

MacLane's Theorem for Graph-Like Spaces

by

Brendan Rooney

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2008

©Brendan Rooney 2008

AUTHOR'S DECLARATION

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Brendan Rooney

Abstract

The cycle space of a finite graph is the subspace of the edge space generated by the edge sets of cycles, and is a well-studied object in graph theory. Recently progress has been made towards extending the theory of cycle spaces to infinite graphs.

Graph-like spaces are a class of topological objects that reconcile the combinatorial properties of infinite graphs with the topological properties of finite graphs. They were first introduced by Thomassen and Vella as a natural, general class of topological spaces for which Menger's Theorem holds. Graph-like spaces are the natural objects for extending classical results from topological graph theory and cycle space theory to infinite graphs.

This thesis focuses on the topological properties of embeddings of graph-like spaces, as well as the algebraic properties of graph-like spaces. We develop a theory of embeddings of graph-like spaces in surfaces. We also show how the theory of edge spaces developed by Vella and Richter applies to graph-like spaces. We combine the topological and algebraic properties of embeddings of graph-like spaces in order to prove an extension of MacLane's Theorem. We also extend Thomassen's version of Kuratowski's Theorem for 2-connected compact locally connected metric spaces to the class of graph-like spaces.

Acknowledgements

This thesis would not have been possible without the guidance and support of my supervisor Bruce Richter. Bruce had a tremendous influence on this research project; and his comments on this thesis have helped enormously to shape the final draft. His interest, encouragement and effusive enthusiasm have kept me motivated through the more difficult and tedious parts of this program. I will remember fondly our long, rambling, and thoroughly enjoyable discussions on graph theory and mathematics in general.

I am also appreciative to Carsten Thomassen for sharing some of his insights into graph-like spaces. His comments and ideas were invaluable in developing the material in Chapter 6.

I would like to thank my committee, Chris Godsil and Ian Goulden, for taking the time to read my thesis and provide feedback.

Financial support was provided by the Natural Sciences and Engineering Research Council of Canada.

I have been spared the task of wrestling with \LaTeX , and worries about typesetting requirements, thanks to the University of Waterloo thesis template created and made publicly available by Matthew Skala.

Finally, I am grateful to my friends and family for their support, and for making my time in Waterloo bearable.

Table of Contents

List of Figures	xi
Notation	xiii
1 Introduction	1
1.1 Infinite Graphs and Graph-Like Spaces	2
1.2 Outline	3
2 Topological Spaces	5
2.1 Topological Spaces	5
2.2 Graph-Like Spaces	8
2.3 A Topological Lemma	14
2.3.1 Recursive Procedure	16
2.3.2 Proof of Lemma 2.13	23
3 Embeddings and Face Boundaries	33
3.1 Topological Properties of Embeddings	34
3.2 Edges	36
3.3 Face Boundaries	46
3.4 Connectedness and Embeddings	71
4 The Thin Cycle Space and the Face Boundary Space	73
4.1 Algebraic Edge Spaces and Topological Edge Spaces	74
4.2 A Property of Surfaces	76
4.3 The Face Boundary Space	78
4.4 The Face Boundaries of a Disconnected Graph-Like Space	84
4.5 The Thin Cycle Space	87

5	MacLane’s Theorem	93
5.1	Face Boundaries in the Plane	94
5.2	2-Bases and MacLane’s Theorem	95
6	Thumbtacks and Kuratowski’s Theorem	101
6.1	Topological Properties of Embeddings	102
6.2	The 2-Connected Subspaces of a Graph-Like Space	106
6.3	Embedding 2-Connected Graph-Like Spaces in the Plane	109
6.4	Embedding Connected Graph-Like Spaces in the Plane	116
6.5	MacLane’s Theorem and Kuratowski’s Theorem	124
7	Conclusion	127
7.1	Future Research	127
	Bibliography	131
	Index	133

List of Figures

2.1	The 2-way infinite ladder L	10
3.1	Arcs l_1 and l_2	54
3.2	Arcs l'_1 and l'_2	56
3.3	Arcs σ_1 and σ_2	56
3.4	Arcs a_1, a_2, a'_1 , and a'_2	57
3.5	Arcs σ'_1 and σ'_2	58
3.6	Simple closed curves τ'_1 and τ'_2	58
3.7	Construction of τ'_3 in Case #2.	60
3.8	Construction of l'_3 in Case #2.	61
3.9	Construction of τ'_3 and τ''_3 in Case #3.	63
3.10	Construction of l'_3 and l''_3 in Case #3.	64
5.1	Graph-like space G' embedded in the torus.	96
6.1	Graph-like space W embedded in \mathbb{S}^2	110
7.1	Edge space E embedded in \mathbb{R}^2	128

Notation

The notation in this thesis is standard. Specific symbols and functions are defined as needed throughout, and the standard meaning of basic symbols and functions is assumed. The following is a list of symbols that appear frequently in the text, and the meanings they are intended to convey.

\emptyset	The empty set.
$a := b$	a is defined as b .
$x \equiv_n y$	x is congruent to y modulo n .
$x \not\equiv_n y$	x is not congruent to y modulo n .
$A - B$	The set $\{a \in A : a \notin B\}$.
$ S $	The cardinality of the set S .
$f _S$	The restriction of the function f to the set S .
$\{x_i\}_{i \in I}$	Given an ordered index set $I = \{i_1, i_2, \dots\}$, the sequence x_{i_1}, x_{i_2}, \dots
$\text{Cl}(S)$	The closure of the topological space S .
$\text{Bd}(S)$	The boundary of the topological space S .
$d(x, y)$	The distance from x to y .
$\text{diam}(S)$	The diameter of the set or space S .
$B(v, \epsilon)$	The open neighbourhood of v consisting of all points at distance less than ϵ from v .
$B(0, 1)$	The open unit disk in the plane, centred at the origin.
\mathbb{S}^1	The unit circle, equivalent to $\text{Bd}(B(0, 1))$.
\mathbb{S}^2	The sphere.
\mathbb{R}^2	The plane.

Chapter 1

Introduction

Graphs are combinatorial objects that encode the notion of adjacency between elements of a given set. Graphs can also be viewed as topological objects and as algebraic objects. There are deep connections between graph theory and both topology and algebra. In this thesis we extend MacLane's Theorem, a classical result in graph theory that combines the topological and algebraic properties of finite graphs.

A *graph* is a pair of finite sets $G = (V, E)$. The set V consists of the *vertices* of G and the set E , the set of *edges* of G , is a set of unordered pairs of vertices. Alternatively, we can define G as a topological space. A *graph* is a set V of vertices, together with a set E of disjoint arcs that connect the vertices in pairs. This definition of a graph allows for topological questions to be addressed. For instance we can consider the topological spaces in which a graph G can be embedded. Much work has been done in the area of embeddings of graphs, particularly embeddings in surfaces and the plane. Classical results include characterizations of planarity, colouring results such as the 4-colour Theorem and Heawood's Theorem, and the deep results of the graph minors project of Robertson and Seymour.

From a graph G we can also consider several associated algebraic objects. For instance we can consider the vector space formed by subsets of E . The characteristic vector of a subset $A \subseteq E$ is the vector $a \in 2^E$ where $a_e = 1$ if and only if $e \in A$. The edge space 2^E of G is a vector space consisting of all characteristic vectors of subsets of E , where vector addition is componentwise addition modulo 2. Vector addition corresponds to the operation of symmetric

difference on the subsets of E . Together with the space 2^E , we consider subspaces generated by subgraphs of G with specific properties. For example, we can consider the space $\mathcal{C}(G)$ generated by the edge sets of cycles of G , and the space $\mathcal{E}(G)$ generated by the edge cuts of G . These spaces are orthogonal to one another.

The vector space $\mathcal{C}(G)$ is related to embeddings of G . Given any embedding of G in some surface, the face boundaries of G form a subspace $\mathcal{B}(G)$ of 2^E , which is a subspace of $\mathcal{C}(G)$. The theorems of MacLane, Tutte and Whitney give connections between the cycle space of a graph and embeddings in the plane.

1.1 Infinite Graphs and Graph-Like Spaces

In the above definition of a graph we considered finite sets V and E . Infinite graphs are the objects that result from altering the definition of a graph to allow for infinite vertex and edge sets. This thesis considers the topological properties of infinite graphs, together with the vector spaces on the edge set of an infinite graph. Our goal is to extend the classical theory of cycle spaces of finite graphs to infinite graphs. However, if we take the definition of the cycle space of an infinite graph to be the same as the definition of the cycle space of a finite graph, the classical results concerning cycle spaces are no longer true. Difficulties arise from the fact that infinite graphs contains no infinite cycles, and the sum of infinitely many cycles may have undesired properties. In order to deal with these difficulties, two approaches have been taken.

Diestel and Kühn developed a theory of cycle spaces for infinite graphs in [6], [7] and [8]. They consider the Freudenthal compactification of a natural topological space associated with a locally finite infinite graph. Their approach focuses on retaining the combinatorial structure of the original graph, while adding sufficiently many topological requirements to allow for a theory of cycle spaces that parallels the classical cycle space theory for finite graphs. Diestel and Kühn are able to extend much of the classical theory to their class of objects, however their approach restricts the types of infinite graphs they can consider.

Vella and Richter approached the cycle space of an infinite graph from a different perspective in [19]. They introduced topological edge spaces as a large class of topological spaces in which the algebraic arguments regarding cycle spaces still apply. In a topological edge space an edge is not an arc, but an open

singleton, and vertices are not necessarily points. The properties of topological edge spaces are designed to be minimal so that graph-theoretic connection corresponds to topological connection. The papers of Vella and Richter, [19], and Casteels and Richter, [3], use connectivity together with algebraic techniques to develop a theory of cycle spaces in topological edge spaces. Since topological edge spaces can be naturally derived from graphs and compactifications of infinite graphs, this gives a generalization of cycle spaces in finite graphs and unifies the work of Diestel and Kühn with that of Bonnington and Richter [1].

Graph-like spaces were introduced by Thomassen and Vella [18]. They were designed as a class of topological spaces in which a version of Menger's Theorem holds. Graph-like spaces give a natural context for studying infinite graphs. Every compactification of an infinite graph is a graph-like space. In particular, the topological space Diestel and Kühn associate with an infinite graph is, almost always, a graph-like space. Furthermore, the theory of topological edge spaces developed by Richter and Vella applies to graph-like spaces. Throughout this thesis we restrict our attention to graph-like spaces, as they are the natural objects for extending classical results from topological graph theory and classical cycle space theory to infinite graphs.

1.2 Outline

The chapters of this thesis can be broken into two parts. The first part consists of Chapters 2, 3 and 4. These chapters focus on developing a theory of embeddings of graph-like spaces, and a theory of face boundary spaces for graph-like spaces. Chapter 2 provides some of the topological background required, and surveys the work of Thomassen and Vella on graph-like spaces. The only new material in Chapter 2 is Section 2.3, wherein a new lemma is proved about decomposing topological spaces. In Chapter 3 we develop a theory of embeddings of graph-like spaces. We prove several technical and foundational results on embedding graph-like spaces in surfaces. These results demonstrate that embeddings of graph-like spaces behave very similarly to embeddings of graphs, and allow us to use combinatorial arguments when considering embedded graph-like spaces. Chapter 4 provides a specialization of the Vella-Richter theory of edge spaces to graph-like spaces. We provide a brief overview of edge spaces, and show that graph-like spaces can be viewed as edge spaces. In this context, we are able to

prove that the face boundaries of an embedded graph-like space form a subspace of the algebraic edge space, and of the cycle space of a graph-like space.

The second part of the thesis consists of Chapters 5 and 6. These chapters contain applications of the theory developed in the preceding chapters. In Chapter 5 we give a proof of MacLane's Theorem for 2-connected graph-like spaces. The proof relies on Thomassen's version of Kuratowski's Theorem for 2-connected metric spaces. The argument is surprisingly short, and pleasant in that it follows very closely the standard proof of MacLane's Theorem for finite graphs. In Chapter 6 we extend Thomassen's version of Kuratowski's Theorem to the class of graph-like spaces. We identify an additional class of forbidden spaces, and use them to characterize the class of graph-like spaces that can be embedded in the plane. We also note that our version of MacLane's Theorem in Chapter 5 can be extended to all graph-like spaces using this new result. Finally, in Chapter 7 we provide some ideas for future research in the theory of embeddings of graph-like spaces.

