying the requirements for curve accessibility. Such access arcs allow us to construct curves to demonstrate several results in this chapter.

4.2 Pathological curves

It is a poor choice to rely on intuition when it comes to the study of Jordan curves. For example, there are many fractal, nowhere differentiable Jordan curves, such as the Koch snowflake, or Julia sets of certain complex polynomials [27]. Such curves do not admit a tangent bundle, and so local arguments are made trickier. Furthermore, there are curves of infinite length. Such curves also demonstrate that two Jordan curves may easily intersect at an infinite number of points, even an uncountable number of times. There are even Jordan arcs A such that $\mathbb{R}^3 \backslash A$ is not simply connected - as described by Fox & Artin [28]. Furthermore, there are nowhere differentiable Jordan curves with positive two-dimensional Lebesgue measure, generally referred to as *Osgood curves* [29][30]. It is possible to construct an Osgood curve of any two-dimensional Lebesgue measure $l \in (0, 1)$, via Knopp's triangle-elimination method [30].

As an aside, no Jordan curve may be space-filling; such curves cannot be injective as a consequence of Netto's Theorem [30].

4.3 Distinguishing $\mathcal{C}^2$ and $\mathcal{J}$ as graphs

Two distinct circles on $\mathbb{S}^2$ may only have a region count of three or four. But we may find two distinct curves γ, λ such that $\mathrm{Reg}(\{\gamma, \lambda\}) = n$ for any $n \geq 3$. This tells us that the intersection graphs of circles and Jordan curves are not isomorphic - recalling that region count is ultimately a property of the graphs. In and of itself this does not guarantee their automorphism groups differ, but instead motivates our investigation.

4.4 Generalising incidence and incident cliques

We begin by introducing generalised notions of incidence and incident cliques.

Definition 4.1 (Types of intersection) Suppose we have two curves γ and λ. Then they are *point-incident* if $\gamma \cap \lambda = \{x\}$ and *arc-incident* if $\gamma \cap \lambda$ is a Jordan arc, that is, they agree on some non-singleton, connected subset of the sphere. We say two curves intersect *transversely* at a point x if there is a homeomorphism $\phi\colon U \to \mathbb{D}^2$ such that γ and λ are mapped to two straight crossing lines, for U a neighbourhood of x. Equivalently, we say they intersect transversely at a point if we may find points of each curve in both sides of the other for all neighbourhoods of the point. Note that transverse intersection points must come in pairs, as long as we have only a finite number of such points.

On the sphere there is a choice of which region bounded by a curve should be considered its interior. In general our arguments will be symmetric whichever choice is made, but we will occasionally assume curves are orientated to ensure conciseness of exposition. Also note that what we term point-incidence is generally referred to as "tangency" or "touching" in the literature [31]. However, we continue to use the terminology of

Georgakopoulos in the name of consistency. Let us also call a curve *circular* if it is the non-empty intersection of a non-tangential plane with $\mathbb{S}^2$.

We now begin proving results that allow us to construct a proof of Theorem 4.2.

Lemma 4.1. *Two non-disjoint curves divide $\mathbb{S}^2$ into exactly three regions if and only if they are point- or arc-incident. Furthermore, these regions are each homeomorphic to the open disc.*

Proof. We first prove sufficiency of the condition. Take two such curves γ, λ and assume without loss of generality that both curves are orientated anti-clockwise. Then γ bounds two connected regions, call the region to the left of γ O_1 and the other O_2. By the generalised Schönflies theorem, both O_1 and O_2 are homeomorphic to the open disc. Then λ is contained entirely within O_1 or O_2 except for the incident point or arc. Note that λ also divides the sphere into two open discs, call the region to the left of λ O_3 and the right region O_4. If λ is contained in O_1, apart from the incident point or arc, then we have the following relations:

(i) $O_1 \cap O_2 = \varnothing$

(ii) $O_3 \cap O_4 = \varnothing$

(iii) $O_3 \subsetneq O_1$ (which implies $O_3 \cap O_2 = \varnothing$)

(iv) $O_2 \subsetneq O_4$

(v) $O_1 \cap O_4 \neq \varnothing$

These relations together imply that only O_3, $O_4 \cap O_1$ and O_2 are non-empty in $\mathbb{S}^2 \backslash (\gamma \cup \lambda)$. As stated earlier, we see that O_2 and O_3 are homeomorphic to the disc.

Let us now make use of Janiszewski's theorem [32] to show that $O_4 \cap O_1$ is connected. The theorem states that for two closed sets in $\widehat{\mathbb{R}}^2$ that any two points which may be connected via a path avoiding either set may be connected via a path avoiding both sets. We may apply the theorem to γ and λ on $\mathbb{S}^2$ making use of the fact that $\widehat{\mathbb{R}}^2$ is homeomorphic to the 2-sphere. Take two points $x, y \in O_4 \cap O_1$. Then there is certainly a path connecting x and y that avoids γ, as both points are in O_1. Hence we may find a path that connects x to y which intersects neither γ nor λ. Hence all three non-empty regions are indeed connected.

We will omit proof of the fact that $O_4 \cap O_1$ is homeomorphic to the open disc as we do not make use of this fact - however, a short argument via Alexander duality suffices. In fact, the use of Alexander duality provides a much shorter proof of the entire lemma.

An analogous argument shows we have three regions if λ minus the incident point or arc is instead contained in O_2.

To show that incidence is necessary, assume we have γ and λ such that $\mathbb{S}^2 \backslash (\gamma \cup \lambda)$ has three components. Consider $D = \gamma \backslash \lambda$, which is an open set. Thus we may express D as a countable union of open intervals on $\mathbb{S}^1$, as a Jordan curve is homeomorphic onto its image and any open set in $\widehat{\mathbb{R}}$ may be expressed as a countable disjoint union of intervals. Now suppose $\gamma \cap \lambda$ is not connected. Then choose a component I of D. Let p and q be the points of $\overline{I} \backslash I$. The curve λ splits into two arcs with endpoints p and q, neither of which are subsets of γ, and at least one such arc intersects γ only at these points. If there is only one such arc A, then $\eta = A \cup \overline{I}$ is a Jordan curve entirely contained in $\gamma \cup \lambda$ which bounds on one side a fourth region. If there are two such arcs, then exactly one provides the same result. In both cases we contradict our assumption, and so $\gamma \cap \lambda$ must be connected and thus is either a point or an arc. $\square$

We can then construct further curves incident to both curves at x within each component homeomorphic to the disc, as in the following lemma.

Lemma 4.2. *For a given Jordan curve γ and a point $x \in \gamma$, there exists another Jordan curve λ such that $\gamma \cap \lambda = \{x\}$.*

Proof. Given γ and x, apply a homeomorphism ϕ of $\mathbb{S}^2$ which transforms γ into a circle, by the Jordan-Schönflies theorem. Now take three distinct points a, b, c contained in a single open disc bounded by $\phi(\gamma)$. Then we may construct arcs $\phi(x)b$ and $\phi(x)c$ which intersect each other only at $\phi(x)$ and intersect $\phi(\gamma)$ only at $\phi(x)$. We then construct arcs ab and ac that intersect only at a and intersect the two initial arcs only at those points whose letters they share. The union of these four arcs, λ is a Jordan curve. Now apply ϕ^{-1}, and γ and $\phi^{-1}(\lambda)$ are two Jordan curves, incident at x. $\square$

Note that that λ as constructed in the proof of Lemma 4.2 can always be of finite length, but a homeomorphism ϕ^{-1} need not preserve rectifiability.

Lemma 4.3. *Let γ be a Jordan curve, and A be a Jordan arc such that A intersects γ only at its endpoints. Then $\mathbb{S}^2 \backslash (\gamma \cup A)$ has three connected components.*

Proof. The result is symmetric in the choice of interior of γ. Let us assume without loss of generality that γ is anti-clockwise orientated and that A lies in the closure of the region left of γ, and intersects γ at points b and c. Let L_1 and L_2 be the arcs in γ with endpoints b and c. Then $\lambda_i := A \cup L_i$ is a Jordan curve for $i = 1, 2$. Furthermore, λ_1

and λ_2 are arc-incident along A. Hence, $\mathbb{S}^2 \backslash (\gamma \cup A) = \mathbb{S}^2 \backslash (\lambda_1 \cup \lambda_2)$ has three connected components by Lemma 4.1. □

The previous lemma finds its use as a way to determine region count in later results. In particular, given a curve and an arc intersecting that curve only at its endpoints and which intersects no other curves, the arc divides the region it is contained in into two. We will refer to this as "repeated use" of the lemma, if multiple such arcs exist.

Lemma 4.4. *Two curves which intersect at exactly two points, and intersect at these points transversely, divide $\mathbb{S}^2$ into four regions.*

Proof. The result is a consequence of Lemma 4.3 applied twice. In particular, let the two points of intersection be x and y. Then γ and λ our curves, are formed of two arcs each; we assume our curves are anti-clockwise orientated without loss of generality. Let A_1 and A_2 be the arcs between x and y that are to the left or right of γ respectively. Let A_3 and A_4 be the arcs to the left and right of λ respectively, such that $\gamma = A_3 \cup A_4$ and $\lambda = A_1 \cup A_2$. Then $A_2 \cup A_4$ is a Jordan curve and A_1 is an arc satisfying the conditions of Lemma 4.3, when applied to $A_2 \cup A_4$. Subsequently, $A_1 \cup A_2$ and A_3 again satisfy the conditions of Lemma 4.3. This implies $\mathrm{Reg}(\{\gamma, \lambda\}) = 4$. □

It is now clear to see that the only way in which two intersecting Jordan curves can divide $\mathbb{S}^2$ into three is if they are point- or arc-incident. Point-incident pairs then suggest a clear path forwards for constructing homeomorphisms of $\mathbb{S}^2$, as they define a single point of intersection; the next question is whether this remains so under a graph automorphism of $\mathcal{J}$. To show this, we first show transverse pairs are preserved.

Lemma 4.5. *A pair of curves that intersect at exactly two points and do so transversely is preserved by any $h \in \mathrm{Aut}(\mathcal{J})$.*

Proof. Suppose that two curves γ and λ intersect at exactly two points x and y, and do so transversely at these points. In particular, this gives us four regions by the previous lemma, and so we must have four regions in the image under h. The only other way two curves may partition $\mathbb{S}^2$ into four regions is by being twice incident, that is $h(\gamma)$ and $h(\lambda)$ are incident to each other at two points, two arcs or a point and an arc. Let us consider the region graphs of these two configurations as illustrated in Figure 4.1. We see that $R_{\{\gamma,\lambda\}} = C_4$, while $R_{\{h(\gamma),h(\lambda)\}} = K_4 - e$ (the complete graph on four vertices minus an edge). But $C_4 \not\cong K_4 - e$ as graphs. We know from Lemma 3.4 that h should preserve region graphs, and so we have a contradiction. Thus h must preserve transverse intersection.

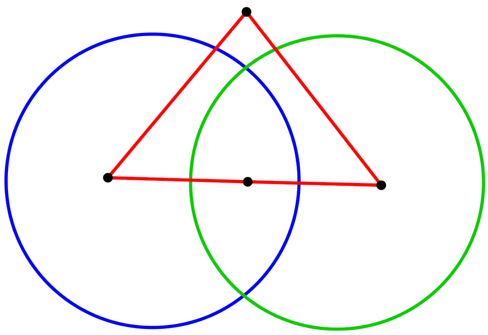

(a) Transverse pair

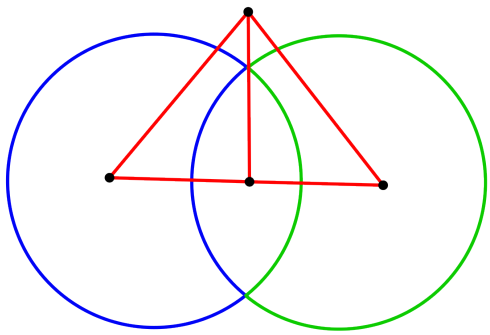

(b) Twice incident pair

Figure 4.1: Region graphs of a transverse pair and a double incident pair are not isomorphic. To see this, imagine removing the green or blue curve from each configuration, and add edges between regions which have merged.

□

Lemma 4.6. *Two curves incident at a single point or arc remain so under any* $h \in \mathrm{Aut}(\mathcal{J})$.

Proof. By Lemma 4.1 and the fact that h preserves adjacency, an incident pair must be sent to an incident pair. To expand on this, we know the only configurations of two adjacent curves that partition the sphere into three regions are an arc- or point-incident pair. As h must preserve region count, we see that neither pair can be mapped to an intersecting pair not arc- or point-incident. Secondly, such a pair cannot be mapped to a disjoint pair, as h is an automorphism of the intersection graph of such curves, and so must preserve intersection.

We now show the statement of the lemma. Suppose an arc-incident pair, γ, λ is sent to a point-incident pair by a graph automorphism h. Take two points, a, b, one in each of the two regions which have the incident arc as part of their boundary. Pick two points x, y on the arc. Then take four Jordan arcs ax, bx, ay and by such that $bx \cap ay = ax \cap by = \varnothing$ and bx, by and ax, ay intersect only at their common endpoints. To justify the existence of such arcs, we use the Jordan-Schönflies theorem [26]. In particular, there is a homeomorphism sending γ to a circle, and this homeomorphism may be extended to the entirety of $\mathbb{S}^2$. We may then construct two such arcs within the region bounded entirely by the circular image of γ. By performing a similar process with λ we see four such arcs exist. Then the union of these four arcs, η, is a Jordan curve intersecting the incident arc at two points, and intersecting γ and λ only at these two points - note that these are transverse intersections. This curve adds two new regions to the preimage by Lemma 4.3, so that $\text{Reg}(\{\gamma, \lambda, \eta\}) = 5$. But there is no such curve $h(\eta)$ in the image. To show this, firstly observe that $h(\eta)$ must intersect $h(\gamma)$ and $h(\lambda)$ at two distinct points each, and these intersections must be transverse, by Lemma 4.5. We recall that by hypothesis $h(\gamma)$ and $h(\lambda)$ are point-incident - now consider the possible cases.

(i) None of the intersection points of $h(\eta)$ with $h(\gamma)$ and $h(\lambda)$ are the incident point. In this case $\text{Reg}(\{h(\gamma), h(\lambda), h(\eta)\}) = 7 \neq 5$.

(ii) One of the intersection points is the incident point. Then $\text{Reg}(\{h(\gamma), h(\lambda), h(\eta)\}) = 6 \neq 5$.

In all cases we are using Lemma 4.3 repeatedly to calculate region count. In case (i) $h(\eta)$ has four distinct intersection points, and in case (ii) has only three. In all cases preservation of region count under h is violated, and thus an arc-incident pair must be mapped to an arc-incident pair. It follows that point-incident pairs are preserved. $\square$

We can now define generalisations of the incident tuples defined in Chapter 3.

Definition 4.2 (Incident tuples) A point-incident n-tuple at x is a collection of n curves where pairwise intersections are all equal to x. An arc-incident n-tuple on an arc A is a collection of n curves where pairwise intersections are all equal to A. We also define a *subarc-incident* n-tuple. This is a collection of n curves such that there is a maximum pairwise intersection $\gamma_i \cap \gamma_j = A$ and all other pairwise intersections are subarcs of A, or a point in A. Hence all arc-incident tuples are also subarc-incident tuples, and a point-incident tuple is simply a subarc-incident tuple wherein the maximal pairwise intersection is a single point. See Figure 4.2 for an example of each. We term a subarc-incident tuple that is neither a point- nor arc-incident tuple *proper*.

Note that all the incident n-tuples defined here have region count $n+1$. Furthermore, the region graph of a point-incident tuple is a *starlike tree* on $n+1$ vertices [33], a tree where at most one vertex is of degree greater than two (in Figure 4.2(a), we get P_{n+1}. In Figure 4.3 we instead get $K_{1,11}$, and Observation 4.1 makes clear why we get the starlike trees.). The region graph of an arc-incident n-tuple is P_{n+1}, and the region graph of a proper subarc-incident n-tuple is a tree on $n+1$ vertices. This aligns with the fact that all incident tuples are subarc-incident, as all region graphs here are themselves trees on $n+1$ vertices.

We first show that point-incident tuples are preserved by automorphisms of $\mathcal{J}$.

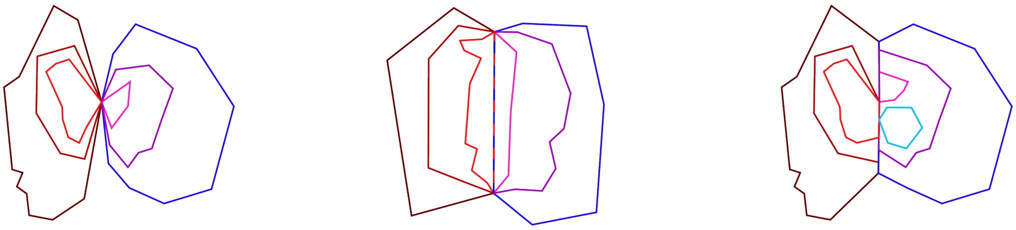

Figure 4.2: Example (a) is a point-incident 6-tuple, see Figure 4.3 for another example. Example (b) is an arc-incident 6-tuple and example (c) is a proper subarc-incident 7-tuple.

Lemma 4.7. *A point-incident n-tuple is preserved by any automorphism h of $\mathcal{J}$.*

Proof. We see that such an n-tuple must give us $n+1$ regions, by the logic of Lemma 4.1. The only configurations of n curves that divide the sphere into $n+1$ regions are the union of some or all of the following:

(i) d disjoint curves - meaning not adjacent to any other curve in the union.