Chapter 2

Topological Spaces

The goal of this thesis is to extend well-known graph-theoretic results to infinite graphs. Infinite graphs can be viewed both as combinatorial and topological objects. In a combinatorial setting, infinite graphs do not have the properties required to extend many graph-theoretic results directly. The natural framework for considering these questions is topology.

In this chapter we discuss the topological spaces and properties that will be used in subsequent chapters. Section 2.1 introduces the standard topological facts we will use throughout this thesis. In Section 2.2, we introduce the graph-like spaces of Thomassen and Vella [18], which are the focus of this thesis. We conclude in Section 2.3 with a decomposition theorem, that shows how to decompose a continuous function from a circle into a topological space into a collection of embeddings of the circle. This is the most technical part of the thesis, and the main result is applied only once, in the proof of Theorem 4.13.

2.1 Topological Spaces

This section provides a brief list of the topological terms that will appear in the remainder of the thesis. We include this material as a reminder of the properties that we will use most often, and to serve as a reference. We will be considering embeddings of topological spaces, so we will also need some basic properties of surfaces.

The spaces encountered in this thesis will, for the most part, be compact, Hausdorff metrizable spaces. A topological space X is called *Hausdorff* if for

each pair x_1, x_2 of distinct points of X , there exist neighbourhoods U_1 and U_2 , of x_1 and x_2 respectively, that are disjoint. The space X is said to be *compact* if every open covering \mathcal{A} of X contains a finite subcollection that also covers X . A *separation* of X is a pair of non-empty open sets U and V such that $U \cap V = \emptyset$ and $U \cup V = X$. A space X is *connected* if no separation of X exists. If A and B are subsets of a topological space X then we say that A and B can be *separated* in X if there is a separation, U, V , of X so that $A \subseteq U$ and $B \subseteq V$.

We will make frequent use of the properties of metric spaces. A *metric* on a set X is a function $d : X \times X \rightarrow \mathbb{R}$ having the following properties:

1. $d(x, y) \geq 0$ for all $x, y \in X$;
2. $d(x, y) = 0$ if and only if $x = y$;
3. $d(x, y) = d(y, x)$ for all $x, y \in X$;
4. $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$.

We refer to the value $d(x, y)$ as the *distance* between x and y . Given a metric space X , we can use the metric in order to specify open neighbourhoods of points in X . If X is a topological space and d is a metric on X then we denote by $B_d(x, \epsilon)$ the subset of X consisting of all points y so that $d(x, y) < \epsilon$. We will often write $B(x, \epsilon)$ for $B_d(x, \epsilon)$ since the space and metric under consideration will be clear from the context. Topological space X is *metrizable* if there exists a metric d on the set X that induces the topology of X .

Given a metric space X , the concept of distance can be extended to any subset of X . A subset A of a metric space is *bounded* if there is some M such that $d(x, y) \leq M$ for all $x, y \in A$. We have by Theorem 20.1 in [10] that for any metric space X with metric d , there is a metric d' on X so that d and d' both give the topology of X , and every subset of X is bounded with respect to d' . Given a subset A of a metric space X the *diameter* $\text{diam}(A)$ of A is defined to be $\sup\{d(x, y) : x, y \in A\}$.

For example, consider a graph G with finite vertex set V and finite edge set E . We can construct a natural topological space Y by taking the vertex set V , with the discrete topology, together with a set of disjoint homeomorphs of the open unit interval, each corresponding to an element of E . We specify a basis B for this topology as follows. Basis B consists of every open subset of every edge, together with a set B_v for each vertex, where B_v is the vertex v together with an

open segment of each edge incident with v in G . Now Y is a compact Hausdorff metrizable space. Unsurprisingly, Y is a graph-like space. We discuss graph-like spaces further in Section 2.2.

In working with topological spaces we will assume several standard results on compactness. These can be found in [10] or [4].

In our example space, Y , the notion of graph-theoretic connection translates into topological connection. More specifically if H is a connected subgraph of G , then H corresponds to a subspace Y' of Y that is arcwise connected. Arcs in topological spaces, together with arcwise connection will be very important throughout this thesis.

An *arc* A is a space homeomorphic to the closed unit interval; the *endpoints* of A are the points p and q so that $A-p$ and $A-q$ are connected. A space X is *arcwise connected* if for any $x, y \in X$ there is an arc $A \subset X$ with endpoints x and y . We say that X is *locally arcwise connected* if for any $x \in X$, for every neighbourhood U of x there is a arcwise connected neighbourhood V of x such that $V \subseteq U$.

Finally, we have the following proposition, which appears as an assertion in [10]. It will be useful in Section 2.3 as well as in subsequent chapters, so we provide a proof.

Proposition 2.1

If X is a compact Hausdorff space and \mathcal{A} is a collection of closed connected subsets of X ordered by inclusion, then $Y = \bigcap_{A \in \mathcal{A}} A$ is a non-empty connected subset of X .

Proof For each $A \in \mathcal{A}$, let $U_A = X - A$. If Y is empty, then $\bigcup_{A \in \mathcal{A}} U_A = X$. Since each U_A is open, $\{U_A : A \in \mathcal{A}\}$ is an open cover of X . Since X is compact, we have a finite subcover $\{U_{A_1}, \dots, U_{A_n}\}$. The sets A are ordered by inclusion, so there is some A_i that is inclusion-wise minimal, and $A_i = \bigcap_{j=1}^n A_j$. Since $A_i \neq \emptyset$, there is some $x \in A_i$, and $x \notin \bigcup_{j=1}^n U_{A_j}$. Thus $\{U_{A_1}, \dots, U_{A_n}\}$ does not cover X , a contradiction.

To prove that Y is connected we proceed by contradiction. Suppose that C, D is a separation of Y . Since C and D are closed in Y and Y is closed in X , C and D are closed in X and hence compact.

We claim that for each $x \in C$ we have a neighbourhood C_x of x disjoint from D . This can be seen as follows.

For each $y \in D$, let U_x^y and U_y be disjoint neighbourhoods of x and y in X . Now the set $\{U_y : y \in D\}$ is an open cover of D , so it has a finite subcover $\{U_{y_i} : 1 \leq i \leq n\}$. Thus $D_x = \bigcup_{i=1}^n U_{y_i}$ is an open set containing D . Further D_x is

disjoint from $\bigcap_{i=1}^n U_x^{y_i} = C_x$. Now C_x is an open set containing x and disjoint from $D_x \supseteq D$.

Now we have collections $\{C_x : x \in C\}$ and $\{D_x : x \in C\}$ such that $D \subseteq D_x$ for all x and $x \in C_x$ for all x . Thus $\{C_x : x \in C\}$ is an open cover of C . Therefore we have a finite subcover of C , $\{C_{x_i} : 1 \leq i \leq m\}$. Now

$$\begin{aligned} C &\subseteq \bigcup_{i=1}^m C_{x_i} = C', \quad \text{and} \\ D &\subseteq \bigcap_{i=1}^m D_{x_i} = D' \end{aligned}$$

where C' and D' are both open in X . Furthermore $C' \cap D' = \emptyset$, since $C_{x_i} \cap D_{x_i} = \emptyset$ for $1 \leq i \leq m$.

Consider $\mathcal{A}' = \{A - (C' \cup D') : A \in \mathcal{A}\}$. Since \mathcal{A} is ordered by inclusion, so too is \mathcal{A}' . Also each $A - (C' \cup D')$ is closed, since A and $X - (C' \cup D')$ are closed and

$$A - (C' \cup D') = A \cap (X - (C' \cup D')).$$

Finally, $A - (C' \cup D')$ must be non-empty since otherwise $A \subseteq (C' \cup D')$, contradicting the connectedness of A . Thus

$$\emptyset \neq \bigcap_{A' \in \mathcal{A}'} A' = Y - (C' \cup D') \subseteq Y - (C \cup D),$$

a contradiction. ■

2.2 Graph-Like Spaces

The main objects considered in this thesis are graph-like spaces. Thomassen and Vella introduced graph-like spaces as a natural, general class of topological spaces for which Menger's Theorem holds. In this section we present some results on graph-like spaces from [18].

First we provide the definition of a zero-dimensional space. A topological space is *zero-dimensional* if whenever $x \in V$, and V is open, there is an open set U with empty boundary such that $x \in U \subseteq V$.

Definition 2.2

Given a topological space G , an edge is an open subset of G , homeomorphic to \mathbb{R} , whose closure is a simple arc. A graph-like space is a topological space G equipped

with a collection E of pairwise disjoint edges such that $V = G - E$ is zero-dimensional. We refer to V as the vertex set of G .

In [18] the authors are concerned primarily with compact, Hausdorff, metrizable graph-like spaces. In the remainder of the thesis we take “graph-like space” to mean “compact, Hausdorff, metrizable graph-like space.”

The zero-dimensionality criterion can be somewhat cumbersome to work with. However in compact Hausdorff spaces we can consider totally disconnected spaces instead. A topological space X is *totally disconnected* if the connected components of X are singletons.

Theorem 2.3 ([4], Thm. 6.C.1)

If X is a totally disconnected compact Hausdorff space, then the family of open and closed subsets of X forms a basis for X .

If a set A in a topological space X is both closed and open, then the boundary of A is empty.

Proposition 2.4

If X is a compact Hausdorff space, then X is zero-dimensional if and only if X is totally disconnected.

Proof If X is totally disconnected, then by Theorem 2.3 the family of open and closed subsets of X forms a basis for X . If $x \in V$ and V is an open subset of X , we have a basis element U contained in V so that $x \in U$. Thus there is an open and closed subset U of X that contains x , and is contained in V . Hence X is zero-dimensional.

Suppose that X is zero-dimensional. We wish to show that every connected component of X is a singleton. Towards a contradiction suppose that we have a connected component C of X and distinct points $x, y \in C$. Since X is Hausdorff we have neighbourhoods V_x and V_y of x and y respectively so that V_x, V_y are disjoint. Since V_x, V_y are open and X is zero-dimensional we have subsets U_x and U_y of V_x and V_y respectively so that $x \in U_x$, $y \in U_y$ and U_x, U_y both have empty boundary. Thus U_x and U_y are disjoint closed subsets of X . Theorem 4.A.11 from [4] states that if I and J are disjoint closed subsets of X , and no connected subset of X intersects both I and J , then I and J can be separated in X . Therefore we can find a separation A, B of X so that $U_x \subset A$ and $U_y \subset B$. However, $A \cap C$ and $B \cap C$ give a separation of C , a contradiction. Thus X is totally disconnected. ■

We can redefine graph-like spaces as follows. A topological space G equipped with a collection E of pairwise disjoint arcs is graph-like if G is compact, Hausdorff and metrizable, and $G - E$ is totally disconnected.

Recall that in Section 2.1 we constructed a topological space Y from a finite graph G . The space Y is an example of a graph-like space. We can also construct a graph-like space from any compactification of any locally finite infinite graph. A *locally finite* infinite graph G is an ordered pair (V, E) . Where V is an infinite set of *vertices* of G and E , the set of *edges* of G , is a subset of unordered pairs of vertices. Furthermore, each $v \in V$ is incident with only finitely many edges of G . If we apply the construction given in Section 2.1 verbatim, the resulting space Y will not be compact. For instance, if G is connected, then G will contain a *ray*, a path $R = (v_0, v_1, \dots)$ of infinite length. In order to construct a graph-like space from a connected locally finite infinite graph G , we will need to “compactify” G .

We define an equivalence relation on the rays of G . If $R = (v_0, v_1, \dots)$ is a ray, a *tail* of R is any ray $R' = (v_i, v_{i+1}, \dots)$ for some finite i . Given rays R_1 and R_2 , R_1 is equivalent to R_2 if and only if, for each finite set $S \subset V$, R_1 and R_2 both have a tail that lies in the same component of $G - S$. We call the equivalence classes of rays the *ends* of G , and we let Ω be the set of ends. For example consider the 2-way infinite ladder, L , shown in Figure 2.1. The graph L consists of two disjoint doubly infinite rays, $R_1 = (\dots, v'_2, v'_1, v_0, v_1, v_2, \dots)$ and $R_2 = (\dots, u'_2, u'_1, u_0, u_1, u_2, \dots)$, together with edges between each v_i and u_i , and each v'_i and u'_i . Note that L has two ends, ω_1 and ω_2 .




Figure 2.1: The 2-way infinite ladder L .