(ii) p point-incident tuples, each containing r_p curves

(iii) a arc-incident tuples, each containing m_a curves

(iv) s proper subarc-incident tuples, each containing u_s curves

such that $\sum_{i=1}^{p} r_i + \sum_{i=1}^{a} m_i + \sum_{i=1}^{s} u_i + d = n$. Note that a single point-incident n-tuple is the case wherein $p = 1$, $r_1 = n$ and $a, s, d = 0$. Given that h preserves region count, the image must be such a union. If our n-tuple is mapped to a configuration containing some disjoint curves, then this contradicts that adjacency is preserved under h, so $d = 0$ and exactly one of $p, a, s = 1$ otherwise we contradict Lemma 4.6. Furthermore, by Lemma 4.6, no pair of curves in our point-incident n-tuple may be sent to an arc-incident pair, and thus $a, s = 0$ and $p = 1$. $\square$

Corollary 4.1. *An arc-incident n-tuple is preserved by any graph automorphism h of $\mathcal{J}$.*

Proof. It follows from Lemma 4.7 applied to h and h^{-1} that a subarc-incident n-tuple, S on an arc A is sent to a subarc-incident n-tuple S' on an arc B. We show that if S is in fact an arc-incident n-tuple, that S' is as well. Suppose towards a contradiction that S is an arc-incident n-tuple on the arc A, but that $S' := h(S)$ is only subarc-incident on an arc B. In S' we take a pair γ, λ whose intersection is B, and a curve η such that

$\eta \cap \gamma = \eta \cap \lambda \subsetneq B$. Call this intersection B', which may be an arc or a point. Note that such curves exist by the definition of a proper subarc-incident n-tuple. Now, we have some closed Jordan arc $B'' \subsetneq B \backslash B'$ such that $B'' \cap \eta = \varnothing$ and $B'' \subsetneq \gamma, \lambda$. Take two points $x, y \in B''$, and a third point a in a region whose boundary contains B'' and B' but not the open arc $\eta \backslash B'$. Then such a region is homeomorphic to the open disc by Lemma 4.3, and so we may construct arcs ax and ay as well as xy, so that they intersect only at their endpoints. Call the curve that is the union of these three arcs μ. Then under h^{-1}, μ must be arc-incident to A and disjoint from $h^{-1}(\eta)$. But this is impossible, as $A \subsetneq h^{-1}(\eta)$ due to S being an arc-incident n-tuple. Hence we have a contradiction and S' must be arc-incident. □

Corollary 4.2. *A proper subarc-incident n-tuple is preserved by any automorphism h of $\mathcal{J}$.*

Proof. This follows from Lemma 4.7 and Corollary 4.1. □

Given our results on finite incident collections of curves, we now show similar results for infinite collections.

Definition 4.3 (Point-incident cliques) A *point-incident* clique P_x is a maximal collection of curves pairwise point-incident at a point x.

By maximal, we mean that for any curve γ not in P_x, there exists some curve λ in P_x such that $\varnothing \neq \gamma \cap \lambda \neq \{x\}$. Note that we have infinitely many point-incident cliques at a given point, and these cliques are not disjoint. For example, start with a circle through a point x. We could construct an incident clique at x as in Chapter 2, which is a point-incident clique. Another choice would be to construct a point-incident clique at x using much more jagged curves. Both such cliques contain the circle, but differ everywhere else.

Before showing point-incident cliques offer a well-defined map of points under the action of a graph automorphism, we make an important observation on their structure.

Observation 4.1 In Chapter 2, recalling that each chord defined a unique circular curve, we saw that incident cliques necessarily had a nested structure. At first glance, the same does not appear to be true for general curves, we may have an arbitrary number of non-nested curves meeting at a single point, as demonstrated in Figure 4.3. We show that any point incident clique contains a sequence of at least n nested curves.

Suppose we have a point-incident clique with no nesting, so is formed of a number of spikes meeting at a point. By the logic of Lemma 4.1, at least one side of each of these spikes is homeomorphic to the open disc. We may therefore choose an individual curve, and place a Jordan curve - within the side homeomorphic to the open disc - that is incident at the same point. We may then repeat this process iteratively, placing new curves within the previously placed curve. This demonstrates that the initial arrangement was not maximal, and so did not in fact constitute a point-incident clique. This can be done for every spike, and so every curve in an incident clique will contain some nesting. Consequently, all point-incident cliques have sequences of nested curves of arbitrary length.

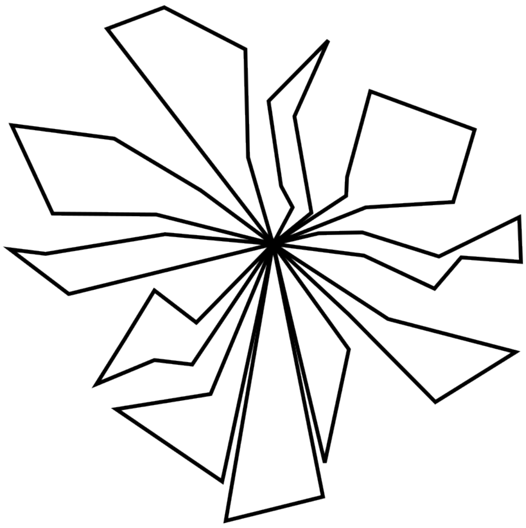

Figure 4.3: Illustration of "spikes" all incident at a point. The spikes may be far more complex than shown of course.

We now show several crucial properties of point-incident cliques, analogous to results in Chapter 3 for incident cliques, but requiring much more careful proof.

Lemma 4.8. *Point-incident cliques are preserved by any $h \in \mathrm{Aut}(\mathcal{J})$. Secondly, h acts injectively on distinct point-incident cliques. Thirdly, h acts surjectively on point-incident cliques. Furthermore, h preserves distinctness or sameness of the incident point; point-incident cliques at different points are mapped to different points, and distinct point-incident cliques at the same point are mapped to the same point.*

To clarify the statement of the lemma, by acting injectively, we mean that for $P_x \neq P_y$ at distinct points x, y, $h(P_x) \neq h(P_y)$ and for $P_x^1 \neq P_x^2$ both at a point x, $h(P_x^1) \neq h(P_x^2)$. By h acting surjectively, we mean given P_y there exists a unique P_x such that $h(P_x) = P_y$. We now proceed to a proof.

Proof. A point-incident clique is preserved as a consequence of Lemma 4.7; every finite subset of the clique must be sent to a point-incident tuple at the same point.

We now show that point-incident cliques P_x and P_y at distinct points cannot be mapped to distinct point-incident cliques at the same point z. Suppose towards a contradiction that this is the case for a given graph automorphism h. Take a curve $\gamma \in P_x$ such that $y \notin \gamma$ and a nested triple $\lambda_3 < \lambda_2 < \lambda_1 \in P_y$, where "$<$" indicates the left curve is nested within the right curve. Such a curve γ exists in P_x; if there were no such curve then all curves in P_x would contain y, contradicting the definition of a point-incident clique (P_x contains at least two curves by Observation 4.1). We justify the existence of a nested triple in P_y as a consequence of Observation 4.1, such a triple is a point-incident 3-tuple at y.

By the Jordan-Schönflies theorem, apply a global homeomorphism sending λ_1 to a circle λ_1'. This homeomorphism preserves the order of nestedness of the triple, and we retain a single common incidence point for the triple $\lambda_1', \lambda_2', \lambda_3'$ in the image. Call this point y'. Note that γ' - the image of γ - cannot contain y', and that this image curve is a closed set. So $y' \in \mathbb{S}^2 \backslash \gamma'$ which is an open set, and so we may find an open disc D centred on y' that does not intersect γ'. We may then find a small circular Jordan curve, which we shall call η', contained in this open disc which does not intersect γ' and which is incident to the image triple at y'. To justify this claim, λ_1' is a circle, and the boundary ∂D is also a circle. Hence the regions bounded within ∂D whose boundaries do not meet λ_2' - which exist due to preservation of nestedness - are homeomorphic to the open disc. Note that the image triple and η' together form a point-incident 4-tuple. Now apply the inverse homeomorphism and call the image of η' under this map η. Now we find that the λ_is are pairwise incident to η at y and that $\eta \cap \gamma = \varnothing$. Now apply h. The image of the pairwise incident 4-tuple consisting of the λ_is and η is incident at z. Furthermore, $z \in h(\gamma)$ as both $h(P_x)$ and $h(P_y)$ are equal to some point-incident cliques P_z^1, P_z^2 by supposition. But this provides a contradiction, as $h(\eta) \cap h(\gamma) \ni z$ as both are curves through z, which contradicts that h preserves adjacency of curves. Hence we conclude that the images of P_x and P_y under h must be point-incident cliques at distinct points.

This implies that h maps point-incident cliques at different points injectively, and as a consequence of the same being true for h^{-1} we conclude that every point-incident clique P_y is the image of some point-incident clique. Thusly, we see that h is bijective in its action on point-incident cliques at distinct points.

It remains to show that distinct point-incident cliques P_x^1 and P_x^2 are mapped to distinct point-incident cliques at a point y. Given a graph automorphism h, we see from the first part of the proof that h^{-1} cannot map point-incident cliques at different points to P_x^1 and P_x^2. Thus their images must be of the form P_y^1, P_y^2 for some point y. For distinctness, recall that the fact P_x^1 and P_x^2 differ implies, without loss of generality, that there is a curve $\gamma \in P_x^1$ such that $\gamma \notin P_x^2$. In particular, given that $x \in \gamma$, this means there exists $\lambda \in P_x^2$ such that $\gamma \cap \lambda \neq \{x\}$, otherwise this would contradict the maximality of P_x^2. If both are mapped to the same point-incident clique, then this contradicts Lemma 4.6 when applied to γ and λ. Given the surjectivity of h deduced in the first part of the proof, we conclude that every P_y at a fixed y is the image of some P_x at a fixed x.

We have shown the statement of the lemma, and in particular demonstrated that the action of h on point incident cliques is well-defined. □

Let us now focus on arc-incident cliques.

Definition 4.4 An arc-incident clique A_B is a maximal collection of curves such that for any pair of curves $\gamma, \lambda \in A_B$, $\gamma \cap \lambda = B$, where B is a Jordan arc.

Note that Observation 4.1 also applies to arc-incident cliques in general, so we can find nested sequences of arbitrary length in such a clique. We now show a result very similar to Lemma 4.8 for arc-incident cliques.

Lemma 4.9. *Arc-incident cliques are preserved by any $h \in \mathrm{Aut}(\mathcal{J})$. Given two distinct arc-incident cliques A_{B_1} and A_{B_2}, the following hold:*

(i) If $B_1 \cap B_2 = \varnothing$, then the same is true under h.

(ii) If $B_1 \neq B_2$ and $B_1 \cap B_2 \neq \varnothing$ then the same is true under h.

(iii) If $B_1 = B_2$ then the same is true under h.

and h maps (i), (ii) and (iii) bijectively.

Proof. The fact that an arc-incident clique is sent to an arc-incident clique by a graph automorphism follows from Corollary 4.1. We refer to the three arrangements of the incident arcs as type (i), type (ii) and type (iii). Our strategy is to show that

(1) Type (i) cannot be sent to type (iii).

(2) Type (ii) cannot be sent to type (iii)

(3) Type (i) cannot be sent to type (ii).

in this order. The proofs for all three share much with that of Lemma 4.8. Our arguments also apply to h^{-1}, which suffices to show the reverse direction of each case.

(1): In a similar vein to Lemma 4.8, take a nested triple $\lambda_1, \lambda_2, \lambda_3 \in A_{B_1}$ forming an arc-incident 3-tuple. Again, such a triple exists by the logic of Observation 4.1. Also take a curve $\gamma \in A_{B_2}$ such that $\gamma \cap B_1 \neq B_1$. Such a curve certainly exists, otherwise every curve in A_{B_2} intersects along B_1, contradicting the definition of an arc-incident clique.

Now apply a global homeomorphism ϕ which sends λ_1 to a circle λ_1' by the Jordan-Schönflies theorem, where we use the same notation as Lemma 4.8. Then γ' does not contain the incident arc of the image triple. Call this incident arc B_1' and take a point $x \in B_1' \backslash \gamma'$. Again we may take an open disc centred on x that does not intersect γ' as the complement of γ' is open. We can then construct a small curve incident to B_1' at a single point and which does not intersect γ'. Call this curve η'. Now apply ϕ^{-1} and h, supposing that h maps our cliques onto a single arc, B_3. Firstly, note that $\lambda_1, \lambda_2, \lambda_3, \eta$ form a proper subarc-incident 4-tuple, which is preserved by ϕ^{-1} and also by h, due to Corollary 4.2. Hence $h(\eta)$ is incident to B_3. But $B_3 \subset h(\gamma)$, which implies that $h(\eta) \cap h(\gamma) \neq \varnothing$, which contradicts that h preserves non-adjacency.

Hence type (i) and type (iii) cannot be interchanged under h.

(2): The proof is very similar to that of (1) and so we only provide a brief overview. Take a nested triple in A_{B_1} and a curve $\gamma \in A_{B_2}$ which does not contain B_1. We apply a global homeomorphism sending a member of the triple to a circle, and construct a small curve incident to our triple and disjoint from γ. This forms a proper subarc-incident 4-tuple which is then preserved by the inverse homeomorphism and h. If both cliques are mapped onto the same incident arc, then the constructed curve must intersect $h(\gamma)$, which is a contradiction.

Hence types (ii) and (iii) cannot be interchanged under h.

(3): Again, the proof is similar to that of (1).

Thus each type must be preserved by h. All that remains is to show h is bijective in its action. Clearly, h is certainly injective on types (i) and (ii), as they cannot be sent to arc-incident cliques on the same arc. Furthermore, h must be injective on type (iii) by the following argument. Take $\gamma \in A_B^1$ with $\gamma \notin A_B^2$, and $\lambda \in A_B^2$ such that $\gamma \cap \lambda \neq B$, so γ and λ are not arc-incident. But if $h(A_B^1) = h(A_B^2)$ then $h(\gamma)$ and $h(\lambda)$ must be arc-incident, contradicting Lemma 4.6. Surjectivity is then guaranteed by the injectivity of both h and h^{-1}. $\square$

A subarc-incident clique is a maximal collection of curves such that some pair of curves define a maximal intersection, and all other pairwise intersections are subarcs or points in the maximal intersection. We term such a clique *proper* if it is neither a point- nor arc-incident clique. Note that we have covered the case where both subarc-incident cliques are point-incident cliques in Lemma 4.8, and the case where both are arc-incident in Lemma 4.9. Preservation of such cliques is already a consequence of Corollary 4.2. The arguments of Lemmas 4.8 and 4.9 also suffice to show similar results on proper subarc-incident cliques, and so we will omit these to avoid repetition.

Before defining boundary cliques, we first show a stronger result on intersection preservation under graph automorphisms.

Lemma 4.10. *If two curves γ, λ are incident at n-points and m-arcs and intersect transversely at k points, then $h(\gamma), h(\lambda)$ under any $h \in \mathrm{Aut}(\mathcal{J})$ do so as well.*

Proof. Observe that the intersection points and arcs are cyclically ordered, as they occur around two simple closed curves. Let us first enumerate the transverse intersection points and incident points and arcs as $i_1, i_2, \dots i_{n+m+k}$.

We show that the number of transverse intersection points, the number of incident points and the number of incident arcs each remain unchanged under a graph automorphism h. Observe that $\mathrm{Reg}(\{\gamma, \lambda\}) = n + m + k + 2$. We call the resulting numbers of intersections and incidences under h k', n' and m'. If $n' + m' + k' \neq n + m + k$ then $\mathrm{Reg}(\{h(\gamma), h(\lambda)\}) = n' + m' + k' + 2 \neq n + m + k + 2$, a contradiction. So let us assume $n' + m' + k' = n + m + k$. If $k' \neq k$ then $R_{\{\gamma,\lambda\}}$ is not isomorphic to $R_{\{h(\gamma),h(\lambda)\}}$. While it is not especially difficult to see this immediately, we go into much more detail in Lemma 6.1.

Now we show $n' = n$ and $m' = m$, assuming $k' = k$, by considering cases. If $n' + m' \neq n + m$, then $\mathrm{Reg}(\{h(\gamma), h(\lambda)\}) = n' + m' + k + 2 \neq n + m + k + 2$, a contradiction. Therefore, we assume without loss of generality that $n' = n + a$ and $m' = m - a$, where $a \leq m$. We construct what we term *linking curves*. For now, let us assume m is even and $n, k \neq 0$. Construct curves $M_{j,j+1}$ where $j, j + 1 \mod n + m + k$ such that:

(i) $M_{j,j+1} \cap M_{j-1,j} = i_j$

(ii) $M_{j,j+1} \cap M_{j+a,j+a+1} = \varnothing$ for $a \not\equiv 1, -1 \mod n + m + k$.

(iii) $M_{j,j+1} \cap \gamma = i_j \cup i_{j+1}$

(iv) $M_{j,j+1} \cap \lambda = i_j \cup i_{j+1}$

Condition (i) implies that only linking curves between consecutive intersections or incidences intersect, and that they do so incidentally. Condition (ii) implies that no other pairwise intersections occur. Conditions (iii) and (iv) tell us that the linking curves intersect γ and λ at no other points. See Figure 4.4 for an example.

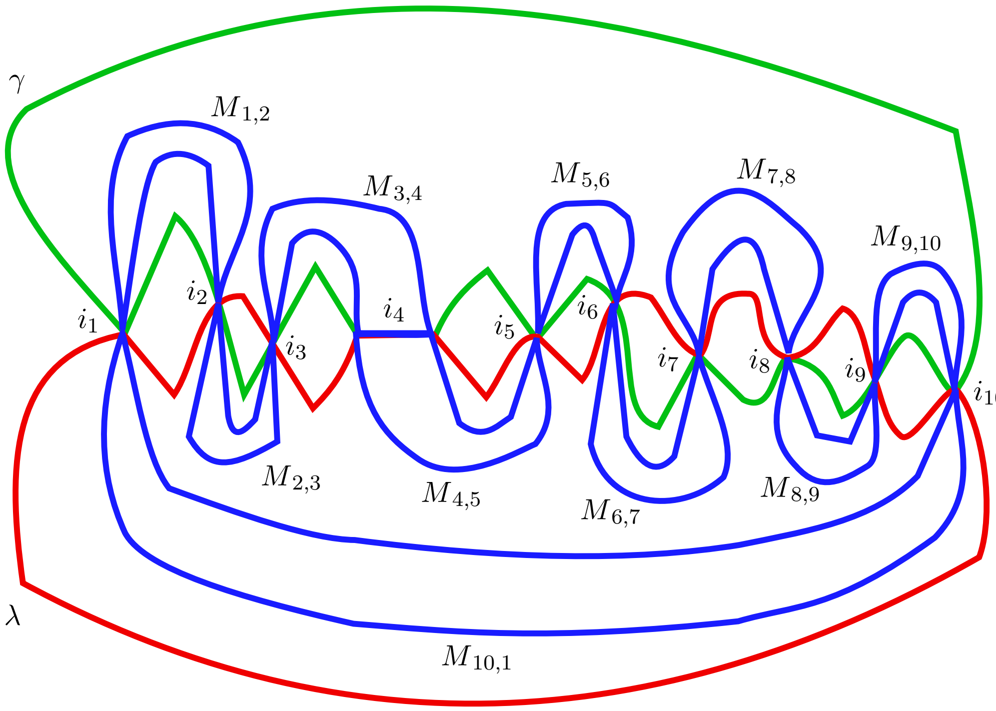

Figure 4.4: Example of linking curves for $k = 4$, $n = 5$ and $m = 1$. Note that linking curves only need to swap sides at an incident arc.