Now we define the *Freudenthal compactification* of a locally finite infinite graph G . For each $\omega \in \Omega$ and any finite subset $S \subset V$ we define the basic open set

of ω with respect to S as follows. Let C be the component of $G - S$ containing the tails of the rays in ω . Then the open set is $C \cup \Omega(C) \cup \delta(C)$, where $\Omega(C) \subseteq \Omega$ is the set of ends ω of G so that each ray in ω has a tail in C , and $\delta(C)$ is a set of half open intervals, one for each edge between C and S . We define the basic open sets for a vertex $v \in V$, and the basic open sets for an edge $e \in E$ as before. These basic open sets give us a compact topological space Y . The space Y is graph-like. Note that unlike graph-like space corresponding to a finite graph, the induced topology on the vertex set of Y is not discrete.

Note that the Freudenthal compactification is only one of many possible compactifications of a locally finite infinite graph. Each end of G needs to be identified with a point of the topology, however, we do not need to add a new point for each end. For instance, once we have the Freudenthal compactification of G , we can identify all of the points in Ω with a single point ω . This is a quotient map from the Freudenthal compactification of G to the *Alexandroff compactification* or *1-point compactification* of G . In general we can identify any closed subset of Ω with any point $v \in V$ or any other point $\omega \in \Omega$ to obtain a compactification of G . These spaces are all graph-like.

We now present some of the properties of graph-like spaces. We have the following lemma (this argument appears as a claim in the proof of Theorem 2.1 from [18]).

Lemma 2.5

Suppose that G is a graph-like metric space. Then, for every $\epsilon > 0$, there are only finitely many edges with diameter larger than ϵ .

Proof Suppose that for some $\epsilon > 0$ we have an infinite set of edges $\{e_1, e_2, e_3, \dots\}$, each with diameter larger than ϵ . Then we have points p_i and q_i such that $p_i, q_i \in e_i$ and $d(p_i, q_i) > \epsilon$ for each i . Since G is compact there are convergent subsequences $\{p_i\}_{i \in I}$ and $\{q_i\}_{i \in I}$ that converge to p and q respectively.

Since edges are open p and q must be vertices of G . Since the vertex set of G is zero-dimensional, it can be partitioned into disjoint closed sets P and Q such that $p \in P$ and $q \in Q$. Since V is closed, P, Q are closed sets of G . Since G is normal, there are disjoint open sets P', Q' in G such that $P \subseteq P'$ and $Q \subseteq Q'$. Now since the sequence of p_i 's converges to p , and the sequence of q_i 's converges to q , there is some index i such that $p_j \in P'$ and $q_j \in Q'$ for all $j > i$.

Let $I' = \{j \in I : j > i\}$. We have that e_j is a connected subset of G for each

edge e_j . Now

$$(P' \cap e_j) \cup (Q' \cap e_j) \neq e_j,$$

and, for each $j \in I'$, we have a point $z_j \in e_j - (P' \cup Q')$. By compactness there is some subsequence of $\{z_j\}_{I'}$ that converges. Say $\{z_j\}_{j \in J}$ converges to z , for some vertex z . But since P, Q is a partition of V we have that $z \in P \cup Q$ and $z \in P' \cup Q'$. Thus $z_k \in P' \cup Q'$ for some index $k \in J$, and hence $k \in I'$, a contradiction. ■

Lemma 2.5 gives us a natural corollary.

Corollary 2.6

If G is a graph-like metric space then G has countably many edges.

Proof We simply apply Lemma 2.5 for $\epsilon \in \{1, 1/2, 1/4, \dots\}$. For each i , there are only finitely many edges with diameter greater than $1/2^i$ and less than or equal to $1/2^{i-1}$. This sequence partitions the edge set of G by length into countably many parts, each of which contains finitely many edges. Thus G has countably many edges. ■

Corollary 2.6 is important as it will allow us to use induction as a method of proof. Also, if we return to our example space Y , we can use the same construction to derive graph-like spaces from infinite graphs with arbitrary degree, provided that they have only countably many edges.

We have the following results on graph-like spaces. A topological space X is *locally connected* if for each $x \in X$ and neighbourhood V of x there is a connected set $U \subset V$ with $x \in U$.

Theorem 2.7 ([18], Thm. 2.1)

Let G be a metric graph-like space. Then G is locally connected.

Proposition 2.8 ([18], Prop. 2.2)

Suppose X is a graph-like metric space, and that H is a closed connected subset of X . Then H is a graph-like space.

The space X is *hereditarily locally connected* if every closed connected subset of X is locally connected.

Corollary 2.9 ([18], Cor. 2.3)

Suppose X is a graph-like metric space. Then X is hereditarily locally connected.

We will not need the full strength of Corollary 2.9, but we will use the fact that if G is graph-like then G is locally connected. We also have the following corollary.

Corollary 2.10 ([18], Cor. 2.4)

Suppose X is a graph-like metric space. Then every closed, connected subspace of X is arcwise connected.

We can use Corollary 2.10 to prove the following strengthening of Lemma 2.5.

Lemma 2.11

Let $\{C_i\}_{i \in \mathbb{N}}$ be a sequence of edge-disjoint connected closed subsets of the graph-like space G . Then $\{\text{diam}(C_i)\}_{i \in \mathbb{N}} \rightarrow 0$.

Proof Suppose that for some $\epsilon > 0$ we have an infinite subsequence $\{C_i\}_{i \in I}$, each with diameter larger than ϵ . Then we have points p_i and q_i such that $p_i, q_i \in C_i$ and $d(p_i, q_i) > \epsilon$ for each $i \in I$. Since G is compact there are convergent subsequences $\{p_i\}_{i \in I}$ and $\{q_i\}_{i \in I}$ that converge to p and q respectively.

Since edges are open p and q must be vertices of G . Since the vertex set of G is zero-dimensional, it can be partitioned into disjoint closed sets P and Q such that $p \in P$ and $q \in Q$. Since V is closed, P, Q are closed sets of G . Since G is normal, there are disjoint open sets P', Q' in G such that $P \subseteq P'$ and $Q \subseteq Q'$. Now since the sequence of p_i 's converges to p , and the sequence of q_i 's converges to q , there is some index $i \in I$ such that $p_j \in P'$ and $q_j \in Q'$ for all $j > i$.

Let $I' = \{j \in I : j > i\}$. For each C_j we have an arc $\alpha_j \subset C_j$ with ends p_j and q_j . Since α_j is a connected subset of G , for each $j \in I'$,

$$(P' \cap \alpha_j) \cup (Q' \cap \alpha_j) \neq \alpha_j,$$

and there is a point $z_j \in \alpha_j - (P' \cup Q')$. Furthermore, since P and Q partition V , each z_j lies in some edge. Since the C_j are edge-disjoint each z_j lies in a distinct edge. By compactness there is some subsequence of $\{z_j\}_{j \in I'}$ that converges. Say $\{z_j\}_{j \in J}$ converges to z , for some vertex z . But since P, Q is a partition of V we have that $z \in P \cup Q$ and $z \in P' \cup Q'$. Thus $z_k \in P' \cup Q'$ for some index $k \in J$, and hence $k \in I'$, a contradiction. ■

Finally we have the following result due to Thomassen (private communication).

Proposition 2.12

If G is a graph-like space, then G has only finitely many connected components.

Proof Suppose that G has infinitely many components G_i . Choose any set of points $\{x_i\}$ so that each $x_i \in G_i$. This set has a convergent subsequence converging to some point x . Since G is compact $x \in G$, and $x \in G_i$ for some i . But now for any neighbourhood, N , of x , N contains infinitely many of the points x_i . Thus N intersects infinitely many of the G_i non-trivially. Therefore no neighbourhood of x is connected. This contradicts the local connectedness of G . ■

Since graph-like spaces have only finitely many connected components, we will concern ourselves primarily with connected graph-like spaces. Our arguments will tend to generalize easily to graph-like spaces with more than one connected component.

2.3 A Topological Lemma

In this section we prove a new lemma on decomposing the continuous image of a circle into homeomorphs of circles. In order to motivate this seemingly obscure goal, consider a closed walk in a graph G .

Given any walk, W , of any graph G , W is called a *trail* if every edge of G appears at most once in W . If W is a closed trail, then we can derive from W a set of cycles in G so that every edge is contained in exactly one cycle. We achieve this by a simple recursive procedure. If $W = (v_0, e_1, v_1, e_2, v_2, \dots, v_n, e_{n+1}, v_0)$ then we walk along W starting at v_0 until we reach a vertex v_i that we have already visited. Then we have an index $j < i$ so that $v_j = v_i$. Now the subwalk of W , $(v_j, e_{j+1}, \dots, e_i, v_i)$ is a cycle C_1 in G . Furthermore, when we remove C_1 from W we are left with a closed trail W' in G , so we can repeat this process on W' . Eventually we have a sequence C_1, \dots, C_k of cycles in G so that each edge of W appears in exactly one C_i .

Now consider the graph-like space, X , obtained from a finite graph G as in Section 2.1. Any path in G corresponds to a homeomorph of the closed unit interval in X ; any walk in G corresponds to the continuous image of the closed unit interval in X ; any cycle in G corresponds to a homeomorph of a circle in X ; and, any closed walk in G corresponds to the continuous image of a circle in X . In our graph-like space X , every continuous image of a circle, C , in which every

edge of X appears at most once can be broken into homeomorphs of circles so that each edge of C appears in exactly one circle. This argument does not hold for graph-like spaces in general, as it relies on the fact that G is finite. However, we can expand this recursive procedure to prove an analogous result in more general topological spaces.

We define a *closed curve* to be the continuous image of a circle, and a *simple closed curve* to be a homeomorph of a circle. Given a closed curve τ in a topological space X with finitely many points of “self-intersection” one can convince oneself that τ can be broken into a finite number of simple closed subcurves, homeomorphs of the circle, such that every point of τ that is not a point of self-intersection is covered by exactly one subcurve. In this case there is a natural graph associated with τ and the proof provided above applies. Given a closed curve τ in X with arbitrary points of intersection, we would like to be able to identify a set of simple closed subcurves of τ that cover τ so that every point that is not a point of self-intersection is covered exactly once.

Here we prove that given a “nice” closed curve τ in X such a covering always exists, and we provide a recursive method to construct such a cover. In order for our construction to work we will require that the closure of the set of self-intersections of τ be totally disconnected. Since τ is the continuous image of a circle in the plane we have a continuous surjection $f : C \rightarrow \tau$ that maps a clockwise traversal of the circle to a “clockwise” traversal of τ . For instance if we choose a point $p \in C$, and $q = f(p) \in \tau$ then we can think of f as tracing around the circle from p to p clockwise as we simultaneously trace around τ from q arbitrarily back to q , such that we trace over each point that is not a point of self-intersection exactly once. We will use the function f to construct the desired cover, g , in the following lemma.

First we define a “vertex set” for our curve τ . Given a closed curve τ in topological space X we define the set $W(\tau) \subseteq \tau$ as,

$$W(\tau) := \{p \in \tau : |f^{-1}(p)| \neq 1\}.$$

The set $W(\tau)$ consists of the points $p \in \tau$ that f “visits” more than once in its traversal. Now we define the set $V(\tau) = \text{Cl}(W(\tau))$. We also define C^* to be an arbitrarily large collection of distinct circles.

Lemma 2.13

Suppose that $f : C \rightarrow \tau$ is a continuous surjection from the circle C to a closed curve τ in the topological space X such that $V(\tau)$ is a totally disconnected subset of τ and f traverses τ as described above. Then there exists a continuous surjection $g : C^* \rightarrow \tau$ such that $g|_{C'}$ is a continuous injection for each circle $C' \in C^*$.

We will prove Lemma 2.13 by providing a transfinite recursive procedure that constructs the function g . In order for such a procedure to work we will need the points in $\tau - V(\tau)$ to form a countable set of connected components. The following proposition implies this fact. Also note that in a graph-like space this proposition implies that every closed curve contains at most countably many edges. This connection is interesting despite its vacuity.

Proposition 2.14

Let A be a collection of closed intervals of the circle with unit circumference so that any two elements of A have at most their endpoints in common. Then only countably many elements of A have non-zero length.

Proof Consider $S = \sum_{I \in A} l(I)$ where $l(I)$ is the length of I . Since A consists of intervals in C that overlap only at their endpoints we have that $S \leq 1$. Since our sum is bounded, we can only have finitely many $I \in A$ with $1/2 \leq l(I)$. Further, we can only have finitely many $I \in A$ with $1/4 \leq l(I) < 1/2$. Continuing in this way we have that only finitely many $I \in A$ have $1/2^{i+1} \leq l(I) < 1/2^i$ for each $i \in \mathbb{N}$. Now if $0 < l(I)$ for some $I \in A$, then $1/2^{i+1} \leq l(I) < 1/2^i$ for some $i \in \mathbb{N}$. Thus we have only countably many $I \in A$ such that $0 < l(I)$. ■

2.3.1 Recursive Procedure

Now we present the recursive procedure that we use to construct our function g as in Lemma 2.13. At each step i we construct the following objects. A set D_i of clockwise ordered closed intervals of the original circle C such that:

- D1. If $[a, b] \in D_i$, then $f(a), f(b) \in V(\tau)$.
- D2. If $x \in C$, then $x \in I$ for some $I \in D_i$.
- D3. If $x \in C$, then x is not in more than two intervals of D_i .
- D4. If $x \in C$ with $x \in I_1, I_2$ for $I_1, I_2 \in D_i$, then x is an endpoint of both I_1 and I_2 .

A partition P_i of D_i such that:

P1. For each $A \in P_i$, $B := \cup_{I \in A} I$ is a closed subset of C .

P2. For each $A \in P_i$, $f(B)$ is a closed connected subset of τ .

A set of “reference” functions $\{\rho_j\}_i$ such that each ρ_j is a continuous order-preserving surjection $\rho_j : B_j \rightarrow C_j$ for $C_j \in C^*$. Finally, a continuous surjection $f_i : C^* \rightarrow \tau$ that will always be defined as $f_i(x) := f(\rho_j^{-1}(x))$ where $x \in C_j$.

Recursive Procedure

Initial State

$f_0 := f$ is a continuous surjection $f_0 : C^* \rightarrow \tau$ by assumption.

$V_0 := \{v \in V(\tau) : |f_0^{-1}(v)| \geq 2\}$.

If $V_0 = \emptyset$ then we set $g := f_0$ and $g : C^* \rightarrow \tau$ satisfies Lemma 2.13.

Else we perform the recursion.

Let $v \in V_0$ be an arbitrary point, and $p \in f^{-1}(v)$. Let I denote the clockwise interval consisting of all of C , starting and ending at p . Set $D_0 = \{I\}$, and $P_0 = \{\{I\}\}$. Finally set $\rho : I \rightarrow C$ to be a natural continuous order-preserving surjection that maps the interval I onto C .

Step $\#\beta + 1$ for ordinal β

From step $\#\beta$ we have the following:

- D_β a set of clockwise cyclically ordered subintervals of C satisfying properties D1 through D4.
- P_β a partition of D_β satisfying properties P1 and P2.
- A set $\{\rho_i\}_\beta$ of reference functions, where each $\rho_i : B_i \rightarrow C_i$ is a continuous order-preserving surjection.
- $f_\beta : C^* \rightarrow \tau$ a continuous surjection defined as $f_\beta(x) := f(\rho_i^{-1}(x))$ where $x \in C_i$.
- V_β a non-empty subset of $V(\tau)$, where for each $v \in V_\beta$, $|(f_\beta|_{C'})^{-1}(v)| \geq 2$ for some $C' \in C^*$.

From these objects we construct our continuous surjection $f_{\beta+1}$.

Take $p, q \in (f_\beta|_{C'})^{-1}(v)$ for some $v \in V_\beta$, and some $C' \in C^*$ such that $|(f_\beta|_{C'})^{-1}(v)| \neq 1$.

Then we have $A \in P_\beta$ such that $\rho : B \rightarrow C'$ is a continuous surjection for $B := \cup_{I \in A} I$.

Let $I_1 = [p, q]$ and $I_2 = [q, p]$ be the clockwise intervals of C' between p and q . We will “split” C' into two new circles, C_1 and C_2 as follows.

If $A = \{I\}$, then we set

$$\begin{aligned} D_{\beta+1} &= (D_\beta - I) \cup \{I_1, I_2\}, \quad \text{and} \\ P_{\beta+1} &= (P_\beta - A) \cup \{\{I_1\}, \{I_2\}\}. \end{aligned}$$

Suppose that $|A| > 1$. We will use I_1 and I_2 to construct $D_{\beta+1}$ from D_β . We will do this in two steps, which are essentially the same, one for p and one for q . In the first step we will use p to construct an intermediate set D from D_β , then in the second step we will use q construct $D_{\beta+1}$ from D .

We consider the possibilities for the position of p . We have two cases.

Case #1:

p lies in the interior of an interval I for some $I \in A$.

If p is in the interior of $I = [a, b]$ then we define new clockwise intervals $I'_1 := [a, p]$ and $I'_2 := [p, b]$. We let

$$D := (D_\beta - \{I\}) \cup \{I'_1, I'_2\}.$$

Case #2:

Now suppose p is the endpoint of one or two intervals in A . We again consider cases.

Case #2.1:

We have $I'_1 = [a, p]$ and $I'_2 = [p, b]$ in A , where either a , b or both may be equal to p . (Note that if $[p, p]$ is an interval in D_β , then $[p, p]$ consists of the point p , not the entire circle C' .) In this case we take $D := D_\beta$.

Case #2.2:

We have $I' = [p, p]$ is the only interval containing p . Then we define $I'' := [p, p]$ and $D := D_\beta \cup \{I''\}$.

From p and D_β , we have constructed the intermediate set D . We now repeat the above with q in place of p and D in place of D_β , and we obtain $D_{\beta+1}$ in place of D .

We note that the intervals I_1 and I_2 partition C' into two closed subsets. This partition induces a natural partition of A (plus the new intervals defined above) into new parts A_1 and A_2 . We have that A_1 consists of all of the intervals in A that lie between p and q in clockwise order, plus any new intervals we have introduced. For instance if point p falls into Case #1, $p \in [a, b]$, and point q falls into Case #2, $[a', q] \in A$ and $[q, b'] \in A$, then A_1 consists of the interval $[p, b]$, the intervals of A between b and a' , and the interval $[a', q]$. Similarly, A_2 consists of all of the intervals in A that lie between q and p in clockwise order, plus any new intervals we have introduced. To carry on with our example, A_2 consists of the interval $[q, b']$, all of the intervals of A between the first interval and a , and the interval $[a, p]$. Now we take

$$P_{\beta+1} = (P_\beta - \{A\}) \cup \{A_1, A_2\}.$$

Let $\theta_1 : I_1 \rightarrow C_1$ and $\theta_2 : I_2 \rightarrow C_2$ be the natural order-preserving continuous surjections from the two intervals of C' to the new circles C_1 and C_2 . We define $\rho_1 : B_1 \rightarrow C_1$ and $\rho_2 : B_2 \rightarrow C_2$ as $\rho_1 := \rho \circ \theta_1$ and $\rho_2 := \rho \circ \theta_2$. Now we define

$$\{\rho_i\}_{\beta+1} := (\{\rho_i\}_\beta - \{\rho\}) \cup \{\rho_1, \rho_2\}$$

to be our set of reference functions.

Define $f_{\beta+1} : C^* \rightarrow \tau$ as $f_{\beta+1}(x) := f(\rho_i^{-1}(x))$ where $x \in C_i$.

Finally we define the set

$$V_{\beta+1} := \{v \in V(\tau) : |(f_{\beta+1}|_{C'})^{-1}(v)| \neq 1 \text{ for some } C' \in C^*\}.$$

If $V_{\beta+1} = \emptyset$ then we have $g = f_{\beta+1}$ satisfies Lemma 2.13 and we stop our recursion. Otherwise we continue.

Step # α for limit ordinal α

For each ordinal $\beta < \alpha$ we have:

- D_β a set of clockwise cyclically ordered subintervals of C satisfying properties D1 through D4.

- P_β a partition of D_β satisfying properties P1 and P2.
- A set $\{\rho_i\}_\beta$ of reference functions, where each $\rho_i : B_i \rightarrow C_i$ is a continuous order-preserving surjection.
- $f_\beta : C^* \rightarrow \tau$ a continuous surjection defined as $f_\beta(x) := f(\rho_i^{-1}(x))$ where $x \in C_i$.
- V_β a non-empty subset of $V(\tau)$, where for each $v \in V_\beta$, $|(f_\beta|_{C'})^{-1}(v)| \geq 2$ for some $C' \in C^*$.

We need to construct a continuous surjection $f_\alpha : C^* \rightarrow \tau$.

First we define D_α by considering chains of nested intervals,

$$D_\alpha := \{\cap_{\beta < \alpha} I_\beta : I_\beta \in D_\beta \text{ and } I_\beta \supseteq I_{\beta'} \text{ for all } \beta' \geq \beta\}.$$

Now we define P_α . First consider the P_β for $\beta < \alpha$. Each is a partition of D_β . Equivalently we can view each P_β as the equivalence classes of an equivalence relation \sim_β on D_β . We define the equivalence relation \sim_α on D_α as follows. For $I = \cap_{\beta < \alpha} I_\beta$ and $I' = \cap_{\beta < \alpha} I'_\beta$ in D_α , $I \sim_\alpha I'$ if and only if $I_\beta \sim_\beta I'_\beta$ for all $\beta < \alpha$. Now we define

$$P_\alpha := \{A : A \text{ is a } \sim_\alpha\text{-equivalence class of } D_\alpha\}.$$

We have that P_α partitions D_α automatically.

Now we will construct our set $\{\rho_i\}_\alpha$ of reference functions, where each $\rho_i : B_i \rightarrow C_i$ is a continuous order-preserving surjection. In order to accomplish this we consider three cases.

Case #1:

For $A \in P_\alpha$, $A \in P_\beta$ for some $\beta < \alpha$.

This is the simplest case. We have a continuous order-preserving surjection $\rho : B \rightarrow C'$ in $\{\rho_i\}_\beta$. We simply re-use this reference function at step α .

For the next two cases we consider the subset $A' \subseteq A$ defined as

$$A' := \{[a, b] \in A : a \neq b\}.$$

Case #2:

For $A \in P_\alpha$, $A \notin P_\beta$ for any $\beta < \alpha$, and $|A'|$ is finite.

If A consists entirely of singleton intervals (i.e. intervals of the form $[x, x]$), then $f(B) \in V(\tau)$. This follows since $f(B)$ is connected, and $f(x) \in V(\tau)$ for each $x \in B$. Since $V(\tau)$ is totally disconnected we must have that $f(B) \in V(\tau)$. In this case the intervals in A are of no further interest to our recursive process, nor to the construction of our function g . Thus we are free to do with them what we will. We simply define $\rho : B \rightarrow C'$ as $\rho(x) = C'$ for all $x \in B$.

If $|A'| > 0$ then we consider the set A' and set $B' = \cup_{I \in A'} I$. Now we can construct a continuous surjection $\rho' : B' \rightarrow C'$ as follows. Suppose $|A'| = n$. Then we map each $I \in A'$ to a clockwise interval of length $1/n$ on C' so that the cyclic order of A' is preserved, and consecutive intervals overlap in exactly one point. We define $\rho : B \rightarrow C'$ as $\rho(x) = \rho'(x)$ if $x \in B'$ and if $x \in B$ with $[x, x] \in A$ we let $\rho(x) = \rho'(y)$ where $[a, y] \in A'$ is the first non-singleton clockwise predecessor of $[x, x]$ in A .

Case #3:

For $A \in P_\alpha$, $A \notin P_\beta$ for any $\beta < \alpha$, and $|A'|$ is not finite.

We have by Proposition 2.14 that A' is countable. Thus we can take an arbitrary enumeration of the elements of A' , and further we can require that I_1, I_2 , the first two intervals, are non-consecutive. Now we map $\rho' : B' \rightarrow C'$ continuously where the circumference of C' is 1. We accomplish this as follows. Take the interval $I_1 \in A'$ and map I_1 to any interval of C' of length $1/2$. Take the interval $I_2 \in A' - \{I_1\}$, recall that I_2 is not consecutive with I_1 , and map I_2 to the interval of C' with length $1/4$ that is diametrically opposite $\rho'(I_1)$. Now we have two intervals $J_{1,2}$ and $J_{2,1}$ the intervals of C' between $\rho'(I_1)$ and $\rho'(I_2)$, each with length $1/8$. We now continue to map the remaining intervals in A' to intervals of C' as follows.