Under h, each linking curve remains incident only to those it intersected in the preimage, as h preserves adjacency. Furthermore, the type of incidence is preserved by Lemma 4.6, two linking curves incident on an arc remain so and the same is true for those incident at a point. Linking curves then allow us to show that $m' = m$ as follows. Firstly, note that two arc-incident linking curves can only be arc-incident along an incident arc of γ and λ. If this were not the case then we necessarily contradict condition (iii) or (iv). In the preimage we have two linking curves arc-incident at every incident arc. If $m' = m - a > 0$ then we must have more than two linking curves arc-incident to some incident arc in the image by the pigeonhole principle, which contradicts that non-consecutive linking curves are disjoint. If $m' = 0$ then we contradict that arc-incidence between linking curves is preserved. Hence $m' = m$ and so $n' = n$; similar logic applies if $n' = n - a$ and $m' = m + a$ for $a \leq n$.

If m is instead odd and both $n, k = 0$, then we cannot construct linking curves satisfying our conditions. At each incident arc we must swap sides to avoid transverse intersection, and so swapping an odd number of times leaves a curve having to intersect at least one other curve transversely. Instead, we may construct linking curves avoiding a single incident arc. Doing this for every such subset of the incident arcs shows that $m - 1 = m' - 1$ on all such subsets, and so $m = m'$ must hold. □

4.5 Generalising boundary cliques

As in Chapter 3, we could choose a representative point-incident clique at each point x to then construct a homeomorphism from. We instead proceed to use boundary cliques to streamline the proof at the end of this section.

Definition 4.5 (Boundary clique) We define a boundary clique to be the union of all point-incident cliques at x.

By our earlier results on incidence, we may find a curve point-incident to any curve through x, at x. Thus every curve through x is indeed in a point-incident clique at x. Hence a boundary clique at x does contain every curve through x, and so is unique. We could equivalently define a boundary clique at x as the collection of all curves containing x, but this definition highlights that point-incident cliques are the fundamental building blocks of boundary cliques.

Lemma 4.11. *Boundary cliques are preserved by graph automorphisms, and are mapped bijectively by any $h \in \mathrm{Aut}(\mathcal{J})$.*

Proof. Let us take a boundary clique K_x. Then $h(K_x) = K_y$ for some y, as $h(K_x) = h(\bigcup P_x) = \bigcup h(P_x) = K_y$. Given K_y, take an incident pair $\gamma, \lambda \in K_y$. Then $h^{-1}(\{\gamma, \lambda\})$ is contained in some point-incident clique at a point x. It follows that $h^{-1}(K_y) = K_x$, and so surjectivity is satisfied. Furthermore, $h(K_x) \neq h(K_z)$ for all $x \neq z$, by an argument analogous to that in the proof of Lemma 4.8. In particular, we can take a nested point-incident 3-tuple in K_x, and a curve K_y which does not contain x. Then we can construct a small curve incident to the triple which does not intersect the curve in K_y.

If $h(K_x) = h(K_y)$ then this implies these two curves intersect, which is a contradiction. Hence we have shown bijectivity of h on boundary cliques. $\square$

We see that boundary cliques still offer us a well-defined map of $\mathbb{S}^2$. We now utilise this fact to prove the main theorem of this chapter.

Theorem 4.2. $\pi\colon \operatorname{Aut}(\mathbb{S}^2) \to \operatorname{Aut}(\mathcal{J})$ *is an isomorphism.*

Note that a homeomorphism g necessarily sends a Jordan curve to a Jordan curve, as $g(\gamma) : \mathbb{S}^1 \to \mathbb{S}^2$ is injective and continuous.

Proof. To see that π is injective, observe that $\ker(\pi) = \mathrm{id}$. Furthermore, to show that π is surjective, we construct a homeomorphism h' such that $\pi(h') = h$, for a given $h \in \operatorname{Aut}(\mathcal{J})$. Construct such an h' from h by its action on boundary cliques; $h'(x) = y$ for $h(K_x) = K_y$.

We must show h' is indeed a homeomorphism. Firstly, h' is bijective by Lemma 4.11 and so we may consider its inverse h'^{-1}. For continuity, we show the preimages of open discs on $\mathbb{S}^2$ are open, this suffices as the discs form a basis for the standard topology on $\mathbb{S}^2$. So take an open disc, D, on $\mathbb{S}^2$. Take the boundary of the disc, ∂D, which is a simple closed curve itself. We see that $h'^{-1}(\partial D)$ is a simple closed curve, and is thus a closed set, and hence divides the 2-sphere into two open regions. Then by our earlier results, the open disc D must be sent to one of these regions by h'^{-1}. Hence we see that $h'^{-1}(D)$ is open and as $\mathbb{S}^2$ is Hausdorff and compact, h' is a homeomorphism.

Finally, for g', h' homeomorphisms such that $\pi(g') = g$ and $\pi(h') = h$ for $g, h \in \operatorname{Aut}(\mathcal{J})$ and $\gamma \in V(\mathcal{J})$, we have that:

$$\pi(g' \circ h')(\gamma) = (g' \circ h')(\gamma) = g'(h'(\gamma)) = g(h(\gamma)) = (g \circ h)(\gamma)$$

which demonstrates that π respects composition. $\square$

Chapter 5

Smooth Jordan curves on the 2-sphere

In this chapter we consider the intersection graph, $\mathcal{S}^\infty$, of the simple smooth - meaning C^∞ - closed curves on the 2-sphere. We then consider the intersection graph $\mathcal{S}^k$ of the C^k simple closed curves on $\mathbb{S}^2$. Such curves are the image of an injective, continuous, C^k function $\gamma\colon \mathbb{S}^1 \to \mathbb{S}^2$, where $k \in \mathbb{Z}^+ \cup \{\infty\}$. Throughout this chapter, "curve" will refer to a simple closed curve, where its differentiability class is clear from context. We often refer to C^∞ curves as "smooth". We show the group of homeomorphisms naturally associated to the automorphisms of $\mathcal{S}^k$ is strictly larger than the group of C^k-homeomorphisms of $\mathbb{S}^2$ for $k \in \mathbb{Z}^+ \cup \{\infty\}$.

5.1 Smooth manifolds

In order to have a notion of smoothness on the 2-sphere, we require that we consider $\mathbb{S}^2$ as a *smooth manifold*. This is a rich and fascinating area, but is not the focus of this thesis and so we only give a basic overview based on John M. Lee's *Introduction to Smooth Manifolds*[17]. In the simplest terms, a topological space M is a manifold if it is Hausdorff, second countable, and is locally homeomorphic to Euclidean $\mathbb{R}^d$. To clarify the third point, this means that every point $p \in M$ has an open neighbourhood, U, homeomorphic to some open subset of $\mathbb{R}^d$. A pair (U, ϕ) is called a *chart*, where ϕ is the aforementioned homeomorphism and U is called the chart's domain. Two charts, (U, ϕ), (V, ψ), are *smoothly compatible* if either $U \cap V = \varnothing$, or $U \cap V \neq \varnothing$ and the map $\psi \circ \phi^{-1}$ is smooth as a map on $\mathbb{R}^d$. An *atlas*, $\mathscr{A}$ is a collection of charts whose domains cover M, and is smooth if any two charts in $\mathscr{A}$ are smoothly compatible. We can then define a maximal smooth atlas on M, a smooth atlas not contained in any larger smooth atlas. Then a smooth manifold is the pair $(M, \mathscr{A})$, where M is a topological manifold, and $\mathscr{A}$ is a maximal smooth atlas. In general it is unnecessary to specify a maximal smooth atlas on M, as such an atlas is unique for a given smooth atlas.

For $\mathbb{S}^2$, a common smooth atlas is given by two stereographic projections from the north and south poles. That is, we take an open cover consisting of the two sets $\mathbb{S}^2 \backslash (0,0,1)$ and $\mathbb{S}^2 \backslash (0,0,-1)$, and homeomorphisms as stereographic projections from each removed point. Then this gives us two charts of $\mathbb{S}^2$ and it can be easily checked that these charts are smoothly compatible.

Given a map $f : M \to N$ where M and N are smooth manifolds, f is smooth if, for every $p \in M$, there exist smooth charts $(U, \phi) \ni p$ and $(V, \psi) \ni f(p)$ such that $f(U) \subseteq V$ and $\psi \circ f \circ \phi$ is smooth from $\phi(U)$ to $\psi(V)$. In our case, we are looking at maps $\gamma : \mathbb{S}^1 \to \mathbb{S}^2$ and so we care about smoothness of $\psi \circ \gamma \circ \phi^{-1} : \mathbb{R} \to \mathbb{R}^2$, where ϕ and ψ are stereographic projections identifying open neighbourhoods of $\mathbb{S}^d$ with $\mathbb{R}^d$ for $d = 1, 2$ respectively. Hence, when we describe a Jordan curve γ as C^k, for $k \in \mathbb{Z}^+ \cup \{\infty\}$, we ultimately mean the prior composition is C^k as a map from $\mathbb{R}$ to $\mathbb{R}^2$.

5.2 Distinguishing $\mathcal{C}^2$, $\mathcal{J}$ and $\mathcal{S}^\infty$ via region counting

Using the argument of §4.3, we see that $\mathcal{S}^\infty$ and $\mathcal{C}^2$ are not isomorphic as graphs.

Furthermore, $\mathcal{S}^\infty$ and $\mathcal{J}$ are not isomorphic. Let us take a point-incident n-tuple of smooth curves, with $n \geq 5$. Then the region graph of such a collection is P_{n+1}, the path on $n+1$ vertices. However, this need not be true for general curves. Consider the example in Figure 5.1, which is a point-incident 5-tuple with region graph $K_{1,5}$.

5.3 The intersection graph of C^∞ curves on $\mathbb{S}^2$

As our curves are smooth, we may define the tangent space of a given curve. We say that two smooth curves intersect *transversely* at a point if the tangent lines are linearly independent and *tangentially* at a point if the tangent lines are equal. If two curves intersect exclusively transversely, we will say they intersect transversely, otherwise we will refer to their intersection as tangential. When we refer to an arc, we mean an open curve. All properties proved for automorphisms of $\mathcal{J}$ are true of automorphisms of $\mathcal{S}^\infty$. Note that all properties of $h \in \mathrm{Aut}(\mathcal{J})$ when applied to general curves are true of $h \in \mathrm{Aut}(\mathcal{S}^\infty)$ when applied to C^∞ curves. We therefore list all such properties but

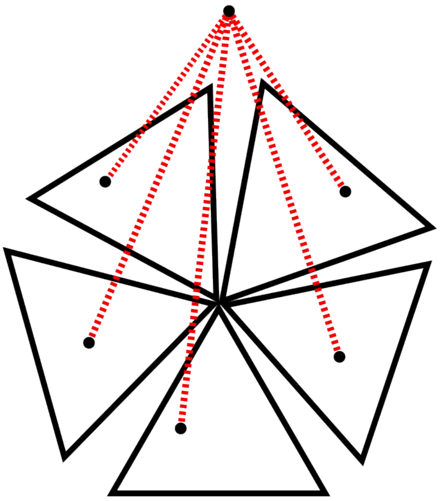

Figure 5.1: Example of a non-path region graph for general curves

omit proofs, except those that demonstrate the differences between smooth and general curves.

Lemma 5.1. *Suppose $\gamma \neq \lambda$ are two simple smooth closed curves, with $\gamma \cap \lambda \neq \varnothing$. Then γ, λ partition $\mathbb{S}^2$ into three regions if and only if $\gamma \cap \lambda$ is path connected.*

Proof. See Lemma 4.1. □

Lemma 5.2. *All graph automorphisms h of $\mathcal{S}^\infty$ preserve arc-incident pairs.*

Proof. See Lemma 4.6. □

Corollary 5.1. *Graph automorphisms of $\mathcal{S}^\infty$ preserve point-incidence.*

Proof. See Lemma 4.6. □

Corollary 5.2. *Graph automorphisms of $\mathcal{S}^\infty$ preserve transverse intersection points.*

Proof. See Lemma 4.5. □

We then define *point-incident cliques* and *arc-incident cliques*.

Definition 5.1 (Point-incident cliques) Given a point $x \in \mathbb{S}^2$, a point-incident clique at x is a maximal collection of curves such that $\gamma_i \cap \gamma_j = \{x\}$ $\forall i, j \in I$ our indexing set.

Definition 5.2 (Arc-incident cliques) Given a smooth arc $A \in \mathbb{S}^2$, an arc-incident clique at A is a maximal collection of curves such that $\gamma_i \cap \gamma_j = A$ $\forall i, j \in I$ our indexing set.

Lemma 5.3. *The image of either clique under an automorphism of $\mathcal{S}^\infty$ is another such clique at either a point or an arc respectively.*

Proof. See Lemmas 4.8 and 4.9 $\square$

Lemma 5.4. *Two smooth curves intersecting transversely at k points, incidentally at n points and at m arcs do so under $h \in \mathrm{Aut}(\mathcal{S}^\infty)$.*

Proof. See Lemma 4.10. $\square$

We then define our *boundary cliques*, and show our standard results.

Definition 5.3 A boundary clique at x, K_x, is the union of all point-incident cliques at x.

Clearly a boundary clique contains all curves through x as every such curve is in some point-incident clique at x.

Lemma 5.5. *An automorphism of $\mathcal{S}^\infty$, h, maps boundary cliques bijectively.*

Proof. Follows the logic of the proof of Lemma 4.11 $\square$

We now denote the group of homeomorphisms of $\mathbb{S}^2$ that preserve smooth curves in both the forward and backward direction as $\mathrm{Aut}^\infty(\mathbb{S}^2)$. We then prove the main theorem of this section,

Theorem 5.1. *The map $\pi\colon \mathrm{Aut}^\infty(\mathbb{S}^2) \to \mathrm{Aut}(\mathcal{S}^\infty)$ is an isomorphism, where $\mathrm{Aut}^\infty(\mathbb{S}^2)$ is the group of homeomorphisms of $\mathbb{S}^2$ such that h and h^{-1} preserve the class of C^∞ curves, for $h \in \mathrm{Aut}^\infty(\mathbb{S}^2)$.*

and we explicitly show that there are homeomorphisms of $\mathbb{S}^2$ that preserve all smooth curves but which are not smooth themselves, via an argument of Le Roux and Wolff [7]. In particular, we denote the group of smooth homeomorphisms of $\mathbb{S}^2$ as $\mathrm{Diff}^\infty(\mathbb{S}^2)$, and show the following:

Proposition 5.1. $\mathrm{Diff}^\infty(\mathbb{S}^2) \subsetneq \mathrm{Aut}^\infty(\mathbb{S}^2)$

Let us first prove Theorem 5.1.

Proof of Theorem 5.1. Firstly, $\ker(\pi) = \mathrm{id}$, and so π is injective. We know from Lemma 5.3 that for a boundary clique at x, K_x, $h(K_x) = K_y$ for some unique $y \in \mathbb{S}^2$. Define a map h' by $x \mapsto y$. Then h' is injective, no two boundary cliques are mapped to the same point. Furthermore, h' is surjective, as every K_y is the image of a K_x under h. To show h' is continuous, use the argument on open discs as in the proof of Theorem 4.2. Hence as $\mathbb{S}^2$ is Hausdorff and compact, h' is a homeomorphism. We therefore find that h' is a homeomorphism which preserves smooth Jordan curves. Furthermore, $\pi(h') = h$, and so π is surjective. The map π also respects composition as in the proof of Theorem 4.2. $\square$

Let us now show Proposition 5.1.

Proof of Proposition 5.1. We modify a construction of Le Roux and Wolff [7]. Let $g\colon \mathbb{R} \to \mathbb{R}$ be a smooth diffeomorphism supported in the segment $[1/2, 2]$. That is to say, $g(x) = x$ for all $x \notin [1/2, 2]$; we suppose however that $g(1) \neq 1$. We consider the map $F\colon \mathbb{R}^2 \to \mathbb{R}^2$ defined as

$$F(x,y) = \begin{cases} (x, xg(y/x)) & \text{if } x \neq 0 \\ (x,y) & \text{if } x = 0 \end{cases} \tag{5.1}$$

This map, as well as its inverse, is obviously smooth in restriction to $\mathbb{R}^2 \backslash (0,0)$. But we find that $dF_{(0,0)}(1,0) = (1,0)$, $dF_{(0,0)}(0,1) = (0,1)$ but $dF_{(0,0)}(1,1) = (1, g(1)) \neq (1,1)$. Hence we see that the differential of F at $(0,0)$ is not linear, and so F is not differentiable at the origin.

We introduce a bump function, allowing us to transfer F onto $\mathbb{S}^2$. Define $\chi\colon \mathbb{R}^2 \to [0,1]$ such that $\chi = 1$ on a disc $\mathbb{D}(0,r) \subset \mathbb{R}^2$ and $\chi = 0$ on $\mathbb{R}^2 \backslash \mathbb{D}(0,2r)$ and χ is smooth (See [17] for explicit constructions). Then define $F_\chi(x,y) = \chi(x,y)F(x,y) + (1 - \chi(x,y))(x,y)$. Then F_χ is smooth except at the origin, equal to F within a disc, and equal to the identity outside a larger disc, furthermore, it transitions smoothly between the two.

Now let $(U, \phi\colon \mathbb{S}^2 \to \mathbb{R}^2)$ be a stereographic projection where U contains the north pole, as in §5.1. Then define the map

$$\tilde{F}(p) = \begin{cases} \phi^{-1} \circ F_\chi \circ \phi(p) & \text{if } p \in U \\ p & \text{if } p \notin U \end{cases} \tag{5.2}$$

Then this gives us a well-defined homeomorphism of $\mathbb{S}^2$, smooth everywhere but the north pole. Now let $\gamma\colon \mathbb{S}^1 \to \mathbb{S}^2$ be a smooth curve. If the image of γ does not contain the north pole, N, then $\tilde{F} \circ \gamma$ is clearly still a smooth curve. Now suppose without loss of generality that $\gamma(0) = N$, where we are parametrising $\mathbb{S}^1$ via $\widehat{\mathbb{R}}$. If $\gamma'(0) \notin [1/2, 2]$ then $\tilde{F} \circ \gamma$ and γ are locally the same around N, in U. That is, for $\mathbb{S}^1 \ni t \to 0$, $\tilde{F} \circ \gamma(t) = \gamma(t)$

for small enough t. Otherwise, we may write $\gamma(t) = (t, \alpha(t))$ up to reparametrisation, for t near 0, where α is a smooth map, and $\alpha(0) = 0$. Then we have

$$\tilde{F} \circ \gamma = (t, tg(\alpha(t)/t))$$

for small t. So as the map $t \to \alpha(t)/t$ for $t \neq 0$ and $t \to \alpha'(0)$ when $t = 0$ is smooth, we see that $\tilde{F} \circ \gamma$ is smooth.

So we have found a function that maps all smooth curves on $\mathbb{S}^2$ to smooth curves, but which is not smooth itself. So we may conclude that $\mathrm{Diff}^\infty(\mathbb{S}^2) \subsetneq \mathrm{Aut}^\infty(\mathbb{S}^2)$. $\square$

5.4 The intersection graph of C^k curves on $\mathbb{S}^2$

As in the previous section, all results proven in Chapter 4 hold for C^k curves, for $k \in \mathbb{Z}^+$. So in particular, we may define boundary cliques, and these boundary cliques are mapped bijectively. We may distinguish each $\mathcal{S}^k$ $(k \geq 1)$ from $\mathcal{J}$ and $\mathcal{C}^2$ by the argument of §5.2. We have that $\mathcal{S}^0 = \mathcal{J}$, and so we will assume $k \geq 1$ for the remainder of this chapter. Within this section, we prove the following theorem

Theorem 5.2. *The map* $\pi\colon \mathrm{Aut}^k(\mathbb{S}^2) \to \mathrm{Aut}(\mathcal{S}^k)$ *is an isomorphism, where* $\mathrm{Aut}^k(\mathbb{S}^2)$ *is the group of homeomorphisms of* $\mathbb{S}^2$ *such that* h *and* h^{-1} *preserve the class of* C^k *curves, for* $h \in \mathrm{Aut}^k(\mathbb{S}^2)$.

We then prove the following proposition

Proposition 5.2. $\mathrm{Diff}^k(\mathbb{S}^2) \subsetneq \mathrm{Aut}^k(\mathbb{S}^2)$

where we denote by $\mathrm{Diff}^k(\mathbb{S}^2)$ the group of C^k homeomorphisms of the 2-sphere.

Proof of Theorem 5.2. The proof is exactly that of Theorem 5.1. $\square$

Proof of Proposition 5.2. Observe that the map $\tilde{F}$ defined in the proof of Proposition 5.1 also preserves all C^k curves, but is not itself C^k. Hence this provides an example of a map in $\mathrm{Aut}^k(\mathbb{S}^2) \setminus \mathrm{Diff}^k(\mathbb{S}^2)$. $\square$

Let us now prove a further result which demonstrates that each $\mathcal{S}^k$ is distinct.

Proposition 5.3. *For every* k*, there is a function that preserves* C^k *curves, but not* C^{k+1} *curves. Hence* $\mathrm{Aut}^{k+1}(\mathbb{S}^2) \subsetneq \mathrm{Aut}^k(\mathbb{S}^2)$ *for all* k*.*

Proof. We modify the proof of Proposition 5.1. Let us first define $f_k(x) = x^k|x|$, for $k \geq 1$. Then f_k is C^k but not C^{k+1}, which can be deduced from the one sided derivatives for $x \geq 0$ and $x < 0$. Let us then modify $F(x,y)$, defining

$$F_k(x,y) = \begin{cases} (x,(x+f_k(x))g(y/x)) & \text{if } x \neq 0 \\ (x,y) & \text{if } x = 0. \end{cases} \tag{5.3}$$

Then F_k is easily seen to be continuous along $x = 0$, and by similar arguments to those applied to F, we see that F_k preserves C^k curves through the origin, but not C^{k+1} curves. As an explicit example, we can take a C^{k+1} curve which is equal to $y = mx$ near 0 such that $m \in [1/2, 2]$. Then $F_k(x,y) = (x, xg(m)) + (0, f_k(x)g(m))$. The second term is not C^{k+1}, and so not all C^{k+1} curves are preserved by F_k. We can then apply a bump function and a stereographic projection to get an analogous map $\tilde{F}_k$ on $\mathbb{S}^2$. $\square$

As a consequence of this proposition, each $\mathcal{S}^k$ is not isomorphic to $\mathcal{S}^l$ for any finite $l \neq k$; graphs with differing automorphism groups cannot be isomorphic. Furthermore $\mathcal{S}^\infty$ is not isomorphic to any $\mathcal{S}^k$ for finite k, as each $\tilde{F}_k$ is in $\mathrm{Aut}^k(\mathbb{S}^2)$ but not $\mathrm{Aut}^\infty(\mathbb{S}^2)$. From this and the earlier sections distinguishing our different intersection graphs of $\mathbb{S}^2$, we see that no pair of $\mathcal{C}^2$, $\mathcal{J}$, $\mathcal{S}^\infty$ or $\mathcal{S}^k$ are isomorphic as graphs for any $k \geq 1$.

Chapter 6

On region graphs

Region graphs provide a novel and interesting method of construction for a graph. We investigate the properties of region graphs on different classes of curves on the 2-sphere. Throughout we assume that all such graphs are finite, meaning that collections of curves are finite, and no two curves intersect infinitely many times. We classify some simple region graphs on collections of general curves which serves to demonstrate their incredible complexity. Thereafter we focus on region graphs of circles and demonstrate that this is a more tractable problem. When we discuss region graphs on $\mathcal{C}^2$, we rely on the fact that every chord in $\mathcal{C}^2$ as defined in Chapter 2 defines a unique circle embedded in $\mathbb{S}^2$.

6.1 Motivating the study of region graphs

Recall that region graphs are defined by "adding" removed curves back into the intersection graph of some class of curves on $\mathbb{S}^2$. In the geometric sense, this corresponds to taking a collection of closed curves $\Gamma = \{\gamma_i\}$ with $i \in I$ and taking a vertex for every topological component of $\mathbb{S}^2 \backslash \Gamma$. Now remove each γ_j with replacement and add edges between all pairs of topological components which have merged. Repeat this for every element of Γ, recalling that the graph should remain simple. The resulting graph is the region graph R_Γ.

We demonstrated in Lemma 3.4 that each automorphism of the larger intersection graph induces an automorphism of the region graph. It is thus useful to study region graphs as invariants of this action.

Region graphs are also interesting from a purely graph-theoretic perspective; we will see definite structure within these graphs, especially within those on circular curves. We show several interesting facts about region graphs, demonstrating that not all graphs are region graphs, and classifying certain simpler classes of region graphs. Furthermore, it shall be seen that there are ways to iteratively construct region graphs of $\mathcal{C}^2$, which is an unexpected result.

6.2 Necessary conditions for adjacency in region graphs

As we will later see, region graphs may be extremely complicated. It is therefore helpful to understand their structure. To do so, we introduce a necessary condition for two vertices to be adjacent in a region graph.

6.2.1 Codes

Suppose we have a finite collection of r distinct curves $\{\gamma_i\}_{i\in[1,r]}$. Every vertex of a region graph corresponds to a distinct topological component of $\mathbb{S}^2$. We recall that the interior of a curve may be chosen arbitrarily on the 2-sphere. Make such a choice for all curves in the collection. Then every region can be identified by an r-digit binary number. Such a code looks like $1001\ldots0110$; this particular code tells us that the region is in the interior of γ_1, the exterior of γ_2 and so on. In particular, if two regions' codes differ at n digits, then they do so in all choices of interior; exchanging the interior and exterior of γ_j corresponds to flipping the j^{th} digit of every code. Different regions may have the same code, but for two regions to merge, we must be able to make their codes the same by removing a single curve. Removing the curve γ_j corresponds to deleting the j^{th} index of every vertex's code. So a necessary condition for two vertices being adjacent is that their codes differ in at most one index.

6.2.2 Codes are not sufficient

We have shown that two regions are adjacent only if their codes differ by at most one digit. The converse is false, as it is possible to have two regions with the same code that are not adjacent, see Figure 6.1a.

In Figure 6.1a, the two hatched regions have the same binary code. To see this, suppose that this arrangement of curves is contained in a small neighbourhood on $\mathbb{S}^2$ homeomorphic to $\mathbb{R}^2$. Then choose interiors as for Jordan curves on the plane. Let us call the purple curve γ_1, the blue γ_2, the green γ_3 and the red γ_4; then the hatched regions have code 1100. As shown above this implies they have the same code for any choice of interiors. We demonstrate that they cannot be adjacent in the region graph on the four curves. If we remove the red curve, then they are still separated by the green

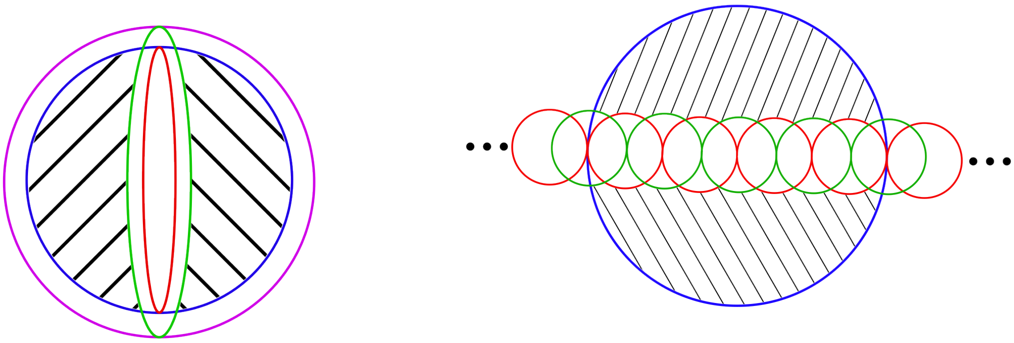

(a) Example of two non-adjacent regions with the same code

(b) Example of two non-adjacent regions with the same code for curves in $\mathcal{C}^2$

Figure 6.1: Arrangements with non-adjacent regions that share a code.

curve and vice versa. If we remove the purple curve, the two regions are still separated by the red and green curves. If we remove the blue curve, then the two regions are contained in two separate regions, bounded by the green and purple curves. So we have demonstrated that no single curve can be removed to merge the two hatched regions.

One notices that we have used non-circular curves in Figure 6.1a, and so the question arises as to whether we can run into the same problem using only circles on $\mathbb{S}^2$. The answer is yes, by considering the arrangement of circles on $\mathbb{S}^2$ in Figure 6.1b. Form an incident chain of circles around an equator. That is to say, we have a sequence of n circles (C_n) around the equator such that $C_i \cap C_{i+1}$ is a singleton with i mod n, and all other pairwise intersections are empty. Then duplicate this sequence and shift it so that the centre of each circle in the duplicate sequence is an incident point in the initial sequence. Now take a larger circle across the equator. Then there are two regions within the larger circle with same code.

In Figure 6.1b the red and green circles represent the previously described sequences on an equator, and the hatched regions have the same code but are not adjacent in the corresponding region graph. If we remove a single red or green circle then the hatched regions are still separated by the chain of the opposite colour. If we remove the blue circle, then both chains still separate the hatched regions.

6.3 Region graphs on $\mathcal{J}$

We begin our results by analysing region graphs on simple closed curves on $\mathbb{S}^2$, as in Chapter 4. In particular, we classify all region graphs on two curves, and discuss a few graph classes which are contained in the class of region graphs on a finite number of closed curves. A finite collection of curves corresponds to a finite induced subgraph of $\mathcal{J}$.

6.3.1 On two curves

Throughout this subsection, we analyse the region graphs of two curves $\gamma, \lambda \in \mathcal{J}$. We denote the region graph of two curves that are incident at n points and m arcs, and intersect transversally at $2k$ points $R_{n+m,k}$. All such incidences and intersections must be disjoint, otherwise we have a self-intersection. See Lemma 4.10. It will become clear that on general curves, region graphs may be incredibly complicated, and so we only prove results in a few particular cases.

In order to classify $R_{n+m,k}$, we need to introduce the *join* of two graphs. Given $G = (V_1, E_1)$ and $H = (V_2, E_2)$, $G \times H$ is the graph with vertex set $V_1 \cup V_2$ with all edges in $E_1 \cup E_2$ and all vertices of G adjacent to all vertices of H. Furthermore, we denote by $G + H$ the disjoint union of two graphs, the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. We denote the disjoint union of n identical graphs G as nG

Lemma 6.1. *We classify all finite region graphs on two curves.*

- $R_{0,0} = P_3$
- $R_{0,k} = 2K_k \times 2K_1$ *for* $k > 0$
- $R_{n+m,0} = K_{n+m} \times 2K_1$*, for* $n + m > 0$
- $R_{n+m,k} = (K_a + K_b) \times 2K_1$ *for* $n + m > 0$ *and* $k > 0$ *where* $a, b \in \mathbb{N}^+$ *and depends on the cyclic ordering of the incidences and transverse intersections.*

Proof. For the sake of illustration, we orientate our two curves anti-clockwise, as in Figure 6.2. The steps we follow below hold no matter the orientation of our two curves, in general replace "right" with "left" and vice versa for an individual curve if the orientation of that curve is flipped.

If $n = m = k = 0$, then we have two disjoint curves, and a quick check shows the corresponding region graph is the path on three vertices, P_3. Throughout the rest of the proof, we use Lemma 4.3 repeatedly to count the number of regions.

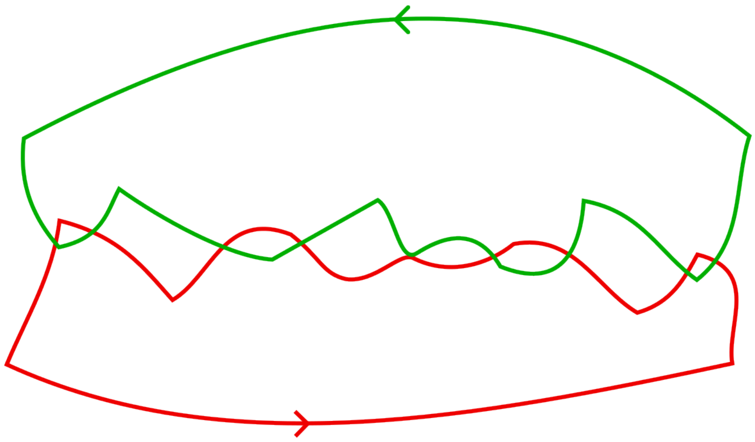

Figure 6.2: Anticlockwise orientation of curves, in this case $k = 4$, $n = 1$ and $m = 0$.

We now tackle the case for $k = 0$. There are exactly $n + m$ regions to the right of both curves. Upon removing either curve, these regions all merge, and hence form a clique K_{n+m} within the region graph. There are then two other regions, one to the left of γ and the right of λ and vice versa. Call these two regions G and L. Upon removing γ, G merges with those regions to the right of both curves, and so is adjacent to all of them in the region graph. Similarly, L merges with all these regions upon removal of λ. However, neither removal merges G and L. Furthermore, there are no regions to the left of both curves. Hence we see that the region graph is $K_{n+m} \times 2K_1$ as claimed.

If $n + m = 0 < k$ then by similar reasoning to the previous case all k regions to the right of both curves form a clique, and all k regions to the left of both curves form a clique. We find regions G and L as before, each to the right of one curve and the left of the other. We find that upon removal of γ, G merges with all regions to the right of both curves, and L merges with all regions to the left of both curves. The opposite is true upon removal of λ. Again G is not adjacent to L, and no region to the right of both curves can be merged with a region to the left of both. Hence we have two separate cliques of size k, and two vertices adjacent to all vertices in both cliques, but not each other. This is exactly $2K_k \times 2K_1$.

For the final case, we may enumerate the transverse intersection points within the natural cyclic ordering of the intersections as in Lemma 4.10. For each transverse intersection point i, we say the number of incident points and arcs occurring between i and $i+1$ in the cyclic ordering is n_i and m_i. We find there are exactly $r := \sum_{\text{even } j}(n_j+m_j+1)$ regions to one side of both curves and $s := \sum_{\text{odd } j}(n_j + m_j + 1)$ to the other side of both. By this we mean we have r regions either to the left of both curves or to the right of both and s regions on the opposite side of both curves. These form two cliques of size r and s which are disjoint. The exact values of r and l depend on the cyclic ordering of the incidences and intersections; for example, if all incidences are between two consecutive transverse intersection points, then we get two cliques of size $n+m+k$ and k. Finally, we again have a region G which is to the left of γ and the right of λ and vice versa for L. Again, both G and L may be merged with both cliques, but not with each other. Similarly, no region in either clique can be merged with a region of another clique. Hence we find the region graph is $(K_r + K_s) \times 2K_1$. □

In fact, the statement of Lemma 6.1 follows from the properties of codes described in §6.2.1. In the case where we have two curves, each code has two digits. Therefore possible codes are 00, 01, 10 and 11. Removing either curve then deletes the first or second indices. Deleting the first or second clearly merges all regions with code 00 and merges all regions with code 11, as their resulting code is 0 or 1. Removing the first curve then shortens the codes 01 and 10 to 1 and 0, and the opposite is true if we remove the second curve. Hence we have demonstrated that these two regions can merge with all regions of code 00 or 11. As regions whose codes differ at more than one digit cannot be adjacent, 10 and 01 are not adjacent and so we are done.

We now show a few properties of the graphs $R_{n+m,k}$. The first result on planarity is not surprising, as Lemma 6.1 demonstrates that large cliques may develop for $n+m+k$ small. Also note the following:

> A graph G is planar on the plane $\mathbb{R}^2$ if and only if it can be embedded on $\mathbb{S}^2$ such that no two edges cross.

This follows via applying a stereographic projection to a planar graph's embedding in $\mathbb{R}^2$. A stereographic projection is Möbius, as we showed in Chapter 2 and so maps non-crossing edges to non-crossing edges, and the meeting point of two edges (a vertex) to a meeting point. This justifies the use of the term planar for graphs embedded in $\mathbb{S}^2$.

Lemma 6.2. *$R_{n+m,k}$ is planar if and only if $n + m + k < 3$.*

Proof. A graph is planar if and only if it contains no K_5-minor or $K_{3,3}$-minor, by Kuratowski's Theorem. Let us first tackle the forwards direction.

If $k = 0$, then $K_{n+m} \times 2K_1$ only contains a K_5 minor if $n + m \geq 3$, and clearly can contain no $K_{3,3}$ minor if $n + m < 3$ as this provides too few vertices.

If $k \geq 1$ and $n+m = 0$ then by Lemma 6.1, we have two disjoint cliques K_k and K_k. To find a K_5 minor, contract all vertices of one K_k to the vertex representing G. This results in a clique of size $c := k + 2$, thus $c \geq 5$ implies $k + 2 \geq 5$ implies $n + m + k \geq 3$ as required.