Consider the interval I_3 . There are three possibilities for I_3 in the cyclic order on C : it is adjacent to neither I_1 nor I_2 ; it is adjacent to one of I_1 and I_2 ; it is adjacent to both I_1 and I_2 . If I_3 is adjacent to neither I_1 nor I_2 , then either: I_3 lies between I_1 and I_2 , in which case we map I_3 to the interval of length $1/16$ centered in $J_{1,2}$; or I_3 lies between I_2 and I_1 , in which case we map I_3 to the interval of length $1/16$ centered in $J_{2,1}$. If I_3 is adjacent to I_1 and lies between I_1 and I_2 then we map I_3 to the interval of length $1/16$ in $J_{1,2}$ that overlaps $\rho'(I_1)$ in exactly one point. We map I_3 analogously if I_3 lies between I_2 and I_1 or if I_3 is adjacent to I_2 . Finally if I_3 is adjacent to both I_1 and I_2 then we map I_3 to either $J_{1,2}$ or to $J_{2,1}$ as is appropriate. We simply repeat this process with each

new interval I_i until we have mapped all of B' onto C' . Note that ρ' may not be surjective; however, $C' - \rho'(B')$ is a totally disconnected set of points.

Now we construct $\rho : B \rightarrow C'$ from $\rho' : B' \rightarrow C'$. For any $x \in B'$ we set $\rho(x) = \rho'(x)$. We want to define $\rho(x)$ for $x \in B - B'$. Consider such an x , and note that if x has a predecessor $I \in A'$ or a successor $J \in A'$ then we map $\rho(x)$ to the appropriate endpoint of I or J . Now suppose that x has no such successor or predecessor. Note that we have a part $Z \in P_\alpha$ with $A \neq Z$. We take arbitrary $y \in Z$. Now for any $x \in B - B'$ we have that x and y partition A' into two parts, A'_l and A'_r , with respect to the clockwise intervals $[x, y]$ and $[y, x]$. Let $B'_l = \cup_{I \in A'_l} I$ and $B'_r = \cup_{I \in A'_r} I$ be defined as usual. Now let $\rho(x) = w$ if w is the unique point of C' that separates $\rho'(B'_l)$ from $\rho'(B'_r)$. If there are two such points, w_1 and w_2 then we let $\rho(x) = w_1$ where w_1 is the accumulation point of $\{\rho'(I_i)\}$ where $\{I_i\}$ is any sequence of intervals in A'_l converging to x (equivalently we could have chosen a sequence $\{I_i\}$ of intervals in A'_r converging to x).

Thus we have a set $\{\rho_i\}_\alpha$ of reference functions, with each $\rho_i : B_i \rightarrow C_i$ a continuous order-preserving surjection. Now we construct $f_\alpha : C^* \rightarrow \tau$, a continuous surjection. As always we define $f_\alpha(x) := f(\rho_i^{-1}(x))$ where $x \in C_i$.

Finally, we define the set V_α . Let C^{**} be the collection of circles in C^* that correspond to classes A with $|A'| \neq 0$. Define

$$V_\alpha := \{v \in V(\tau) : |(f_\alpha|_{C'})^{-1}(v)| \geq 2 \text{ for some } C' \in C^{**}\}.$$

If $V_\alpha = \emptyset$ then we have $g = f_\alpha$ satisfies Lemma 2.13 and we stop our recursion. Otherwise we continue.

This concludes the description of the recursive construction. We spend the balance of this section proving Lemma 2.13. We will accomplish this by proving a succession of claims demonstrating that the objects created by the above recursion have the properties listed. Once we have proven these claims we prove the lemma by demonstrating that our process terminates.

Before we begin the proof we note two properties of our construction. First, at each stage we constructed a set D_i from the sets D_j where $j < i$. Notice that if $j < i$ then D_i is a refinement of D_j in the sense that each interval of D_i is contained in an interval of D_j . Also note that at each step we were careful to construct our reference functions in such a way that they preserve the order of traversal of f . This is of no consequence to our main lemma, and we could have

constructed these functions so that they do not preserve the order of f . However, when we consider the face boundaries of an embedded graph-like space we will use our topological lemma in order to break a traversal into sub-traversals, and order-preservation will be useful.

2.3.2 Proof of Lemma 2.13

Now we provide the proof of Lemma 2.13. The proof is by induction, however since there are many objects under consideration, and many properties to prove, we break the proof into parts. We will prove eight claims that demonstrate that our construction really does work as described, and that the objects D_i , P_i , $\{\rho\}_i$ and f_i have the properties required. Then we will prove that the recursive procedure terminates and returns the desired function.

The first four claims prove that our construction is sound at step $\#\beta + 1$, for ordinal β .

Claim 2.15

$D_{\beta+1}$ satisfies properties D1 through D4.

Proof Recall the construction of $D_{\beta+1}$. We have that $D_{\beta+1}$ satisfies D1 by construction. If we have $x \in C$ then x is covered by D_β and hence also by $D_{\beta+1}$, so $D_{\beta+1}$ satisfies D2. Further for $x \in C$, x is in no more than two intervals of D_β . Thus x is in no more than two intervals of $D_{\beta+1}$ and furthermore x must be an endpoint of those intervals. Therefore $D_{\beta+1}$ satisfies properties D1 through D4. ■

Claim 2.16

$P_{\beta+1}$ satisfies properties P1 and P2.

Proof Recall the construction of $P_{\beta+1}$. We have that $P_{\beta+1}$ clearly partitions $D_{\beta+1}$.

Note that for $A \in P_{\beta+1} \cap P_\beta$, A satisfies P1 and P2 by assumption. Thus we only need consider A_1 and A_2 . We have from step $\#\beta$ that $\rho : B \rightarrow C'$ is a continuous order-preserving surjection. Further I_1 and I_2 are closed subsets of C' . Therefore $\rho^{-1}(I_1)$ and $\rho^{-1}(I_2)$ are closed subsets of C . But $\rho^{-1}(I_1) = B_1$ and $\rho^{-1}(I_2) = B_2$, so B_1 and B_2 are closed subsets of C and $P_{\beta+1}$ satisfies P1.

Now since $\rho(B_1)$ and $\rho(B_2)$ are connected subsets of C' , and $f_\beta|_{C'}$ is continuous, $f_\beta|_{C'}(\rho(B_1))$ and $f_\beta|_{C'}(\rho(B_2))$ are connected subsets of τ . But

$$f_\beta|_{C'}(\rho(B_1)) = f(\rho^{-1}(\rho(B_1))) = f(B_1), \quad \text{and}$$

$$f_\beta|_{C'}(\rho(B_2)) = f(\rho^{-1}(\rho(B_2))) = f(B_2)$$

so $f(B_1)$ and $f(B_2)$ are connected subsets of τ . Note that both C and τ are compact Hausdorff spaces and both B_1 and B_2 are closed. We have by Theorem 26.2 in [10] that every closed subspace of a compact space is compact. Thus each B_i is a closed subspace of C and hence compact. Since f is continuous, Theorem 26.5 in [10] implies that $f(B_i)$ is a compact subspace of τ . Finally by Theorem 26.3 in [10] $f(B_i)$ is a compact subspace of a Hausdorff space, and hence is a closed subspace of τ . Therefore $f(B_1)$ and $f(B_2)$ are closed subsets of τ and $P_{\beta+1}$ satisfies P2. ■

Claim 2.17

The functions $\{\rho_i\}_{\beta+1}$ are continuous order-preserving surjections $\rho_i : B_i \rightarrow C_i$.

Proof Recall that we defined $\theta_1 : I_1 \rightarrow C_1$ and $\theta_2 : I_2 \rightarrow C_2$ to be the natural continuous order-preserving surjections from the two intervals of C' to the new circles C_1 and C_2 . We defined $\rho_1 : B_1 \rightarrow C_1$ and $\rho_2 : B_2 \rightarrow C_2$ as $\rho_1 := \rho \circ \theta_1$ and $\rho_2 := \rho \circ \theta_2$. Thus, since the θ_i and ρ are continuous functions, the ρ_i are continuous order-preserving surjections. ■

Claim 2.18

$f_{\beta+1} : C^* \rightarrow \tau$ is a continuous surjection.

Proof Note that $f_{\beta+1}$ is a surjection since f is a surjection and $D_{\beta+1}$ covers C . In order to prove continuity we consider cyclically monotonic sequences $\{a_i\}$ consisting of points on C_j . By assumption if C_j is a circle other than C_1 or C_2 , $f_{\beta+1}|_{C_j}$ is continuous. Thus we only consider C_1 and C_2 .

Suppose we have a sequence $\{a_i\}$ of cyclically monotonic points converging to a in C_j for $j \in \{1, 2\}$. We have two cases.

Case #1:

$$\{a_i\} \rightarrow a \text{ where } |\theta_j^{-1}(a)| = 1.$$

Then there is a finite index l such that $|\theta_j^{-1}(a_i)| = 1$ for all $i \geq l$. Since θ_j is a continuous bijection on $I_j - \{p, q\}$, θ_j^{-1} is a continuous bijection on $C_j - \{\theta_j(\{p, q\})\}$. Thus $\lim \theta^{-1}(a_i) = \theta^{-1}(a)$. Now since by assumption f_β is continuous we have $\lim f_\beta(\theta_j^{-1}(a_i)) = f_\beta(\theta_j^{-1}(a))$. Since $\rho_j := \rho \circ \theta_j$ we have,

$$\begin{aligned} f_\beta(\theta_j^{-1}(a_i)) &= f(\rho^{-1}(\theta_j^{-1}(a_i))) = f(\rho_j^{-1}(a_i)) = f_{\beta+1}(a_i) \quad \text{and,} \\ f_\beta(\theta_j^{-1}(a)) &= f(\rho^{-1}(\theta_j^{-1}(a))) = f(\rho_j^{-1}(a)) = f_{\beta+1}(a). \end{aligned}$$

Thus $\lim f_{\beta+1}(a_i) = f_{\beta+1}(a)$, as required.

Case #2:

$\{a_i\} \rightarrow a$ where $|\theta_j^{-1}(a)| \neq 1$.

We have $\theta_j^{-1}(a) = \{p, q\}$. Now $\{\theta_j^{-1}(a_i)\}$ is a cyclically monotonic sequence in C' bounded by $\theta_j^{-1}(a)$, so it converges to some $b \in C'$. Since θ_j is a continuous function, $\{\theta_j(\theta_j^{-1}(a_i))\} \rightarrow \theta_j(b)$, or $\{a_i\} \rightarrow \theta_j(b)$. Thus $\theta_j(b) = a$, and $\{\theta_j^{-1}(a_i)\} \rightarrow b \in \theta_j^{-1}(a)$. By assumption, f_β is continuous, so $\{f_\beta(\theta_j^{-1}(a_i))\} \rightarrow f_\beta(b)$. But $f_\beta(\theta_j^{-1}(a_i)) = f_{\beta+1}(a_i)$ and $f_\beta(b)$ is constant for all $b \in \theta_j^{-1}(a)$. Thus $f_\beta(b) = f_{\beta+1}(a)$. Therefore we have $\{f_{\beta+1}(a_i)\} \rightarrow f_{\beta+1}(a)$, as required. ■

The final four claims prove that our construction is sound at step # α , for limit ordinal α .

Claim 2.19

D_α satisfies properties D1 through D4.

Proof D_α covers C , since if $x \in C$ then $x \in I_\beta$ for some $I_\beta \in D_\beta$ for each $\beta < \alpha$. If x is in the interior of I_β then we have some $I_{\beta+1}$ so that $I_\beta \supseteq I_{\beta+1}$ by construction. If x is a boundary point of I_β then we may have two choices for $I_{\beta+1}$. In this case we choose $I_{\beta+1}$ so that $I_{\beta+1}$ is a subinterval of I_β . Thus we have a chain of nested intervals each containing x . Therefore $x \in \bigcap_{\beta < \alpha} I_\beta \in D_\alpha$ and D_α covers C .

For each $x \in C$ we also know that D_α covers x at most twice. Our justification is similar to the previous argument. At each β we have at most two intervals containing x , and once we have made the first choice for I_β the rest of the chain is determined. Thus there are at most two chains whose intersection contains x . This proves that D_α satisfies D2 and D3.

Now suppose we have $x \in C$ such that $x \in I, I'$ for $I, I' \in D_\alpha$. Then there are two valid chains containing x , and x is a boundary point of I_β and I'_β for some $\beta < \alpha$. Then, by assumption, x is a boundary point of I_β and I'_β . Thus x is a boundary point of I and I' and D_α satisfies D4.

Finally if $I = [a, b] \in D_\alpha$ we want to show that $f(a), f(b) \in V(\tau)$. If a or b is an endpoint of I_β for any interval in the chain then $f(a)$ or $f(b)$ is in $V(\tau)$ by assumption. Otherwise we have sequences $\{a_i\}$ and $\{b_i\}$ of the left and right endpoints, respectively, of the I_i and $\{a_i\} \rightarrow a$, $\{b_i\} \rightarrow b$. But since f is a continuous function this means we have $\{f(a_i)\} \rightarrow f(a)$ and $\{f(b_i)\} \rightarrow f(b)$ in τ . Now by assumption $f(a_i), f(b_i) \in V(\tau)$, and $V(\tau)$ is a closed subset of τ . Thus $f(a), f(b) \in V(\tau)$ as required, and D_α satisfies D1 through D4. ■

Claim 2.20

P_α satisfies properties P1 and P2.

Proof Let $A \in P_\alpha$, and let $B = \cup_{I \in A} I$. Suppose $I \in A$. Then, $I = \cap_{\beta < \alpha} I_\beta$ for intervals $I_\beta \in A_\beta$. Let $B_\beta = \cup_{J \in A_\beta} J$. Now $I \subseteq I_\beta \subseteq B_\beta$ for each $\beta < \alpha$. Thus $B \subseteq \cap_{\beta < \alpha} B_\beta$. Furthermore, if $x \in \cap_{\beta < \alpha} B_\beta$, then $x \in B_\beta$ for each $\beta < \alpha$. If $x \in J_\beta$ for each $\beta < \alpha$, then $x \in \cap_{\beta < \alpha} J_\beta = J$. Note that $J_\beta \sim_\beta I_\beta$ for each $\beta < \alpha$, so $J \sim_\alpha I$ and $J \in A$. Hence, $x \in B$ and $B = \cap_{\beta < \alpha} B_\beta$. Thus B is the intersection of the closed sets B_β and therefore is closed. Therefore P_α satisfies P1.

To prove that $f(B)$ is closed we use the same proof as before. Namely, B is closed, so B is compact, so $f(B)$ is compact, so $f(B)$ is closed.

For $f(B)$ connected, we will employ Proposition 2.1. Since for $\beta' \geq \beta$, $B_{\beta'} \subseteq B_\beta$, we have that $f(B_{\beta'}) \supseteq f(B_\beta)$. Thus the sets in the collection $\{f(B_\beta) : \beta < \alpha\}$ are non-empty closed and connected, so by Proposition 2.1, their intersection is connected.

Since f is a function,

$$f(B) = f(\cap_{\beta < \alpha} B_\beta) \subseteq \cap_{\beta < \alpha} f(B_\beta).$$

Suppose $x \in \cap_{\beta < \alpha} f(B_\beta)$. Then, for each $\beta < \alpha$, $x \in f(B_\beta)$ and there is some $y_\beta \in B_\beta$ such that $f(y_\beta) = x$. Let $\{y_i : i \in \Delta\}$ be a convergent subnet of $\{y_\beta : \beta < \alpha\}$, with limit y . Then since f is continuous, $f(y) = x$. Furthermore, since the sets B_β are nested, $y \in B_\beta$ for each $\beta < \alpha$, and $y \in B$. Thus $x \in \cap_{\beta < \alpha} f(B_\beta)$ and $f(B) = \cap_{\beta < \alpha} f(B_\beta)$. Therefore $f(B)$ is a connected subset of τ , and P_α satisfies properties P1 and P2. ■

Claim 2.21

The functions $\{\rho_i\}_\alpha$ are order-preserving continuous surjections $\rho_i : B_i \rightarrow C_i$.

Proof Recall that in constructing $\rho : B \rightarrow C'$ we had to consider three cases.

Case #1:

In this case we re-used a reference function ρ from a preceding step, so there is nothing to prove.

Case #2:

In this case A' is a finite set in P_α .

Since $A - A'$ consists of point intervals, each $[p, p] \in A - A'$ is mapped to a point of $V(\tau)$ by f . Thus $A - A'$ is a totally disconnected set of points in C . Also

ρ maps the points of B to the “right” places in the sense that if $\rho(x) = \rho(y)$ then $f(x) = f(y)$. We constructed ρ to be the natural order-preserving continuous surjection as in the non-limit case, and again there is nothing to prove.

Case #3:

In this case $|A'|$ is not finite. Recall the construction of ρ . In constructing ρ we constructed a continuous function $\rho' : B' \rightarrow C'$. First we show that $C' - \rho'(B')$ is totally disconnected.

Consider the construction of ρ' at step i . At step i we have mapped i of the intervals of A' onto C' . Let $A_i = \rho'(I_1 \cup \dots \cup I_i) \subset C'$ be the union of the images of the first i intervals of A' . We refer to the intervals of C' in $C' - A_i$ as *unassigned* intervals. At step $i + 1$ we either: remove an unassigned interval; reduce the size of an unassigned interval by half; or break an unassigned interval into two unassigned intervals, each of which has length $1/4$ of the original. Furthermore, if we have an unassigned interval at step i then by construction we have some interval of A' that will be mapped in between its endpoints.

Form a rooted tree T consisting of unassigned intervals as follows. The first unassigned interval is C' , which becomes the root vertex. An unassigned interval in $C' - A_i$ is adjacent to the unassigned interval in $C' - A_{i-1}$ that contains it. Let P be a maximal path in T starting at C' . If P is finite, then the last vertex of P corresponds to an interval I of C' in $C' - A_i$ for some i , and we map the interval I_{i+1} onto I . Otherwise $P = (C', J_1, J_2, \dots)$ is a ray and, for infinitely many indices i , J_{i+1} is half the length of J_i . Thus, $\bigcap_{i \geq 1} J_i$ is at most a single point. If $x \in C' - \rho'(B')$, then, for every i , x is in some interval J_i of $C' - A_i$ and, evidently, $\{x\} = \bigcap_{i \geq 1} J_i$. Now suppose x, x' are distinct points of $C' - \rho'(B')$, with corresponding paths $P_x = (C', J_1, J_2, \dots)$ and $P_{x'} = (C', J'_1, J'_2, \dots)$ in T , so that $\{x\} = \bigcap_{i \geq 1} J_i$ and $\{x'\} = \bigcap_{i \geq 1} J'_i$. Then there exists some i so that $J_i \neq J'_i$ (as otherwise $x = x'$). Thus, x and x' are in different intervals in $C' - A_i$. We can construct sets U and V by arbitrarily partitioning $C' - A_i$ into two sets of intervals and letting U and V be the union over each part respectively. Since x, x' are in distinct intervals, we can choose U and V so $x \in U$ and $x' \in V$. Now set

$$\begin{aligned} U' &= U \cap (C' - \rho'(B')), \quad \text{and} \\ V' &= V \cap (C' - \rho'(B')). \end{aligned}$$

We have that U', V' is a separation of $C' - \rho'(B')$ with $x \in U'$ and $x' \in V'$, thus

$C' - \rho'(B')$ is totally disconnected.

Now recall the construction of ρ from ρ' . Note that ρ is a surjection, since for $x \in C'$ either $x \in \rho'(B')$, or x is an accumulation point of intervals $\{\rho'(I_i)\}$, where the intervals are cyclically monotonic. Then $\{I_i\}$ is a cyclically monotonic sequence of intervals in C bounded by I_1 or I_2 and hence converges to some point $y \in C$. But since B is a closed subset of C we must have that $y \in B$. Now by construction $\rho(y) = x$. This also demonstrates that ρ is continuous (as usual consider a sequence $\{a_i\} \rightarrow a$ in C , now $\{\rho(a_i)\} \rightarrow \rho(a)$ follows). ■

Before we can prove that f_α is a continuous surjection we need two propositions.

Proposition 2.22

For any $A \in P_\alpha$ if $[a, b] \in A$ is the predecessor of $[c, d] \in A$ in the cyclic order, then $f(b) = f(c)$.

Proof We proceed by induction. First we note that this statement follows trivially in the non-limit ordinal step. We assume that the Proposition holds for all $\beta < \alpha$ and that $A \in P_\alpha$ is not in P_β for any $\beta < \alpha$.

Assume that $[a, b] = \cap_{\beta < \alpha} I_\beta$ and $[c, d] = \cap_{\beta < \alpha} I'_\beta$. Then $I_\beta \sim_\beta I'_\beta$ for all $\beta < \alpha$. Further we have some $\beta' < \alpha$ such that I_β and I'_β are consecutive in the cyclic order of A_β for all $\beta' \leq \beta < \alpha$ (this is true as otherwise we have a chain of intervals that lie between I_β and I'_β for all $\beta < \alpha$ and hence an interval between $[a, b]$ and $[c, d]$). Now the right endpoints of the I_β and the left endpoints of the I'_β for $\beta' \leq \beta < \alpha$ form two sequences converging to b and c respectively. But by assumption these sequences are mapped to equal sequences in $V(\tau)$ by f . Thus since f is continuous $f(b) = f(c)$ as required. ■

Note that Proposition 2.22 also applies if $|A| = 1$. In that case there is only one interval $[a, b] \in A$ and $f(a) = f(b)$.

Proposition 2.23

Given $[a, b], [c, d] \in A$ for some $A \in P_\alpha$ such that there are no non-singleton intervals between them in the cyclic order on C in A , $f(b) = f(c)$.

Proof If there are no intervals between $[a, b]$ and $[c, d]$ then the result follows from Proposition 2.22. Similarly, if there are finitely many singleton intervals between $[a, b]$ and $[c, d]$ then the result follows by repeatedly applying Proposition 2.22 to the intervals between b and c .

Suppose that $x \in (b, c)$ and $x \notin B$. Since B is closed, $C - B$ is a collection of open intervals. Let (x_1, x_2) be the open interval component of $C - B$ that contains x . Then $x_1, x_2 \in B$, and x_1, x_2 are both endpoints of intervals in A . Thus by Proposition 2.22 $f(x_1) = f(x_2)$, as the intervals corresponding to x_1 and x_2 are consecutive.

Now we define a new function $h : [b, c] \rightarrow \tau$. For $x \in B$, let $h(x) = f(x)$. In order to define h on $[b, c] - B$ note that for all $x \in [b, c] - B$, x is in an open interval (x_1, x_2) as described above. Furthermore the ends of this interval are mapped to the same point in $V(\tau)$ by f , and we set $h(x) = f(x_1) = f(x_2)$.

We claim that h is continuous. Consider a subset $Q \subseteq [b, c]$. From the definition of h we have a subset of B , B' , such that $h(Q) = h(B') = f(B')$. Suppose that $z \in h(\text{Cl}(Q))$. Then there is some $x \in \text{Cl}(Q)$ such that $h(x) = z$. If $x \in Q$ then $z \in h(B')$, and $z \in h(\text{Cl}(B'))$. If $x \notin Q$, then there is a sequence $\{y_i\}$ of points in Q with limit x . We have two cases, either $x \notin B$, or $x \in B$. If $x \notin B$, then x lies in an open interval (x_1, x_2) of $C - B$ as described above. Furthermore, there is some index i so that $y_j \in (x_1, x_2)$ for all $j > i$. However, for all points $y \in (x_1, x_2)$, $h(y) = h(x_1) = h(x_2)$. Since each $y_i \in Q$, either $x_1 \in B'$ or $x_2 \in B'$, so $z \in h(B')$ and $z \in h(\text{Cl}(B'))$.

Suppose that $x \in B$. The sequence of points $\{y_i\}$ is contained in Q , so for each y_i there is some $y'_i \in B'$ corresponding to the definition of h so that $h(y_i) = h(y'_i)$. Since B is closed, any point in the closure of $\{y'_i\}$ lies in B . Assume that y' is a limit point of the sequence $\{y'_i\}$. Then $y' \in \text{Cl}(B')$. We claim that there is a limit point y' so that $h(y') = h(x) = z$, and thus $z \in h(\text{Cl}(B'))$. There are two possibilities. Either there is some component I of $C - B$ and some index i so that $y_j \in I$ for all $j > i$, or there is a sequence of components $\{I_l\}$ of $C - B$ so that each I_l contains some but not all of the points y_i . In the first case x is an endpoint of I . However, by definition y' is also an endpoint of I , and since the endpoints of I are mapped to the same point by h , $h(y') = h(x) = z$ as required. In the second case the points y'_i are endpoints of the intervals I_l . Either we have a subsequence of intervals $\{I_l\}_{l \in L}$ so that x is the limit of the intervals I_l for $l \in L$, or x is the endpoint of infinitely many of the intervals I_l . In the first case the sequence of points $\{y'_i\}_{i \in I}$ that correspond to the endpoints of intervals I_l for $l \in L$ converge to x , so we can take $y' = x$ and y' satisfies our requirements. Otherwise x is the endpoint of infinitely many intervals, and we apply the argument from the first case.