If instead $k \geq 1$ and $n + m > 0$, then by Lemma 6.1 we have two disjoint cliques of size K_l and K_r, where $l + r = n + m + 2k$. Contract the clique K_x to the vertex G, where $x = \min\{r, l\}$, which results in a clique of size c, where $c = \max\{r, l\} + 2$. Observe that $c + x = r + l + 2 = n + m + 2k + 2$. Then $c \geq 5$ implies $n + m + 2k + 2 \geq 5$ implies $n + m + 2k \geq 3$. It only remains to observe that $n + m > 0$ and $k \geq 1$ together imply $n + m + 2k \geq 3$ and so all such graphs must contain a K_5 minor.

For the backwards implication, a manual check shows that $R_{0,0}$, $R_{0,1}$, $R_{0,2}$, $R_{1,0}$, $R_{1,1}$ and $R_{2,0}$ are all planar. $\square$

We now prove a perhaps slightly surprising result.

Lemma 6.3. *$R_{n+m,k}$ is perfect for all n, m, k.*

Proof. Note that cliques are perfect, and the disjoint union of two cliques is also perfect.

For the case $n = m = k = 0$, the region graph is P_3, a brief check shows it is perfect. We now assume $n + m > 0$ or $k > 0$. We therefore have a clique of size r, a clique of size l, and two vertices joined to these two cliques. Let us label vertices based on our choice of orientation. All of the following results hold for other labellings of vertices. We label vertices in the K_l clique as ll, those in the K_r clique as rr and then label G as lr and L as rl. These labellings are simply the side of γ and λ that each vertex's region is located on, as in the proof of Lemma 6.1. Our task is to show that $\chi(R_{n+m,k}) = \omega(R_{n+m,k})$, and that the same equality holds for all induced subgraphs I of $R_{n+m,k}$. We initially show the equality holds for $R_{n+m,k}$. We need r colours to colour the K_r, and l colours for the K_l. We then need two extra colours for G and L, as both are adjacent to all vertices in the l and r cliques. Hence $\chi(R_{n+m,k}) = \max\{r, l\} + 2$. Furthermore, $\omega(R_{n+m,k}) = \max\{r, l\} + 2$. This is because both K_r and K_l form cliques of size $r + 2$ and $l + 2$ with G and L. Suppose without loss of generality that $r > l$. Then adding another vertex to the $r + 2$ clique implies we must add a vertex from the K_l clique. But such a vertex is not adjacent to any vertex of the K_r clique, and so we have indeed found the clique number of $R_{n+m,k}$. We now consider cases for the labelled vertices of induced subgraphs I. Vertices can be labelled:

(i) Entirely rr or entirely ll - in which case I is perfect as a clique

(ii) Either lr or rl - in which case I is perfect as K_1 or $2K_1$

(iii) rr and ll - and so I is a disjoint union of cliques, and is perfect

(iv) Both rr and rl or lr - then I is a clique and is perfect

(v) Both ll and rl or lr - same as case (iv)

(vi) A combination of all three - so we choose a vertices labelled rr, b vertices labelled ll and $p \in \{1, 2\}$ vertices labelled lr or rl. Following the logic for $R_{n+m,k}$ above, $\chi(I) = \max\{a, b\} + p$ and $\omega(I) = \max\{a, b\} + p$. Hence I is perfect.

We have shown that $\chi(I) = \omega(I)$ for all induced subgraphs I of $R_{n+m,k}$, and so $R_{n+m,k}$ is perfect. □

It is true in fact that the join of two perfect graphs is perfect, and so Lemma 6.3 is not surprising, given the structure of such graphs. We chose not to rely on this property of the join to help demonstrate the structure of induced subgraphs of $R_{n+m,k}$.

6.3.2 On more than two curves

We now focus on results involving region graphs on more than two curves, $R_{\{\gamma_i\}}$. We begin by demonstrating several well-known graphs which are themselves region graphs.

Lemma 6.4. *A path on n vertices is a region graph.*

Proof. Take $n - 1$ nested, disjoint curves. □

In fact, we can show a stronger result.

Lemma 6.5. *Any finite tree is a region graph.*

Proof. Given a tree, we may assume it is rooted. We construct the tree iteratively. Begin with an empty set of curves, which has region graph K_1, which we shall assume corresponds to the root. For each child of the root, take a disjoint curve with no nesting. The resulting region graph is the root with its children. At the k^{th} step, for a vertex v added in the previous step, let $C(v)$ denote the number of children of v in the given tree. Add $C(v)$ disjoint non-nested curves to the region corresponding to v. Repeat this for each such vertex. The region graph resulting from this process is exactly the given tree. □

All paths are trees, and so this proof suffices to show Lemma 6.4. In fact, this tells us that there are at least as many planar region graphs on n curves as the number of unlabelled trees on n vertices [34].

Lemma 6.6. *Every cycle C_{2l} where $l \in \mathbb{N}^{\geq 0}$ is a region graph.*

Proof. Take l circles such that pairwise intersections are all equal to $\{x, y\}$. That is to say, all circles in the collection should intersect each other at the same two points, and nowhere else. See Figure 6.3 for an example for C_6. □

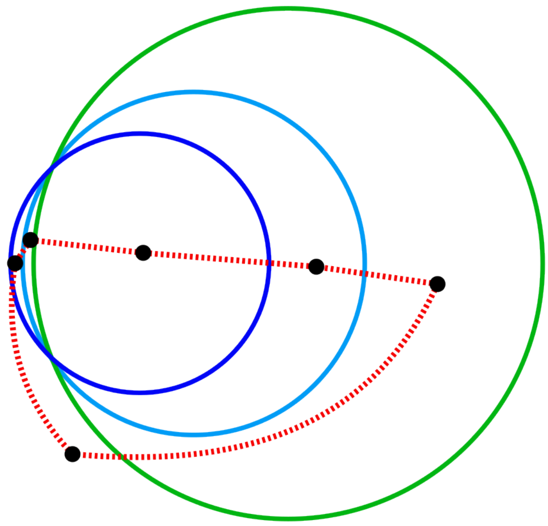

Figure 6.3: Arrangement of circles with region graph C_6

Lemma 6.7. *The only complete region graphs are K_1 and K_2.*

Proof. Firstly, $K_1 = R_\varnothing$. Secondly, $K_2 = R_{\{\gamma\}}$, the region graph of a single curve. Now suppose that $v \geq 3$ and K_v is a region graph on r curves, so in particular we have v regions. Then $r \leq v - 1$ as r curves divide $\mathbb{S}^2$ into at least $r + 1$ regions. We see that $r = 1$ is impossible, as the only region graph on a single curve is P_2. We see by Lemma 6.1 that $r = 2$ is also impossible. Now suppose $3 \leq r \leq v - 1$, and our curves are $\{\gamma_1, \gamma_2, \dots, \gamma_r\}$. Then the fact every region is adjacent to every other implies that the code for a given region can differ from every other region's code in at most one place. Suppose we take a region with code $x_1 x_2 \dots x_{r-1} x_r$ where $x_i \in \{0, 1\}$. Then without loss of generality, assume the code for every other region is of the form $y x_2 \dots x_{r-1} x_r$, where $x_1 \neq y \in \{0, 1\}$. But this implies every region is on the same side of every curve

$\gamma_2, \ldots, \gamma_r$. But this is impossible, every curve γ_i divides $\mathbb{S}^2$ into exactly two non-empty regions, and so we must have at least two regions which differ at a given index in their code. Hence we have a contradiction, and so K_v cannot be a region graph for a collection of $r \leq v-1$ curves, and so cannot be a region graph at all. Ergo, we have shown the only complete region graphs on Jordan curves embedded in $\mathbb{S}^2$ are K_1 and K_2. $\square$

Note that for the above proof it is not necessary to classify all region graphs $R_{n+m,k}$ on two curves, or even those for one curve. Every region has a code of 1 or 2 digits in either case, and it is guaranteed that there are at least two regions which differ in the i^{th} index by the Jordan curve theorem.

Lemma 6.7 implies that region graphs do not form a hereditary or minor-closed graph class. For example, $R_{3,0}$ contains K_3 as an induced subgraph and as a minor, but K_3 is not a region graph.

In the examples of §6.2.2 we used at least four curves to construct an arrangement containing vertices with the same code that are not adjacent in the corresponding region graph. Are there any ways to construct non-adjacent vertices with the same code for three or fewer curves? For three, the answer is yes as Figure 6.4 demonstrates. For two and below, the answer is no, as Lemma 6.8 shows.

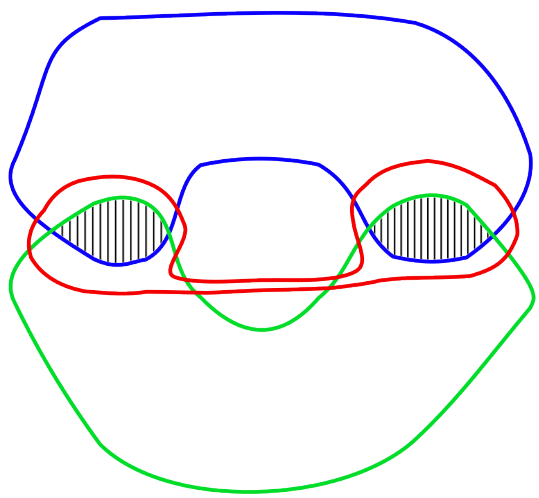

Figure 6.4: An arrangement of three curves such that the two hatched regions are non-adjacent and have the same code.

Lemma 6.8. *A region graph on r curves may contain non-adjacent vertices with the same code if and only if $r \geq 3$.*

Proof. We have already seen examples for $r = 3, 4$. To construct a region graph on $r > 4$ vertices with non-adjacent vertices with the same code, add $r - 3$ disjoint circles intersecting no other curves into a single region of the arrangement in Figure 6.4. This corresponds to adding a tree sharing the vertex of this region. This cannot connect the two non-adjacent regions with the same code, and so we have a graph on r curves with a pair of such regions.

We now tackle the forwards implication. If $r = 0$, then we have a single region and so the result holds trivially. If $r = 1$, then we have exactly two regions which necessarily differ in their codes. If $r = 2$, then we can have many regions with the same code, but as demonstrated in Lemma 6.1, all vertices with the same codes are contained in a clique and so there are no non-adjacent vertices with the same code. □

We end this section by demonstrating that a region graph on $\mathcal{J}$ need not be connected.

Lemma 6.9. *There are disconnected region graphs on $\mathcal{J}$.*

Proof. For a region graph on a class of curves to have multiple connected components, we must have some "double enclosing". By this, we mean that some curve in our collection must lie along at least one other curve at all points, see Figure 6.5a. This ensures that the boundary of some region is "doubled up", in the sense that every point of the boundary of the region is contained in at least two distinct curves. Thus removing any single curve cannot merge the region with any other. In a similar vein we may have some arbitrary finite collection of curves contained entirely within a region which is double enclosed. This then demonstrates we may have non-trivial connected components in such a region graph.

Let us now show that double enclosure is necessary for disconnectedness of a region graph. Let us assume that our region graph has an isolated vertex. Suppose this region is not double enclosed. Then there must be some portion of the boundary of the region that lies on only a single curve. Removing this curve necessarily merges the region with another, contradicting our assumption that the region is isolated. Thus "double enclosure" is necessary and sufficient for a region to be isolated in a region graph.

Furthermore, as we assume our region graph is on a finite collection of curves, we must have some arc-incidence along the boundary; we cannot double up an entire boundary just via point-incidence using only finitely many curves with finitely many intersections. □

As a corollary of Lemma 6.9, we can determine that all finite forests are indeed region graphs.

6.4 Region graphs on $\mathcal{C}^2$

As seen in the previous section, region graphs on general curves are in general extremely complex, and the task of characterising them seems intractable. However, if we instead focus on circles on $\mathbb{S}^2$, there is much we can say about region graphs. To begin with, we must justify what is meant by a region graph on $\mathcal{C}^2$, whose vertices are chords rather than curves. We recall that $\mathcal{C}^2$ is isomorphic to the intersection graph of circles embedded in $\mathbb{S}^2$. Every chord defines a unique embedded circle and vice versa, and hence the intersection graph of circles embedded on $\mathbb{S}^2$ has all of the properties proven in Chapter 2. In this way we justify our terminology around region graphs on $\mathcal{C}^2$. Two circles which are tangent shall be referred to as *incident* as in Chapter 2 and two circles intersecting at two distinct points *overlap*. A finite collection of circles corresponds to a finite induced subgraph of $\mathcal{C}^2$.

We begin with a simple lemma, which demonstrates a crucial difference between region graphs on $\mathcal{J}$ and those on $\mathcal{C}^2$.

Lemma 6.10. *Every region graph on $\mathcal{C}^2$ is connected.*

Proof. By Lemma 6.9, we see that to have double enclosure of a region, we must have arc-incidence along its boundary. Therefore, if our curves are all circles, then we must have arc-incident circles in our collection. But circles that agree along an arc must be the same, contradicting our assumption that our region graph is on distinct curves. Hence such an arrangement is impossible on $\mathcal{C}^2$, and so all such region graphs must be connected. □

Remark We can put bounds on the number of vertices of a region graph of a collection of r circles. In particular, a lower bound is $r+1$ which can be achieved if all r circles are disjoint. The upper bound is $r^2 - r + 2$ which may be achieved if the r circles pairwise overlap and there are no multiple intersection points.

This implies C_5 is not a region graph on any number of circles by the following. For $r < 3$ the upper bound on the number of vertices is less than five. If $r > 4$ the lower

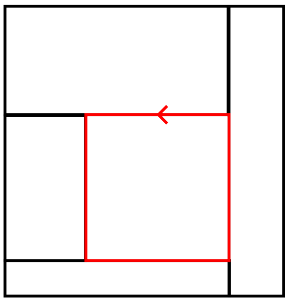

(a) Four black rectangular curves, and an orientated red curve which lies entirely along the boundaries of the other four. The region to the left of the red curve cannot be merged with any other region by removing a single curve, and so is isolated in the corresponding region graph.

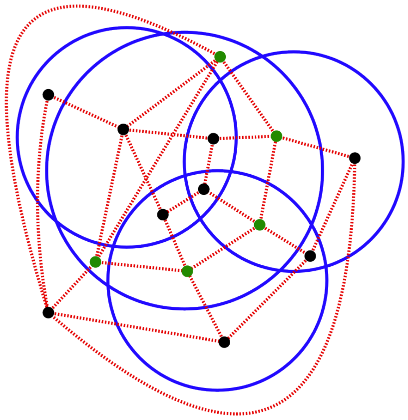

(b) Example of a non-perfect region graph; the green vertices induce a C_5, so by the Strong Perfect Graph theorem, this region graph is not perfect.

Figure 6.5

bound on the number of vertices is greater than five. We will see in our classification of region graphs on three circles in §6.4.1 that C_5 is not a possible region graph - see Figure 6.6. The only way for four circles to produce five regions is if all existing intersections are incidences. The only region graphs of such arrangements are the trees with five vertices, P_5, *chair* and $K_{1,4}$.

Remark Lemma 6.3 begs the question of whether all region graphs are perfect. This is not true, consider Figure 6.5b formed of circles in $\mathcal{C}^2$.

6.4.1 Region graphs of small collections of circles

To classify region graphs on circles, the most obvious method is to find all region graphs on collections of r circles, with any pattern of overlaps. This is possible for small r, but quickly becomes intractable.

Let us list the region graphs of up to three circles.

Zero circles: The only region graph on the empty collection is the single vertex graph, K_1.

One circle: The only region graph on a single circle is K_2.

Two circles: There are exactly two region graphs on two circles: the path P_3, and the 4-cycle C_4.

Three circles: Up to isomorphism, there are eleven region graphs on three circles. These are shown in Figure 6.6.

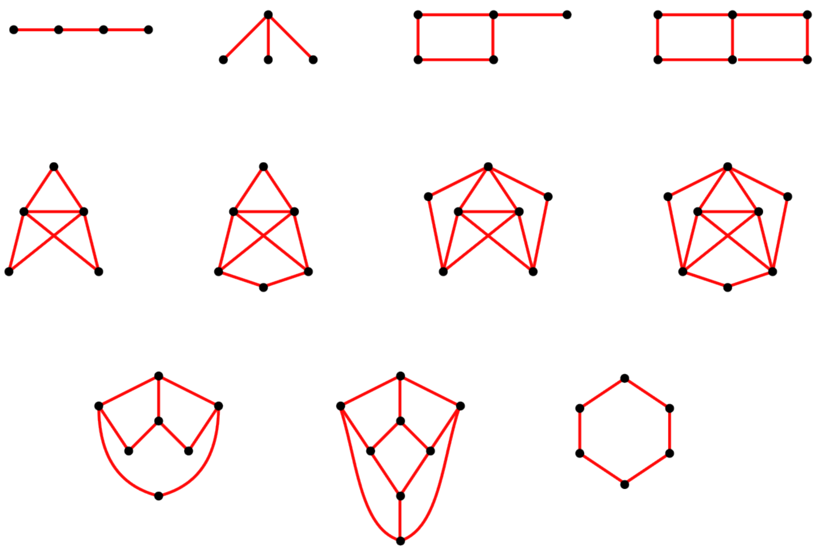

Figure 6.6: Region graphs on three circles.

In order to classify the region graphs on r circles, it is necessary to enumerate all possible arrangements of r circles on $\mathbb{S}^2$ up to Möbius transformation. That is to say, two arrangements are considered the same if there is a Möbius transformation which sends one to the other. This is not a trivial task, and so we would seek to lessen the number of configurations that must be calculated; this will be the focus of the next section. There is one arrangement for a single circle, three for a pair, and nineteen for three circles. See Figure A.1 for an illustration of those arrangements of three circles.

Note that we must allow tangency (incidence) and multiple intersection points in our configurations. It is then a simple computational task to find all non-isomorphic region graphs on such arrangements. In fact, the study of region graphs on circles lends itself to computational methods. Several of the results in this chapter were inspired by the analysis of region graphs of small (< 15) randomly generated collections of circles. We have provided the code used to generate these collections, as well as to calculate the corresponding region graph, in Appendix C. Most analysis of the graphs was done via SageMath.

Given that there are nineteen arrangements of three circles on the sphere, and only eleven region graphs, this implies that region graphs do not encode all topological information of a corresponding arrangement of circles.