We have that if $z \in h(\text{Cl}(Q))$, then $z \in h(\text{Cl}(B'))$. Thus $h(\text{Cl}(Q)) \subseteq h(\text{Cl}(B'))$. Now, since f is continuous, $f(\text{Cl}(X)) \subseteq \text{Cl}(f(X))$ for any $X \subseteq C$. Therefore

$$h(\text{Cl}(Q)) \subseteq h(\text{Cl}(B')) = f(\text{Cl}(B')) \subseteq \text{Cl}(f(B')) = \text{Cl}(h(Q)),$$

and h is continuous.

Now, since $[b, c]$ is connected and h is continuous, $h([b, c])$ is connected. Note that by definition $h(a) \in V(\tau)$ for all $a \in [b, c]$. But $V(\tau)$ is a totally disconnected set, thus $h(a) = h(a')$ for all $a, a' \in [b, c]$. Therefore $f(a) = f(a')$ for all $a, a' \in [b, c] \cap B$. Thus $f(b) = f(c)$, as required. ■

Now we prove that f_α is a continuous surjection.

Claim 2.24

$f_\alpha : C^* \rightarrow \tau$ is a continuous surjection.

Proof Suppose that $C_j \in C^*$ so that there is a function $\rho_j \in \{\rho_i\}_\alpha$ with $\rho_j : B_j \rightarrow C_j$. Consider a cyclically monotonic sequence $\{a_i\}$ of points in C_j converging to a . We show that $\{f_\alpha(a_i)\}$ converges to $f_\alpha(a)$. We have three cases, corresponding to the three cases in the construction of the ρ_j .

Case #1:

$\rho_j \in \{\rho_i\}_\beta$ for some $\beta < \alpha$. In this case $f_\alpha|_{C_j} = f_\beta|_{C_j}$ and hence $\{f_\alpha(a_i)\}$ converges to $f_\alpha(a)$ by assumption.

Case #2:

A'_j is a finite set. Then we have a finite index l such that $a_i \in \rho_j(I)$ for some $I \in A_j$ for all $i \geq l$. Thus Proposition 2.23 and the continuity of f give us that $\{f_\alpha(a_i)\} = \{f(\rho_j^{-1}(a_i))\}$ converges to $f_\alpha(a)$.

Case #3:

A'_j is an infinite set. Then we have $\rho_j : B_j \rightarrow C_j$, a continuous surjection. For each i , let I_i be the interval in A_j containing a_i . Since $\{a_i\}$ is cyclically monotonic, and ρ_j is order preserving, the sequence $\{I_i\}$ is cyclically monotonic and bounded above by an interval I in A_j containing a . Thus the sequence formed by taking the terms of $\{\rho_j^{-1}(a_i)\}$ in their cyclic order is a cyclically monotonic convergent sequence in C converging to b . But then, since ρ_j is continuous, $\{a_i\}$ converges to $\rho_j(b)$. Therefore $a = \rho_j(b)$. Now since f is continuous, $\{f(\rho_j^{-1}(a_i))\}$ converges to $f(b)$. As $b \in \rho_j^{-1}(a)$, Proposition 2.23 implies that f is constant on $\rho_j^{-1}(a)$ for all $a \in C$. Therefore $\{f(\rho_j^{-1}(a_i))\}$ converges to $f(\rho_j^{-1}(a))$, so $\{f_\alpha(a_i)\}$ converges to $f_\alpha(a)$ as required. ■

Thus we have shown that our recursive procedure is well-defined and produces the desired objects. We conclude the proof of Lemma 2.13 by showing that the recursive procedure terminates.

Proof We want to demonstrate that our recursive procedure terminates before some ordinal γ . We begin by taking another look at our procedure. Notice that we can view the entire process as “removing” points on τ that are visited multiple times by f . Say for example that we have $p \in \tau$ such that p is visited exactly twice by f . Then $p \in V(\tau)$ and if $f^{-1}(p) = \{a, b\}$ then at step #1 we may choose to break C at points a and b . The result is that $f_1(C^*)$ has one less self-intersection at p . But this is not always the case. We are not guaranteed to reduce the number of intersections by exactly one at each step.

For each β , each $p \in \tau$ and each $C_i \in C^*$, we let

$$X(\beta, p, C_i) = \{x \in C_j : f_\beta(x) = p\}.$$

Now the set of points of interest on C_i at step β is

$$X(\beta, C_i) = \{x \in X(\beta, p, C_i) : |X(\beta, p, C_i)| \geq 2\}.$$

Finally, set $X(\beta) = \cup_{C_j \in C^*} X(\beta, C_j)$. By construction, $X(\beta + 1)$ is a proper subset of $X(\beta)$, since we remove at least one point of multiple-intersection at each non-limit step of the procedure. We also have that if α is a limit ordinal, then $X(\alpha) \subseteq X(\beta)$ for all ordinals $\beta < \alpha$.

Therefore we cannot iterate as often as γ times, where γ is the smallest ordinal with cardinality $> |X(0)|$. Thus, since our recursive procedure is sound (see preceding claims) the function g exists and Lemma 2.13 holds. ■

We will be able to apply Lemma 2.13 to graph-like spaces in order to prove Theorem 4.13.

Chapter 3

Embeddings and Face Boundaries

The main results of this thesis concern embeddings of graph-like spaces in the plane. In order to prove these results we will need to explore some of the properties of embeddings of graph-like spaces. The arguments presented in this chapter will be applied in subsequent chapters to the planar case. However, these results also hold in arbitrary surfaces.

A *surface* is a compact 2-manifold with no boundary. A *2-manifold* is a Hausdorff space X with a countable basis such that every point $x \in X$ has a neighbourhood that is homeomorphic with an open subset of \mathbb{R}^2 .

The sphere, \mathbb{S}^2 , and the plane will be the only spaces in which we will construct embeddings of graph-like spaces. We will colloquially refer to the plane as a surface, even though it is not compact. When considering embeddings, “sphere” and “plane” are, for the most part, interchangeable. We will use both, and favour one over the other simply for ease of explanations. Even though we will not construct embeddings in more general surfaces, the results in this chapter all hold for arbitrary surfaces.

An *embedding* is an injective continuous map $\phi : X \rightarrow Y$ so that if $Z = \phi(X)$ is the image of X in Y , then the function $\phi' : X \rightarrow Z$ formed by restricting the range of ϕ is a homeomorphism. In other words, an embedding of a space X in the space Y is a subspace Z of Y with the same topology as X . In this chapter we will consider an embedding of a graph-like space G in a surface Σ .

In Section 3.1 we give a brief introduction to embeddings, and define the objects of interest. In Section 3.2 we will consider the edges of an embedded graph-like space. We prove that the edges of such a space have neighbourhoods

that are homeomorphic to open disks. In Section 3.3 we will consider the faces of an embedded graph-like space. We will use the neighbourhoods discussed in Section 3.2 to construct a continuous surjection from the circle to the boundary of any face. Finally Section 3.4 contains a brief discussion of how these results apply to embeddings of disconnected graph-like spaces.

3.1 Topological Properties of Embeddings

Consider an embedding of a graph-like space G in a surface Σ . We refer to the subspace of Σ homeomorphic to G as K . Note that K is also a graph-like space, with its own set of vertices and edges. We also have an additional set of objects associated with K and Σ , the faces of K .

Since K is a subset of Σ it is natural to consider the set $\Sigma - K$. We refer to the connected components of $\Sigma - K$ as the *faces* of K . Since K is compact, and Σ is Hausdorff, K is a closed subset of Σ , so $\Sigma - K$ is an open subset of Σ . Since 2-manifolds are locally connected, each connected component of $\Sigma - K$ is open. We will also be interested in finding arcs, both in K and in the individual faces of K . In this regard, the following proposition will be useful.

Proposition 3.1

If X is locally arcwise connected then every connected open set in X is arcwise connected.

Proof Suppose that C is a connected open set in X . Consider $x \in C$ and let A be the arcwise connected component of C containing x . Now by definition, if U is a neighbourhood of x , then we have a neighbourhood V_x of x such that $V_x \subseteq U$ and $V_x \subseteq A$. Thus $A = \cup_{x \in A} V_x$ and A is an open subset of C . This holds for every arcwise connected component of C . If \mathcal{A} is the collection of arcwise connected components of C , then A and $\mathcal{A} - A$ form a separation of C , unless $C = A$. Since C is connected we must have that $C = A$. Therefore C is arcwise connected. ■

Since surfaces are by definition locally arcwise connected, each face of K is arcwise connected.

Furthermore, we can require that the faces of K be simply connected. Suppose that F is a face of K and σ is a non-contractible simple closed curve in F . If $\Sigma - \sigma$ contains more than one connected component whose intersection with

K is non-empty, then K is not connected. For now we will assume that K is a connected graph-like space. Our results will extend easily to graph-like spaces that are not connected.

Assume that K is connected, and F is a face of K so that there is a simple closed curve $\sigma \subset F$, and σ is a non-contractible curve in Σ . If $\Sigma - \sigma$ is not connected, then there is one connected component Σ' of $\Sigma - \sigma$ so that $K \subset \Sigma'$. Let σ' be the boundary of Σ' . We can construct a new surface Σ'' by identifying σ' with the boundary of a closed disk. The surface Σ'' contains K and has smaller genus than Σ .

Suppose $\Sigma - \sigma$ is connected. Then we have non-contractible simple closed curves σ' and σ'' contained in F , so that $\Sigma - \{\sigma', \sigma''\}$ contains two components. One component A that contains K , and one component B that is homeomorphic to either an annulus or a Möbius strip and contains σ . If A is homeomorphic to an annulus, then we construct a new surface Σ' by identifying the boundary component B_1 of B with the boundary of a closed disk C_1 , and identifying the boundary component B_2 of B with the boundary of a closed disk C_2 . If A is homeomorphic to a Möbius strip, then we construct a new surface Σ' by identifying the boundary of B with the boundary of a closed disk. In either case the surface Σ' contains K and has smaller genus than Σ .

Since Σ has finite genus, we can only repeat this process finitely many times. In the resulting surface Σ^* , we have an embedding of K so that the face F' of K in Σ^* corresponding to the face F of K in Σ is simply connected. Thus we can restrict our attention to embeddings K whose faces are simply connected.

Theorem 16.C.3 in [4] states that any simply connected noncompact 2-manifold is homeomorphic to an open disk. Thus each face F of K is homeomorphic to an open disk. We use the notation $B(x, \epsilon)$ to denote the open disk in Σ centred at x with radius ϵ . In the plane we let $B(0, 1)$ be the unit disk centred at the origin, and $\mathbb{S}^1 = \text{Bd}(B(0, 1))$. For each face F we have a natural homeomorphism $h : B(0, 1) \rightarrow F$.

Now consider the boundary of a face F of K . By definition, $\text{Bd}(F)$ is a closed subset of K . We also have that $\text{Bd}(F)$ is connected.

Lemma 3.2

If Σ is a surface and $S \subseteq \Sigma$ is homeomorphic to an open disk, then $\text{Bd}(S)$ is a closed connected subset of Σ .

Proof $\Sigma - \text{Cl}(S)$ is an open (perhaps empty) subset of Σ . Since S is homeomorphic to an open disk, we have a homeomorphism $h : B(0, 1) \rightarrow S$. Consider $B(0, \epsilon)$ for any $0 < \epsilon < 1$. Then $h(B(0, \epsilon))$ is an open subset of Σ . Thus $\text{Cl}(S) - h(B(0, \epsilon))$ is a closed subset of Σ . Note that

$$\text{Bd}(S) = \text{Cl}(S) \cap (\text{Cl}(\Sigma - S)) = \text{Cl}(S) \cap (\Sigma - S) = \text{Cl}(S) - S,$$

so we have

$$\text{Cl}(S) - h(B(0, \epsilon)) = (S - h(B(0, \epsilon))) \cup \text{Bd}(S).$$

Since $B(0, 1) - B(0, \epsilon)$ is connected, $S - h(B(0, \epsilon))$ is connected and thus $(S - h(B(0, \epsilon))) \cup \text{Bd}(S)$ is a closed connected subset of Σ .