6.4.2 Region graphs on collections with few pairwise intersections

In §6.4.1 it is clear that for $r > 3$ circles, the number of arrangements that must be analysed quickly grows. We can instead focus on finite collections of arbitrary size, where individual circles cannot intersect more than some number of other circles I. This approach lessens the number of arrangements that must be calculated. We say the *degree* of a circle in a collection is the number of other circles it intersects. Recall that we distinguish two arrangements only up to Möbius transformation.

Our aim in this section is to classify the region graphs on collections of circles with maximum degree at most three. We also demonstrate an inductive method to construct a region graph on a collection with maximum degree k.

Let us first define some terminology. A *block* in a graph G is a maximal biconnected subgraph of G. A biconnected graph is a connected graph such that deleting any single vertex does not disconnect it. A *chain* of circles is a finite sequence of circles such that consecutive circles may intersect, either incidentally or by overlapping, and no other pairwise intersections occur. The length of a chain is the number of circles it contains. An *incident* chain is a chain in which the only intersections are incident, and an *overlapping* chain is a chain in which all intersections are overlaps. See Figure 6.7 for some examples. A *loop* is a chain in which the first and last circles may also intersect; the same terminology used for chains also applies to loops. We also insist that loops are of length four or greater.

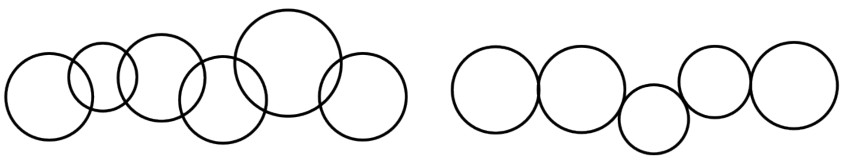

Figure 6.7: An overlapping chain of length six (left) and an incident chain of length five (right). A chain could also contain both overlaps and incidences.

We denote the class of region graphs for finite collections of circles on $\mathbb{S}^2$ where a given circle can intersect at most I others $Circle_I$. It is then clear that for $I = k$, $Circle_0 \subset Circle_1 \subset \dots Circle_{k-1} \subset Circle_k \subset Circle_{k+1} \subset \dots$.

Collections of circles with $I = 0$

The induced subgraph of $\mathcal{C}^2$ of such a collection of circles is an independent set. The only region graphs on r disjoint (as $I = 0$) circles are the trees on r vertices, which can be seen as follows. We have only a single operation for modifying a given collection of disjoint circles. This operation is to add another disjoint circle. Such a circle may contain some portion of the given collection in both sides. Within the region graph, this corresponds to splitting the articulation vertex connecting the collections on either side in two, such that each split vertex is the root of the region graphs of the collections contained in either side. The two vertices should be adjacent to one another. Repeating this iteratively we can obtain any tree. Hence we see that $Circle_0$ is exactly the class of trees.

We also note that the trees are those connected graphs in which all blocks are single edges.

Collections of circles with $I = 1$

In this case, given a circle, it can either be disjoint from all others, be incident to a single other circle, or overlap a single other circle. The induced subgraph of $\mathcal{C}^2$ of such an arrangement is a graph with maximum degree one. Let us start with a collection of disjoint circles. The region graph of such an arrangement is a tree, as shown in the previous section. We iteratively modify the starting collection where the possible ways to add a circle are:

(i) Adding a disjoint circle to a region. This is an operation already available in the case $I = 0$. See Step 6 of Figure 6.8 for an example.

(ii) Adding a circle to a region, incident to a single circle on the boundary of this region. Within the region graph, this turns a K_2 into a P_3, which is equivalent to adding a pendant vertex. So in fact this is equivalent to operation (i) in its modification of the region graph. We could achieve the same change to the region graph by adding a disjoint circle to the same region with the same circles in each hemisphere. See Step 3 of Figure 6.8 for an example.

(iii) Adding a circle that overlaps a single existing member of the collection. Such a circle can only intersect two regions, the regions bounded by the circle it overlaps. Within the region graph, this corresponds to expanding an edge xy into a 4-cycle by adding two new vertices u, v and the edges ux, uv, vy. See Steps 5, 7 and 9 of Figure 6.8 for examples.

Each of these is what we will call an *addition operation* or *operation* for short of $Circle_1$.

In general, for $Circle_k$, all operations in classes $Circle_l$ with $l < k$ are available, but we also have new operations, allowing us to create circles of degree k, which are not available in lower classes; we will call such operations *semi-saturating.* A semi-saturating operation takes a collection which has maximum degree at most $k-1$, and adds a new circle in such a way that either the added circle or one of its neighbours is of degree k. We term a semi-saturating operation such that the added circle and all of its neighbours are of degree k *saturating.*

These may be interpreted as modifications of the intersection graph of a collection of circles. It must be kept in mind that the intersection graph does not encode all topological information about an arrangement of circles; so to recover the region graph we must keep track of the collection of circles itself. Within the intersection graph, a semi-saturating operation of order k consists of adding a new vertex v and at most k edges incident to v, such that at least one vertex incident to a newly added edge is of degree k. A saturating operation is then the same operation, but with the stronger requirement that all vertices incident to the added edges are of degree k.

Before characterising the graphs in $Circle_1$ and determining the structure of graphs in $Circle_2$, let us first demonstrate the utility of semi-saturating operations.

6.4.3 An algorithm for recursively constructing a region graph in $Circle_k$

Given our definitions for semi-saturating and saturating operations, it is natural to ask if all graphs in $Circle_k$ can be obtained from a graph in $Circle_{k-1}$ using just semi-saturating operations or saturating operations. We first show this result is true of semi-saturating operations.

Lemma 6.11. *Any $G \in Circle_k$ can be obtained from a graph $G' \in Circle_{k-1}$ through only semi-saturating operations.*

Proof. Suppose we have a graph $G \in Circle_k$ and a corresponding arrangement of circles. We work in the intersection graph, H, of the collection of circles by a recursive series of vertex deletions.

Firstly, note that $\Delta(H) \leq k$ as a result of the corresponding collection of circles. At a given step delete a vertex of degree k, which corresponds to removing a circle of degree k from the starting collection. Create a list L of these deletions, appending each removal to the end of this list. Such a deletion reduces the degree of any neighbouring vertices by one. Continue to do so until there are no vertices of degree k left, so all remaining vertices are of degree at most $k-1$. This corresponds to a collection of circles of maximum degree $k-1$, and so the associated region graph G' is a member of $Circle_{k-1}$. Also note that this process will terminate at a finite step, as our initial collection is finite by assumption.

Now we reverse this process, by reading L right to left and adding in the required vertices and edges step by step. Each such addition is a semi-saturating operation by definition, as we are adding in a circle of degree k. We do indeed recover G in its entirety through these additions. So this demonstrates that G may be recovered from G' by a series of semi-saturating operations of $Circle_k$. □

The above lemma demonstrates that the only "meaningful" operations of $Circle_k$ are its semi-saturating operations. When applied to $Circle_1$, this means we can start with a tree, and add only incident or overlapping circles to obtain any member.

As may be expected, we can determine an even stronger result from Lemma 6.11.

Corollary 6.1. *Given $G \in Circle_k$, we may obtain G by starting from an empty collection of circles, and applying semi-saturating operations in $Circle_0$, then $Circle_1$, ... and then in $Circle_k$ in that order.*

Proof. We follow a proof by induction. For $k=0$, we have already shown that every tree can be obtained from the single vertex graph by the single operation - which is semi-saturating - of $Circle_0$. Now assume the inductive hypothesis holds for all graphs in $Circle_{k-1}$. Then take a graph $G \in Circle_k$, and a corresponding arrangement of circles. Follow the proof of Lemma 6.11 to remove circles of degree k. Then the maximum degree of any circle in the resulting arrangement is $k-1$, and so the corresponding region graph H is in $Circle_{k-1}$. Then add the removed circles back in to obtain G from H by only semi-saturating operations of $Circle_k$. Thus the statement of the corollary holds. □

The same is not true of saturating operations. A chain of length greater than five, for example, necessarily requires semi-saturating but not saturating operations of $Circle_2$ to construct.

Let us make some observations that are a consequence of Lemma 6.11 and Corollary 6.1.

Observation 6.1 Any sequence of operations (whether non-saturating, semi- or saturating) that results in the same collection of circles, up to Möbius transformation, will give us the same region graph. This is because every collection of circles has a unique region graph - although each region graph does not necessarily have a unique collection of circles, as Figure 6.6 demonstrates. Furthermore, given a particular sequence of operations, we may permute the $Circle_0$ operations, then permute the $Circle_1$ operations and so on. At no point can we create a circle with degree greater than l by performing the $Circle_l$ operations in a different order.

Observation 6.2 To get $G \in Circle_k$ we can start from the single vertex graph. As we perform operations in $Circle_l$, $l < k$, we may label the added edges. In particular, those edges corresponding to adjacency to a region at least partially bounded by a circle of degree l should be labelled. Then labelled edges cannot again be modified, until we start applying operations from $Circle_{l+1}$; this would require adding a circle exceeding degree l. As our operations are semi-saturating, we may not need to label all edges created in such an operation. If instead an operation is saturating, then all new edges must be labelled.

Then assuming semi-saturating operations for $l < k$ are known, if we are able to find all such operations of $Circle_k$, we may construct all associated region graphs. The task of finding all such operations is non-trivial. One must determine all of the distinct arrangements of circles up to Möbius transformation of maximum degree k that achieve this maximum. This is, however, a significant reduction in the number of arrangements that must be determined than in §6.4.1.

Due to Lemma 6.11, we see that graphs in $Circle_1$ have the following characterisation.

> A graph G is in $Circle_1$ if and only if it is connected and all cycles in G are of length four, and any pair of cycles shares at most one vertex.

This is exactly the class of cactus graphs containing only 4-cycles and 2-cycles [35]. Equivalently, these are the connected graphs whose blocks are either an edge or C_4. All such graphs are planar and perfect.

Based on Observation 6.2, we can colour the edges of an added 4-cycle, and colour the edges of a P_3 resulting from adding an incident circle. Adding a disjoint circle corresponds to adding an uncoloured edge. Coloured edges cannot then be enlarged to 4-cycles. We can clearly add an uncoloured edge adjacent to any given vertex, by inserting a disjoint circle in the corresponding region. See Figure 6.8 for an example of building up a graph in $Circle_1$.

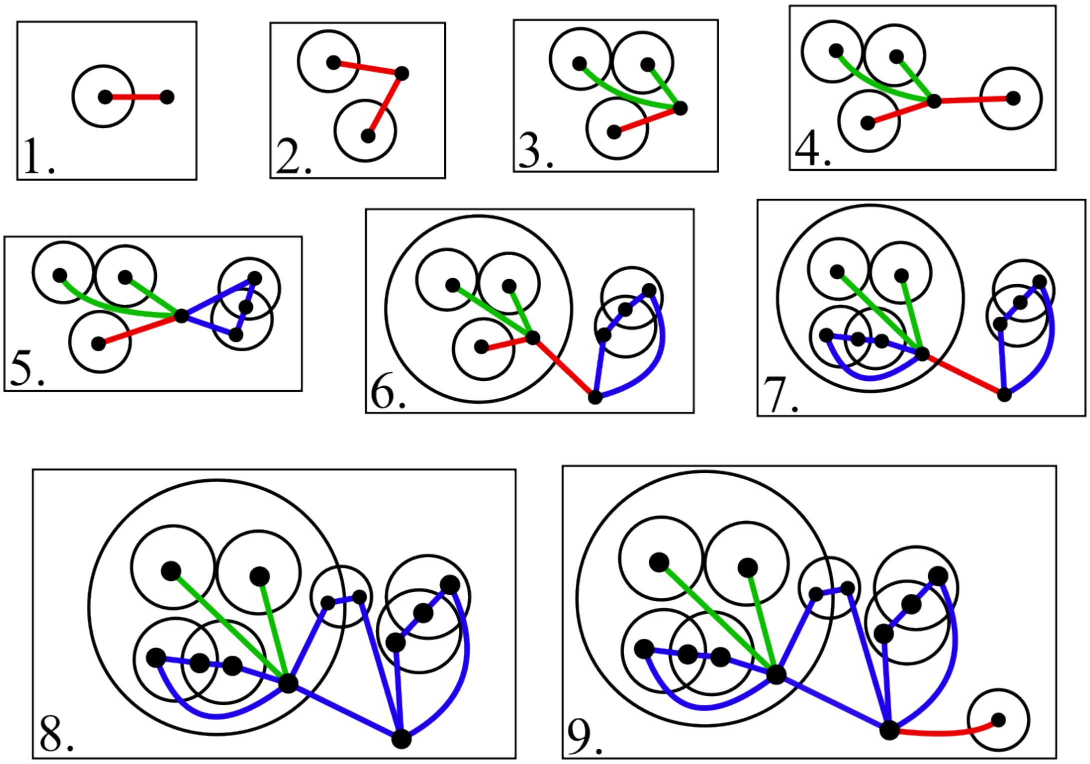

Figure 6.8: In this example we build a graph in $Circle_1$ from a single vertex via operations (i), (ii) and (iii). We colour edges of 4-cycles **blue** and the edges corresponding to an incident pair **green**. Edges corresponding to uncoloured edges are **red**. It is clear from Lemma 6.11 that we could reach step 9 from a different starting tree in only 4 operations.

Collections of circles with $I = 2$

By analysing all semi-saturating operations of $Circle_2$, we can determine the structure of all graphs in this class. We list all such operations in Appendix B as there are a large number of possible operations. Before characterising the graphs of $Circle_2$, let us first define $Loop \subsetneq Circle_2$, the class containing all region graphs of loops.

Definition 6.1 A graph $G \in Loop$ is a graph that may constructed in the following manner:

1. Begin by taking a disjoint union of isolated vertices and paths of odd length, $P_{2k+1} = x_1x_2 \ldots x_{2k+1}$.
2. Add a vertex adjacent to all isolated vertices and alternating vertices $x_1, x_3, \ldots x_{2k+1}$ of each path. The resulting graph is the region graph of a chain of circles, see Figure

6.9. Let us call this added vertex v.

3. Add a new vertex v' such that $N(v') = N(v) \cup \{v\}$.

The result is the region graph of a loop. The odd length paths result in 4-cycles which represent the overlapping circles, while the single vertices represent the incident circles. See Figure 6.9 for an example.

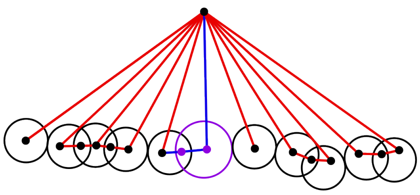

(a) A region graph of a chain. The addition of the purple circle is a saturating operation.

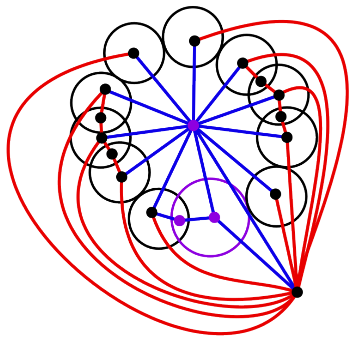

(b) A region graph of a loop. The addition of the purple circle is a saturating operation.

Figure 6.9: In both cases, we may start with the black circles and add the purple circle. New edges and vertices created by this operation are blue and purple respectively.

Also note that $Loop \subset Circle_2$, but $Loop \cap Circle_1 = \varnothing$ and that all graphs in $Loop$ are biconnected. Furthermore, within the region graph of a chain, we get sequences of 4-cycles, with each pair of consecutive 4-cycles sharing a single edge. These correspond to overlapping sequences of circles within the chain. Such a graph is biconnected, and may be expressed as an odd length path with a single vertex adjacent to every other vertex of the path, including the end vertices. Given this, we can now characterise those graphs in $Circle_2$.

Lemma 6.12. *A graph G is in $Circle_2$ if and only if it is connected and its blocks are:*

(i) *K_2 or C_4*

(ii) *One of the graphs listed in Figure 6.6 - excluding P_4, $K_{1,3}$ and 4-pan as these are not biconnected.*

(iii) *An odd length path $x_1 x_2 \dots x_{2k+1}$ with a vertex adjacent to $x_1, x_3 \dots x_{2k+1}$.*

(iv) A graph in Loop.

Proof. The forward direction follows from the list of semi-saturating operations of $Circle_2$ in Appendix B, and that all graphs in (i), (ii), (iii) and (iv) are biconnected. No region graph of a non-overlapping chain is biconnected; we may remove the vertex corresponding to Step 2 of Definition 6.1.

For the backward direction suppose we have a graph with such blocks. Then every such block corresponds exactly to an arrangement of circles with maximum degree two. Where two blocks share a vertex, one arrangement of circles must be contained in the region corresponding to the shared vertex. By these deductions we may then construct an arrangement of circles with region graph exactly the graph we started with. □

Chapter 7

Concluding remarks

We began this dissertation with a short overview of the higher-dimensional analogues of Möbius transformations, and demonstrated several useful properties of such maps. We then proceeded to show their relationship to the d-sphere graph in answering Question 1. Subsequently, we were able to recover the entire automorphism group of $\mathbb{S}^2$ by generalising the allowable curves in our intersection graph, and investigated those homeomorphisms of $\mathbb{S}^2$ that preserve C^k Jordan curves. We ended by analysing the properties of region graphs and explored several more tractable classes of region graphs of circular curves. In this conclusion, we will discuss further work that may be done in these areas.

To begin with, the topic of Möbius transformations is a well-studied area. In this way, what we have demonstrated in the first chapter merely skims the surface. For example, [8] and [9] both go into far more depth on the structure of the d-dimensional Möbius group and its links to hyperbolic geometry. There is also much to be said on the links between Möbius transformations and conformal mappings - those functions that preserve angles, which [9] focuses on. We also made a brief mention of Clifford algebras, and would refer the reader to several papers on the links between Möbius transformations and Clifford algebras [36][37][38].

In Chapter 3, we were able to answer Question 1. However, there are plenty of other open questions concerning sphere graphs.

As an example, in [1], Georgakopoulos was able to demonstrate that $\mathcal{C}_\mathbb{Q}$ is a strongly universal element for the class of countable circle graphs. In this case, a circle graph is an induced subgraph of $\mathcal{C}$ and more generally, a d-sphere graph is an induced subgraph of $\mathcal{C}^d$. The graph $\mathcal{C}_\mathbb{Q}$ is the intersection graph of those chords whose endpoints are rational, where we identify $\mathbb{S}^1$ with $[0, 1]$. A graph U is strongly universal for a graph class $\mathcal{D}$ if every $G \in \mathcal{D}$ is an induced subgraph of U. It is unknown whether there is a strongly universal element for the class of countable d-sphere graphs for every $d > 1$.

Georgakopoulos has also shown that $\mathcal{C}_{\mathbb{Q}}$ is invariant under *local complementation* - swapping the edges and non-edges in $N(v)$ for some vertex v. An open question is whether there are any non-empty, connected, countably infinite and vertex-transitive graphs that are invariant under local complementation, which are not isomorphic to the Rado graph or $\mathcal{C}_{\mathbb{Q}}$. In particular, would a strongly universal element of the class of d-sphere graphs for some $d > 1$ satisfy this property, thereby providing infinitely many examples?

One could also determine whether there are any finite d-sphere graphs which are invariant under local complementation.

Within Chapter 4 we were able to find a natural graph $\mathcal{J}$ such that $\mathrm{Aut}(\mathcal{J})$ is isomorphic to $\mathrm{Aut}(\mathbb{S}^2)$. It seems likely a method of proof similar to ours should generalise to higher dimensions. In particular, one could consider the intersection graph $\mathcal{J}^d$ of the continuous, injective maps $\gamma : \mathbb{S}^{d-1} \to \mathbb{S}^d$, in which case $\mathcal{J} = \mathcal{J}^1$. More care would need to be taken around wild embeddings, such as the Alexander horned sphere.

It would then be another natural generalisation to consider the smooth Jordan curves in higher dimensions. One could also analyse other classes of Jordan curves. For example, in Lemma 6.10, we relied on the fact that two circular curves on $\mathbb{S}^2$ that agree on an arc must be the same. One could then consider the real-analytic Jordan curves, in which the same property holds. Rectifiable Jordan curves, which are those curves of finite length, may also provide a fruitful path of exploration.

Within Chapter 6 we hope to have demonstrated some interesting results that arise from this seemingly novel method of graph construction. There is much more work that may be done on these graphs. For example, we were able to classify $Circle_I$ for $I = 0, 1, 2$ by admissible blocks. Is it true that every such class $Circle_I$ admits such a characterisation via admissible blocks? Furthermore, one may ask whether it is truly necessary to calculate all suitable arrangements of circles up to Möbius transformation to find these region graphs, or whether there exists some more tractable combinatorial method to determine them. It seems likely that the case for $I = 3$ and perhaps for $I = 4$ could be solved through the first method computationally, but beyond this we expect a different approach would be needed.

We also noted that region graphs do not encode all topological information of a collection of curves in general. There are several simple examples of collections of circles with the same intersection graph and the same region graph, but which are topologically distinct. Consider Figure 7.1 for example.

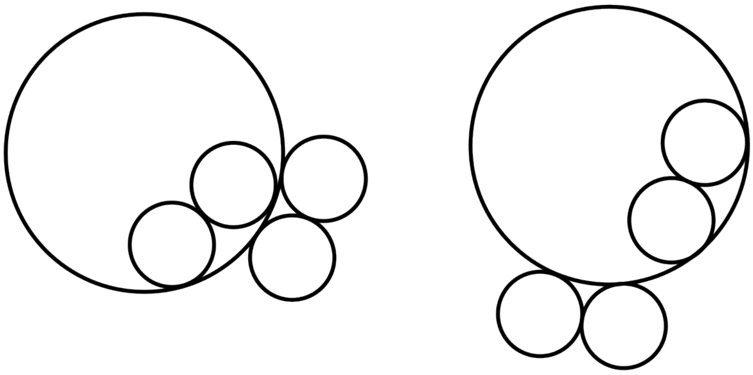

Figure 7.1: Topologically distinct arrangements of circles with the same intersection and region graphs.

In both arrangements, the cyclic order of incidences around the large circle differ. It is immediately clear that the intersection graphs are the same, and one can quickly verify that the region graphs agree. This demonstrates that even the pair of the intersection graph and the region graph does not encode the entire topology in general. Could one find an invariant such as the region graph, that encodes all topological information about a given set of curves? One potential candidate is a multigraph equivalent of the region graph. Again, vertices correspond to topological components. Between each pair of vertices, add edges each corresponding to a distinct sequence of curve removals that merge the two regions. One could reverse the question, and ask if there is much to be said about those arrangements whose topology region graphs do determine.

Appendix A

Arrangements of circles

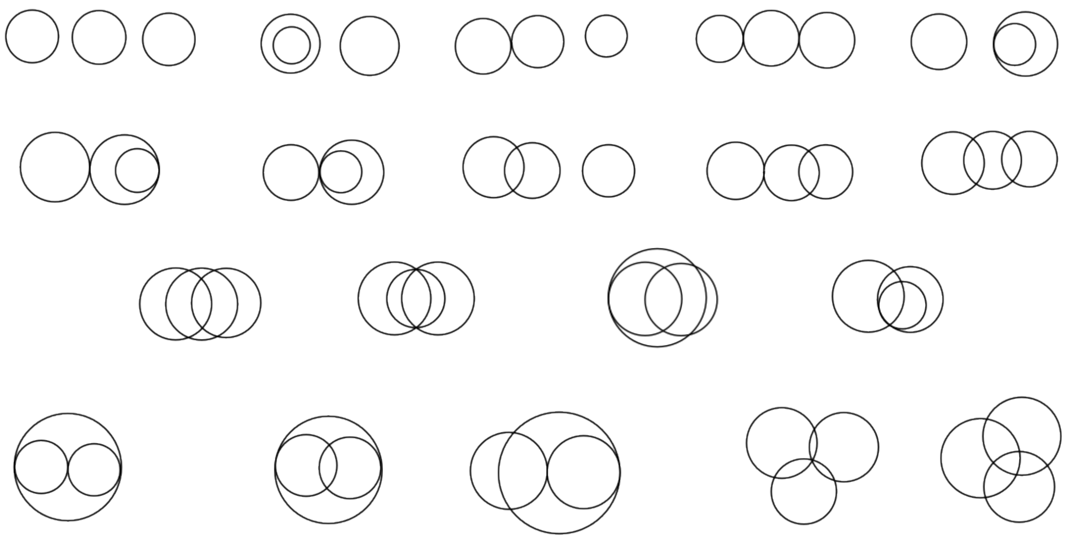

Figure A.1: The 19 arrangements of three circles on the 2-sphere up to Möbius transformation

Appendix B

Semi-saturating operations of $Circle_2$

We list all semi-saturating operations of $Circle_2$. Within our demonstrative diagrams, we consider an existing pair of circles, in black, and add a third purple circle. We then show the change to the corresponding region graph, where any new edges are blue and new vertices are purple. Figures B.1-B.4 demonstrate operations on a disjoint pair of circles, which are all semi-saturating, except for two cases outlined in Figures 6.9a and 6.9b. Figures B.6-B.7 demonstrate operations on an incident pair. Finally, Figures B.8-B.15 demonstrate operations on an overlapping pair. Note that in each case, there may be another way of obtaining the same collection via a different sequence of operations. As we remarked, this does not affect the resulting modification to the region graph.

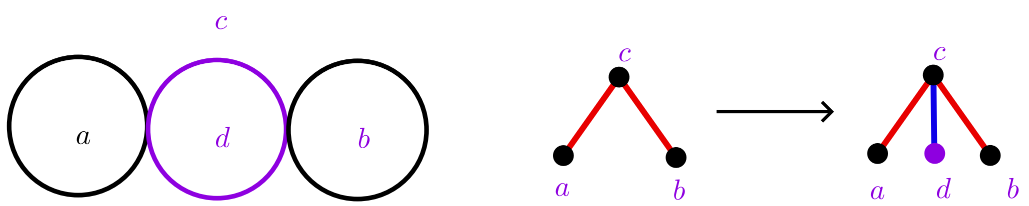

Figure B.1: Operation I: Add a circle incident to both, such that both original circles are on the same side. This adds a pendant vertex to the region bounded by both circles, turning a P_3 into a $K_{1,3}$ by adding a single vertex and a single edge. This is equivalent to adding a single circle incident to an existing circle of degree one.

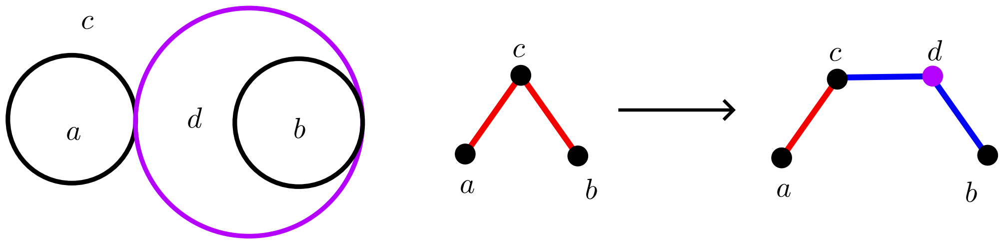

Figure B.2: Operation II: Add a circle incident to both such that each is on a different side. This turns a P_3 into a P_4 by splitting a vertex and adding an edge. This is equivalent to adding a circle incident to an existing circle of degree one, on the opposite side to its existing neighbour.

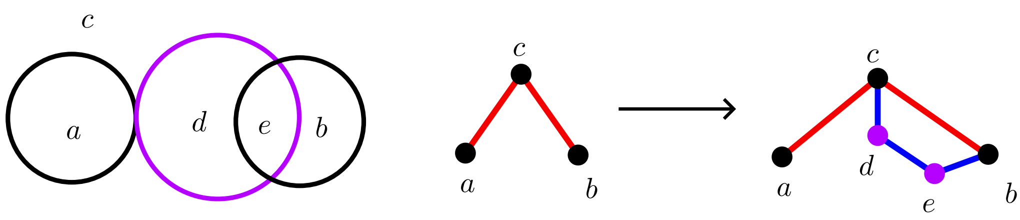

Figure B.3: Operation III: Add a circle incident to one and overlapping the other. This adds a 4-cycle to a single edge of the P_3, by adding two vertices and three edges.

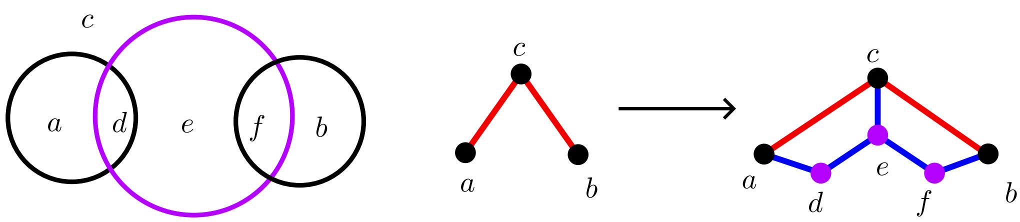

Figure B.4: Operation IV: Add a circle overlapping both. This turns a P_3 into two 4-cycles sharing an edge, by adding three vertices and five edges.

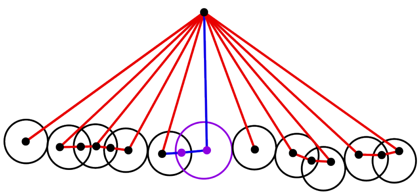

(a) An example of operation (iii) where each disjoint circle is of degree one. This diagram also demonstrates what the region graph of a chain looks like, a single vertex with runs of 4-cycles and pendant vertices.

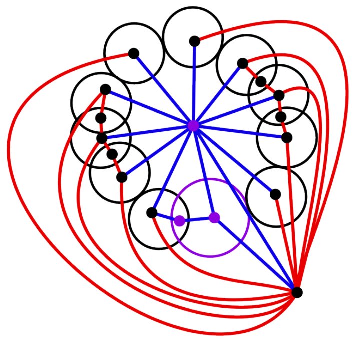

(b) Joining up the ends of a chain to form a loop via operation (iii). Such an operation takes a region graph as in Figure (a) and adds a new vertex corresponding to a new entrapped region. Such a vertex mirrors the upper vertex in Figure (a) and is adjacent to it; removing any single circle from the loop merges the entrapped and outer regions.

Figure B.5

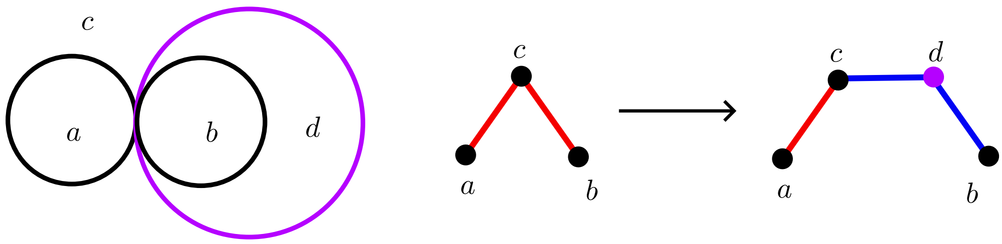

Figure B.6: Operation V: Add a circle incident to both such that there is a triple point. This operation turns a P_3 into a P_4, and so is analogous to operation (ii).

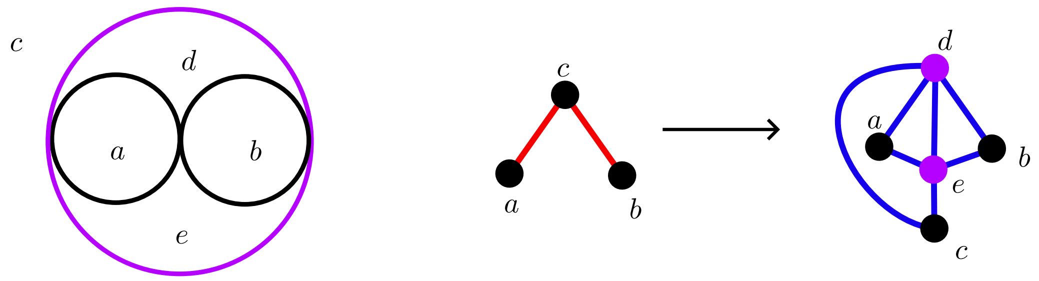

Figure B.7: Operation VI: Add a circle incident to both such that there is no triple point. Such an operation takes a P_3 and adds edges and vertices such that we get three 3-cycles sharing an edge. This is exactly the operation in Figure 6.9b, where we join the ends of an incident chain of length two.

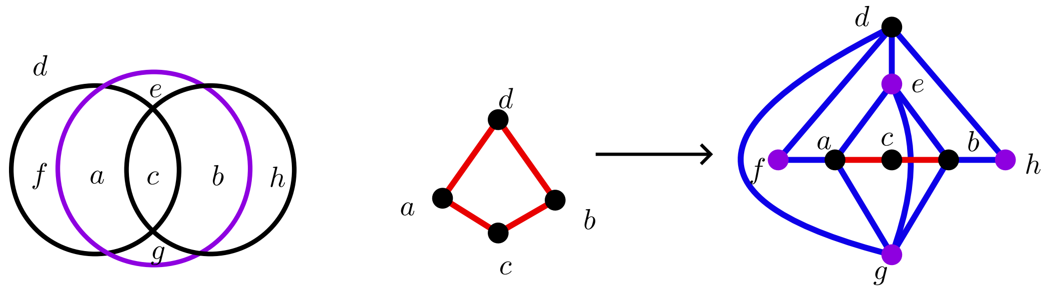

Figure B.8: Operation VII: Add a third which overlaps both and that creates two regions with the same code.

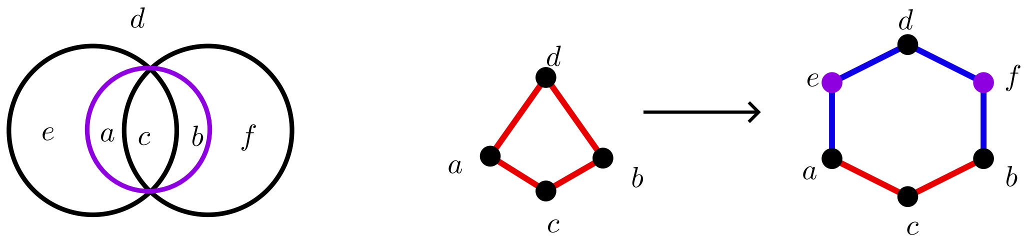

Figure B.9: Operation VIII: A third which overlaps both such that we obtain two triple points.

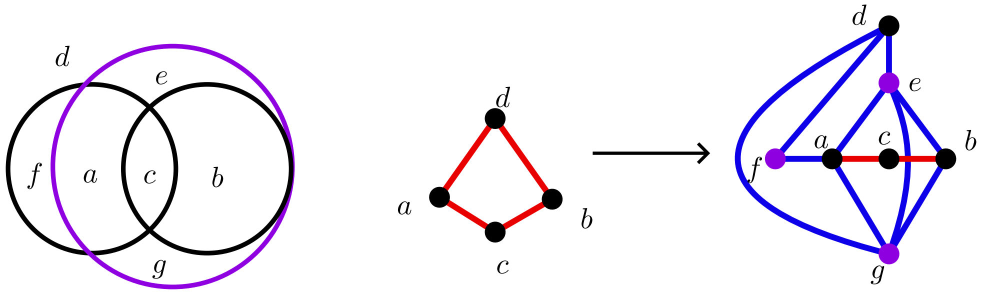

Figure B.10: Operation IX: Add a third overlapping one and incident to the other with no triple points.

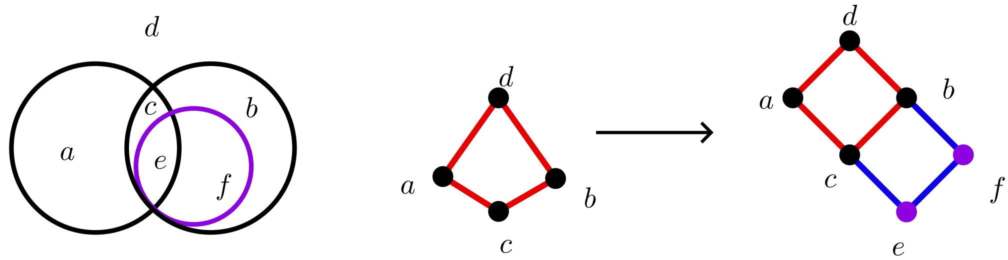

Figure B.11: Operation X: The above, but with a triple point.

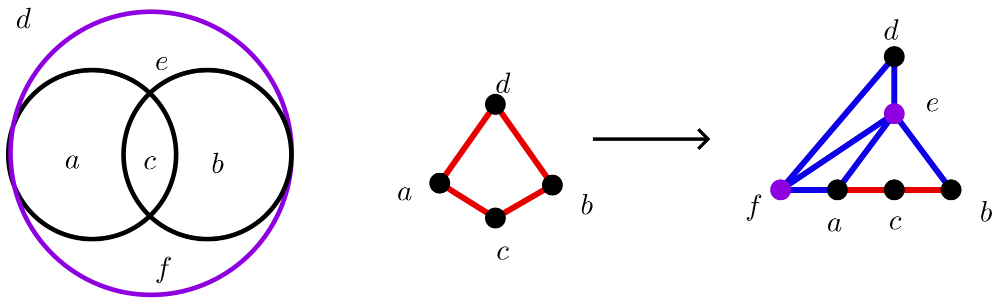

Figure B.12: Operation XI: Add a third incident to both, with both on the same side of the new circle.

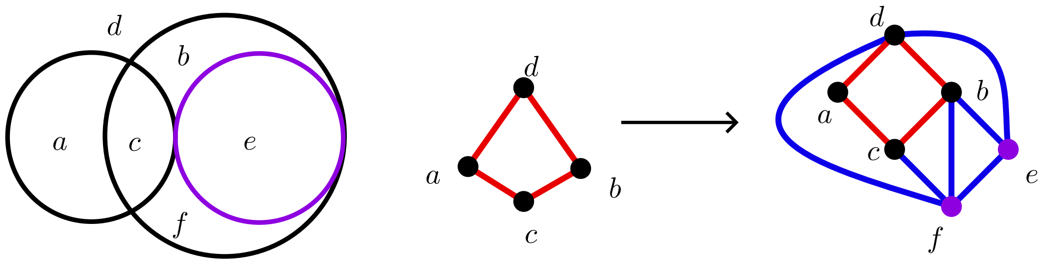

Figure B.13: Operation XII: Add a third incident to both, with the new circle being on opposite sides of the overlapping pair.

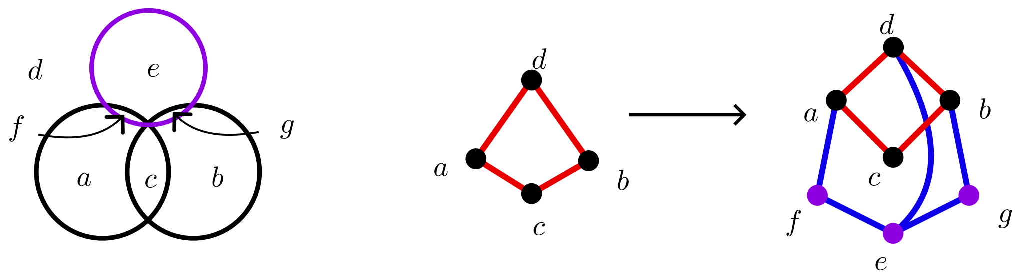

Figure B.14: Operation XIII: Add a third overlapping both such that we obtain a single triple point.

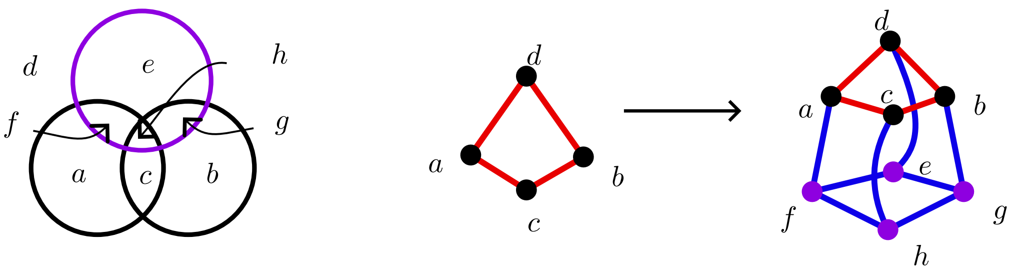

Figure B.15: Operation XIV: Adding a circle such that we get eight regions and no triple points. Observe that this turns a 4-cycle into a cube graph.

Appendix C

Code for generating region graphs on circles

```python
import numpy as np
from numpy import random
import matplotlib.pyplot as plt
import networkx as nx
from shapely.geometry import Point
from shapely.ops import unary_union, polygonize

def generate_circle(centre, radius, num_points=500):
    """Returns a shapely geometry for a circle of given centre and radius."""
    return Point(centre).buffer(radius, resolution=num_points), centre, radius

def generate_circles(n):
    """Generates n centres and radii to be converted to shapely geometries."""
    centres = 1.5 * random.rand(n,2)  # Spread out along x-axis
    radii = random.rand(n, 1)
    centres_radii = []
    for i in range(0, n):
        centres_radii.append([centres[i], radii[i]])
    return [generate_circle(centre, random.rand()) for centre in centres]

def compute_regions_with_outside(circles, buffer=10):
    """Finds the distinct topological regions within the arrangement of circles,
    and returns them as shapely geometries."""
```

```
    boundaries = [circle.boundary for circle in circles]
    merged_lines = unary_union(boundaries)

    regions = [geom for geom in polygonize(merged_lines)]
    return regions

def build_region_graph(circles):
    """Takes a collection of n circles, finds the topological regions
    and then iteratively removes circles to check which regions are
    adjacent in the corresponding region graph."""

    # Adds a circle which allows us to add an "outside" region
    circles[0].append(Point([0, 0]).buffer(10, resolution=200))
    circles[1].append(np.array([0, 0]))
    circles[2].append(10)

    # Finds regions
    full_regions = compute_regions_with_outside(circles[0])

    # Adds regions to region graph
    G = nx.Graph()
    for i, region in enumerate(full_regions):
        G.add_node(i, shape=region)

    # Removes a single circle and calculates new regions
    for idx, removed_circle in enumerate(circles[0][:-1]):
        remaining = []
        for j, c in enumerate(circles[0]):
            if j != idx:
                remaining.append(c)
        new_regions = compute_regions_with_outside(remaining)

        # If a new region contains an old one, it must have merged with another
        for j, new_region in enumerate(new_regions):
            merged = []
            for i, old_region in enumerate(full_regions):
                intersection = new_region.intersection(old_region)

                # Accounts for imprecision of shapely geometries
                if intersection.area / old_region.area > 0.98:
                    merged.append(i)

```

```
            # Add edges to the region graph between regions merged by removal of this circle
            for i in merged:
                for k in merged:
                    if i < k:
                        G.add_edge(i, k, removed_circle=idx)
    return G, full_regions, circles

def plot_regions_and_graph(G, regions, circles):
    """Plot the arrangement of circles, and label distinct topological regions.
    Also draw the corresponding region graph."""
    fig, ax = plt.subplots()
    label_positions = []

    # Plot the arrangement of circles
    for i, circle in enumerate(circles[0]):
        temp = plt.Circle(circles[1][i], circles[2][i], color="b", fill=False)
        ax.add_artist(temp)

    # Colour and label regions, and add node to each region
    for i, region in reversed(list(enumerate(regions))):
        x, y = region.exterior.xy
        ax.fill(x, y, alpha=0.5, label=f"Region {i}")
        if i == len(regions) - 1:
            ax.plot(-1, -1, marker="o", color="black", markersize=4)
            label_positions.append([-1, -1])
        else:
            point = region.representative_point()
            ax.plot(point.x, point.y, marker="o", color="black", markersize=4)
            label_positions.insert(0, [point.x, point.y])

    print("EDGES:", G.edges)

    # Plot line between merged regions
    for u, v in G.edges:
        x1, y1 = label_positions[u]
        x2, y2 = label_positions[v]
        ax.plot([x1, x2], [y1, y2], "k-", alpha=0.5)

    # Configure plot and display
    ax.set_aspect("equal")
```

```python
    ax.set_xlim(-1.5, 2.5)
    ax.set_ylim(-1.5, 2.5)
    plt.title(f"Region Graph of {len(circles[0])-1} circles")
    plt.show()

# Contains generated region graphs as nodes and edges
region_graph_list=[]

# Number of circles region graph is on
circle_num = []

# List of circles region graph is on, as centres and radii
circles_list = []

# Number of circles to generate region graph on, change the for loop
    below to generate desired graphs. Currently set
# to generate 9000 region graphs on increasing number of circles. For 10
    or more circles generation is fairly slow.
n=0
for i in range(0, 9000):
    if 0 <= i <= 199:
        n = 4
    if 200 <= i <= 999:
        n = 5
    if 1000 <= i <= 1999:
        n = 6
    if 2000 <= i <= 2999:
        n = 7
    if 3000 <= i <= 3999:
        n = 8
    if 4000 <= i <= 4999:
        n = 9
    if 5000 <= i <= 5999:
        n = 10
    if 6000 <= i <= 6999:
        n = 11
    if 7000 <= i <= 7999:
        n = 12
    if 8000 <= i <= 8499:
        n = 13
    if 8599 <= i <= 8999:
        n = 14

    # Adds data to lists for export
```

```python
    circle_num.append(n)
    circles_temp = generate_circles(n)
    circles_polygons = [[circles_temp[j][0] for j in range(0, n)], [
    circles_temp[j][1] for j in range(0, n)],
                  [circles_temp[j][2] for j in range(0, n)]]
    circles_list.append([[list(circles_temp[k][1]), circles_temp[k][2]]
    for k in range(0, n)])
    region_graph, regions, circles2 = build_region_graph(circles_polygons
    )
    region_graph_list.append(list(region_graph.edges))

    # Prints most recently generated region graph and associated
    information
    print("Edges", region_graph.edges)
    print("circle_layout", circles_list)
    print(f"GENERATED {i}th GRAPH")

# Formats outputs for analysis in Mathematica
print(region_graph_list)
formatted = str(region_graph_list).replace("[", "{").replace("]", "}").
    replace("(", "{").replace(")", "}")
formatted2 = "\n"+str(circle_num).replace("[", "{").replace("]", "}").
    replace("(", "{").replace(")", "}")
formatted3 = "\n"+str(circles_list)

# Write to text file
with open("GraphsOutputRandom9000.txt", "w") as f:
    f.write(formatted)
    f.write(formatted2)
    f.write(formatted3)
```

Bibliography

[1] Agelos Georgakopoulos. Circle graphs and the automorphism group of the circle, 2025. URL https://arxiv.org/abs/2501.07698. arXiv:2501.07698 [math.CO].

[2] Étienne Ghys. Groups acting on the circle. *Enseignement Mathematique*, 47(3/4):329–408, 2001. URL https://perso.ens-lyon.fr/ghys/articles/groupscircle.pdf.

[3] Andrés Navas. *Groups of circle diffeomorphisms*. University of Chicago Press, 2011. URL https://arxiv.org/abs/math/0607481.

[4] L A Beklaryan. Groups of homeomorphisms of the line and the circle. Topological characteristics and metric invariants. *Russian Mathematical Surveys*, 59(4):599, aug 2004. URL https://dx.doi.org/10.1070/RM2004v059n04ABEH000758.

[5] Terry A. McKee and F. R. McMorris. *Topics in Intersection Graph Theory*. Society for Industrial and Applied Mathematics, 1999. URL https://epubs.siam.org/doi/abs/10.1137/1.9780898719802.

[6] Madhumangal Pal. Intersection graphs: An introduction. *Annals of Pure and Applied Mathematics*, 4(1):43–91, 2014. URL https://arxiv.org/abs/1404.5468.

[7] Frédéric Le Roux and Maxime Wolff. Automorphisms of some variants of fine graphs. *Algebraic & Geometric Topology*, 24(8):4697–4730, 2024. URL http://dx.doi.org/10.2140/agt.2024.24.4697.

[8] Alan F. Beardon. *The Geometry of Discrete Groups*. Springer, 1983. URL https://link.springer.com/book/10.1007/978-1-4612-1146-4.

[9] Frederick W. Gehring, Gaven J. Martin, and Bruce P. Palka. *An Introduction to the Theory of Higher-Dimensional Quasiconformal Mappings*. Mathematical Surveys and Monographs, 2017. URL https://bookstore.ams.org/surv-216/.

[10] J. B. Wilker. Inversive Geometry. In Chandler Davis, Branko Grünbaum, and F. A. Sherk, editors, *The Geometric Vein*, pages 379–442, New York, NY, 1981. Springer New York. URL https://link.springer.com/chapter/10.1007/978-1-4612-5648-9_27.

[11] Brian C. Hall. *Lie Groups, Lie Algebras, and Representations.* Springer Cham, 2015. URL https://doi.org/10.1007/978-3-319-13467-3.

[12] D. J. H. Garling. *Clifford algebras: an introduction*, volume 78 of *London Mathematical Society Student Texts*. Cambridge University Press, 2011. URL https://doi.org/10.1017/CBO9780511972997.

[13] Lars V. Ahlfors. Möbius transformations and Clifford numbers. In *Differential Geometry and Complex Analysis: A Volume Dedicated to the Memory of Harry Ernest Rauch*, pages 65–73. Springer, 1985. URL https://doi.org/10.1007/978-3-642-69828-6_5.

[14] Victor Pambuccian. A logical look at characterizations of geometric transformations under mild hypotheses. *Indagationes Mathematicae*, 11(3):453–462, 2000. URL https://doi.org/10.1016/S0019-3577(00)80009-9.

[15] Baokui Li and Guowu Yao. On characterizations of sphere-preserving maps. *Mathematical Proceedings of the Cambridge Philosophical Society*, 147(2):439–446, 2009. URL https://doi.org/10.1017/S0305004109002291.

[16] Chris Godsil and Gordon Royle. *Algebraic Graph Theory.* Springer, 2001. URL https://link.springer.com/book/10.1007/978-1-4613-0163-9.

[17] J. M. Lee. *Introduction to Smooth Manifolds, second edition.* Springer, 2012. URL https://link.springer.com/book/10.1007/978-1-4419-9982-5.

[18] José Burillo, Sean Cleary, Melanie Stein, and Jennifer Taback. Combinatorial and metric properties of Thompson's group T. *Transactions of the American Mathematical Society*, 361(2):631–652, 2009. URL https://www.ams.org/journals/tran/2009-361-02/S0002-9947-08-04381-X/.

[19] Allen Hatcher. *Algebraic Topology.* Cambridge University Presses, 2002. URL https://pi.math.cornell.edu/~hatcher/AT/AT.pdf.

[20] J. W. Cannon, W. J. Floyd, and W. R. Parry. Introductory notes on Thompson's groups. *L'Enseignement Mathématique*, 42(3–4), 1996. URL https://www.imo.universite-paris-saclay.fr/~emmanuel.breuillard/Cannon.pdf.

[21] Wilhelm Magnus. Residually finite groups. *Bulletin of the American Mathematical Society*, 75(2):305–316, 1969. URL https://projecteuclid.org/journals/bulletin-of-the-american-mathematical-society/volume-75/issue-2/Residually-finite-groups/bams/1183530287.full.

[22] J. W. Alexander. A proof and generalization of the Jordan-Brouwer Theorem. *Trans. Amer. Math. Soc*, 23:333–349, 1922. URL https://webhomes.maths.ed.ac.uk/~v1ranick/jordan/alex5.pdf.

[23] Andrew Putman. The generalised Schoenflies theorem, 2015. URL https://www3.nd.edu/~andyp/notes/Schoenflies.pdf.

[24] Franca Lippert. The Alexander horned sphere and Bing's hooked rug. *Topological Manifolds II Seminar - University of Bonn*, 2021. URL https://www.maths.gla.ac.uk/~mpowell/WriteUp_Franca.pdf.

[25] R. H. Bing. A homeomorphism between the 3-sphere and the sum of two solid horned spheres. *Annals of Mathematics*, 56(2):354–362, 1952. URL https://doi.org/10.2307/1969804.

[26] Carsten Thomassen. The Jordan-Schönflies theorem and the classification of surfaces. *The American Mathematical Monthly*, 99(2):116–130, 1992. URL https://doi.org/10.1080/00029890.1992.11995820.

[27] Manfred Sauter. How twisted can a Jordan curve be? *Ulmer Seminare*, Heft 20:133–140, 2017. URL https://www.uni-ulm.de/fileadmin/website_uni_ulm/mawi.inst.020/sauter/files/Sauter-TwistedJordan.2017.usem20.pdf. Expository note.

[28] Ralph H. Fox and Emil Artin. Some wild cells and spheres in three-dimensional space. *Annals of Mathematics*, 49(4):979–990, 1948. URL https://doi.org/10.2307/1969408.

[29] William F. Osgood. A Jordan curve of positive area. *Transactions of the American Mathematical Society*, 4(1):107–112, 1903. ISSN 00029947. URL http://www.jstor.org/stable/1986455.

[30] Hans Sagan. *Space-filling curves*. Springer, 1994. URL https://link.springer.com/book/10.1007/978-1-4612-0871-6.

[31] János Pach, Natan Rubin, and Gábor Tardos. A crossing lemma for Jordan curves. *Advances in Mathematics*, 331:908–940, 2018. URL https://www.sciencedirect.com/science/article/pii/S0001870818301026.

[32] R. H. Bing. Generalizations of two theorems of Janiszewski. *Bulletin of the American Mathematical Society*, 51:954–960, 1945. URL https://www.ams.org/journals/bull/1945-51-12/S0002-9904-1945-08480-8/. Minor corrections made in "Generalizations of two theorems of Janiszewski II".

[33] Mamoru Watanabe and Allen J. Schwenk. Integral Starlike Trees. *Journal of the Australian Mathematical Society*, 28:120–128, 1979. URL https://doi.org/10.1017/S1446788700014981.

[34] OEIS Foundation Inc. Sequence A000055. *The On-Line Encyclopedia of Integer Sequences*, 2024. Published electronically at http://oeis.org.

[35] H. N. de Ridder et al. Graphclass: cactus. *Information System on Graph Classes and their Inclusions (ISGCI)*, 2025. URL https://www.graphclasses.org/classes/gc_108.html.

[36] A. F. Beardon. Continued fractions, Möbius transformations and Clifford algebras. *Bulletin of the London Mathematical Society*, 35(3):302–308, 2003. URL https://doi.org/10.1112/S0024609302001807.

[37] Jimmie Lawson. Clifford algebras, Möbius transformations, Vahlen matrices and *b*-loops. *Commentationes Mathematicae Universitatis Carolinae*, 51(2):319–331, 2010. URL https://eudml.org/doc/37763.

[38] Johannes Gerrit Maks. *Modulo (1, 1) periodicity of Clifford algebras and generalized (anti-) Möbius transformations.* PhD thesis, Technische Universiteit Delft, 1989. URL https://repository.tudelft.nl/record/uuid:af420b59-3d85-42b6-8224-ea39b07b5bfe.

[39] J. Maks. *Clifford algebras and Möbius transformations*, pages 57–63. Springer Netherlands, 1992. URL https://doi.org/10.1007/978-94-015-8090-8_6.