Now consider the sequence $a_i = 1 - 1/2^i$ for each $i \in \mathbb{N}$. Set

$$S_i = (S - h(B(0, a_i))) \cup \text{Bd}(S).$$

Now each S_i is a closed connected subset of Σ , and $S_i \supseteq S_j$ for all $i < j$. Thus by Proposition 2.1 $S' = \bigcap S_i$ is a non-empty closed connected subset of Σ . But $S' = \text{Bd}(S)$, since for each $x \in S$ we have $i \in \mathbb{N}$ such that $x \notin S_i$. ■

Thus by the propositions in Section 2.2 $\text{Bd}(F)$ is arcwise connected, and that $\text{Bd}(F)$ is a graph-like subspace of K .

3.2 Edges

In this section we consider the edges of the graph-like space K embedded in Σ . Since K is graph-like we have a set E of edges of K that consists of a set of pairwise disjoint arcs in Σ , and a set $V = K - E$ that is a totally disconnected subset of Σ .

We are primarily interested in the faces of K , and the relation between the faces of K and the edges of K . We would like to show that for each edge $e \in E$, e either lies entirely in the boundary of F , or is disjoint from the boundary of F . In order to show this we need an important lemma.

Lemma 3.3

For each edge $e \in E$, e is contained in an open subset of Σ , U_e , such that U_e is homeomorphic to an open disk and $U_e \cap K = e$.

In this section we will prove Lemma 3.3 and provide several useful corollaries.

When we are working in the plane we will make frequent use of the following well-known topological theorem.

Theorem 3.4 (Jordan-Schönflies)

If f is a homeomorphism of a simple closed curve C in the plane onto a closed curve C' in the plane, then f can be extended to a homeomorphism of the entire plane [9].

The Jordan-Schönflies Theorem provides an essential tool for working with embeddings of graphs, as well as embeddings of graph-like spaces. We will also use a version of the Jordan-Schönflies Theorem proven by Thomassen in [16].

Theorem 3.5 ([16], Thm. 3.3)

Let Γ and Γ' be 2-connected plane graphs such that g is a homeomorphism and plane-isomorphism of Γ onto Γ' . Then g can be extended to a homeomorphism of the entire plane.

A *plane isomorphism* is a graph isomorphism between 2-connected plane graphs Γ and Γ' so that if the cycle C is mapped to the cycle C' , then C bounds a face of Γ if and only if C' bounds a face of Γ' .

Lemma 3.3 is the main result of this section. We break the proof of Lemma 3.3 into parts, first proving the following three propositions.

Proposition 3.6

Given an edge $e \in E$ and a point $p \in e$ there is an open neighbourhood N_p of p , such that N_p is homeomorphic to an open disk and $N_p \cap K$ is an open subinterval of e .

Proof Consider the point $p \in e$. We can choose an arbitrary open disk neighbourhood D of p in Σ . Now consider $K \cap D$. Since V is a totally disconnected compact subset of Σ and $p \notin V$ we can rechoose D so that $D \cap V = \emptyset$. Thus $D \cap K$ consists only of points of edges of K . Now consider the homeomorphic image of D in the plane. The connected components of $D \cap K$ map to non-intersecting arcs in the plane. One of these arcs is the subarc τ of e containing p . We can use Theorem 3.4 to construct a homeomorphism h between D and $B(0, 1)$ so that τ is mapped to the horizontal arc $(-1, 1) \times \{0\}$ (i.e. the subarc of the x -axis between $y = -1$ and $y = 1$), and the point p is mapped to $(0, 0)$ (i.e. the origin).

Note that p is not an accumulation point for any sequence $\{p_i\}$ of points $p_i \in e_i$ where $e_i \in E - \{e\}$. Thus there is some $0 < \epsilon < 1$ such that if $e_i \in E - \{e\}$

and $e_i \cap D \neq \emptyset$, then $h(e_i \cap D) \cap B(0, \epsilon) = \emptyset$. Furthermore, τ is an open subarc of e . Therefore there is some $0 < \epsilon' < \epsilon$ such that if $x \in \tau$ and

$$h(x) \in B(0, \epsilon) - ((-1, 1) \times \{0\}),$$

then $d(h(x), p) > \epsilon'$. Thus

$$B(0, \epsilon') \cap h(K \cap D) = (-\epsilon', \epsilon') \times \{0\}.$$

Now $h^{-1}(B(0, \epsilon'))$ is an open neighbourhood of p satisfying our requirements. ■

Proposition 3.6 shows that every interior point of e has a well-behaved neighbourhood. We use this fact, together with compactness to construct well-behaved neighbourhoods of each edge.

Proposition 3.7

Given a compact subspace S of Σ there is a finite open cover, $\{D_i : 1 \leq i \leq n\}$, of S such that $\text{Bd}(D_i)$ is a simple closed curve for each i and $\text{Bd}(D_i) \cap \text{Bd}(D_j)$ is finite whenever $i \neq j$.

In order to prove this proposition we borrow ideas from the proof of Lemma 16.A.4 from [4] and ideas from the proof of Theorem 4.1 from [16].

Proof For each $p \in S$ choose a disk neighbourhood D_p of p . Now for each D_p we have an associated homeomorphism $h_p : B(0, 1) \rightarrow D_p$. Let $C_1(p) = h_p(\text{Bd}(B(0, 1/4)))$ and $C_2(p) = h_p(\text{Bd}(B(0, 3/4)))$. We take U_p to be the open disk $h_p(B(0, 1/4))$ corresponding to p . Note that $\text{Bd}(U_p) = C_1(p)$. Now the set $\{U_p : p \in S\}$ is an open cover of S by open disks whose boundary is a simple closed curve. Since S is compact we have a finite subcover $\{U_{p_i} : 1 \leq i \leq n\}$. We massage this finite open cover in order to obtain an open cover with the desired properties.

We inductively build our open cover by re-choosing each $C_1(p_i)$ in sequence. Start by setting $C(p_1) = C_2(p_1)$, and $U'_{p_1} = h_{p_1}(B(0, 3/4))$. Now $C_1(p_1) \subseteq \text{Cl}(U'_{p_1})$. Suppose that we have chosen simple closed curves $C(p_i)$ for each $1 \leq i \leq j-1 < n$, and let U'_{p_i} be the open disk bounded by $C(p_i)$. Furthermore suppose that the $C(p_i)$ have been chosen so that $C_1(p_i) \subseteq \text{Cl}(U'_{p_i})$ for each i and $C(p_i) \cap C(p_l)$ is finite whenever $i \neq l$ and $1 \leq i < l \leq j-1$. We consider the point p_j , and the curves $C_1(p_j)$ and $C_2(p_j)$. If $C_1(p_j)$ intersects each of the curves $C(p_i)$ in finitely

many points, then we take $C(p_j) = C_1(p_j)$. If $C_2(p_j)$ intersects each of the curves $C(p_i)$ in finitely many points, then we take $C(p_j) = C_2(p_j)$.

Assume that neither curve intersects each of the $C(p_i)$ finitely. In the plane, $C_2(p_j)$ and $C_1(p_j)$ define the boundary of an open annulus, A . Choose points a and b in $A - \cup_{1 \leq i \leq j-1} C(p_i)$ so that a and b lie on distinct radii r_1 and r_2 of $B(0, 1)$ in the plane. Then r_1 and r_2 divide A into two connected components, A_1 and A_2 , each of which is an open disk in the plane.

Next we find an arc α from a to b in A_1 and an arc β from a to b in A_2 such that $\alpha \cap (\cup_{1 \leq i \leq j-1} C(p_i))$ and $\beta \cap (\cup_{1 \leq i \leq j-1} C(p_i))$ are finite. Then $C(p_j) = \alpha \cup \beta \cup \{a, b\}$ satisfies our requirements.

We call an arc τ in A_1 admissible if $\tau \cap (\cup_{1 \leq i \leq j-1} C(p_i))$ is finite. Consider a point $x \in A_1$. If $x \notin \cup_{1 \leq i \leq j-1} C(p_i)$ then x has a small disk neighbourhood in A_1 disjoint from each $C(p_i)$. Any two points in this neighbourhood can be connected by an admissible arc. If $x \in \cup_{1 \leq i \leq j-1} C(p_i)$ then we have two possibilities, either x lies on a unique $C(p_i)$ or x is a crossing point of some of the $C(p_i)$. In the first case we can apply Proposition 3.6 to find a disk neighbourhood of x in A_1 so that its intersection with $C(p_i)$ is a single arc. Any two points in this neighbourhood can be connected by an admissible arc. In the second case we can apply Proposition 3.6 a finite number of times to find a disk neighbourhood of x in A_1 so that its intersection with each $C(p_i)$ is a single arc. Any two points in this neighbourhood can be connected by an admissible arc.

Now consider the ‘‘admissible arc’’-wise connected components of A_1 . Since for each $x \in A_1$ there is an open neighbourhood of x contained in a single ‘‘admissible arc’’-wise connected component, each of these components is both open and closed in A_1 . Thus since A_1 is connected any two points in A_1 can be connected by an admissible arc. The same holds for A_2 . Furthermore, since a and b do not lie on any of the $C(p_i)$ we can take α_1 and α_2 admissible arcs from a to some $x_1 \in A_1$ and from b to some $x_2 \in A_1$ respectively. Therefore α exists. The same reasoning holds for A_2 , so β exists. Thus we have $C(p_j)$ as required. ■

Every edge e of K is an open arc in Σ . We apply the approach used in the proof of Proposition 3.7 to find a well-behaved cover of every closed subarc of an edge e .

Proposition 3.8

If τ is a simple open arc in Σ , and S is a closed subarc of τ , then there is a finite open cover of S , $\{D_i : 1 \leq i \leq n\}$, such that $\text{Bd}(D_i)$ is a simple closed curve for each i

and $\text{Bd}(D_i) \cap \text{Bd}(D_j)$ is finite whenever $i \neq j$, and each $\text{Bd}(D_i)$ intersects τ in exactly two points.

Proof Choose arbitrary points u' and v' on τ . Consider the traversal of τ so that in the order on τ induced by the traversal we have $u' < v'$. Now we refer to the closed subarc of τ between u' and v' as $[u', v']$.

Observe that $S = [u', v']$ is a compact subset of Σ . At this point we can invoke Proposition 3.7 in order to find a finite cover of $[u', v']$ with nice properties. However, for this proof we need to be able to choose a cover with additional properties. We walk through the proof of Proposition 3.7 again.

For each point $p \in [u', v']$ use Proposition 3.6 to choose an open disk neighbourhood D_p of p such that $D_p \cap K$ is a subarc of τ . Now for each D_p we have an associated homeomorphism $h_p : B(0, 1) \rightarrow D_p$. Let $C_1(p) = h_p(\text{Bd}(B(0, 1/4)))$ and $C_2(p) = h_p(\text{Bd}(B(0, 3/4)))$. We take U_p to be the open disk $h_p(B(0, 1/4))$ corresponding to p . Note that $\text{Bd}(U_p) = C_1(p)$. The set $\{U_p : p \in S\}$ is an open cover of S by open disks whose boundary is a simple closed curve that intersects τ in exactly two points. Since S is compact we have a finite subcover $\{U_{p_i} : 1 \leq i \leq n\}$. We massage this finite open cover in order to obtain an open cover with the desired properties.

We build our open cover by re-choosing each $C_1(p_i)$ in sequence. Start by setting $C(p_1) = C_2(p_1)$, and $U'_{p_1} = h_{p_1}(B(0, 3/4))$. Now $C_1(p_1) \subseteq \text{Cl}(U'_{p_1})$. Suppose that we have chosen simple closed curves $C(p_i)$ for each $1 \leq i \leq j-1 < n$, and let U'_{p_i} be the open disk bounded by $C(p_i)$. Furthermore suppose that the $C(p_i)$ have been chosen so that $C_1(p_i) \subseteq \text{Cl}(U'_{p_i})$, $C(p_i) \cap C(p_l)$ is finite whenever $i \neq l$, and $|C(p_i) \cap \tau| = 2$ for each i . We consider the point p_j , and the curves $C_1(p_j)$ and $C_2(p_j)$. If $C_1(p_j)$ intersects each of the curves $C(p_i)$ in finitely many points, then we take $C(p_j) = C_1(p_j)$. If $C_2(p_j)$ intersects each of the curves $C(p_i)$ in finitely many points, then we take $C(p_j) = C_2(p_j)$.

Assume that neither curve intersects each of the $C(p_i)$ finitely. In the plane, $C_2(p_j)$ and $C_1(p_j)$ define the boundary of an open annulus, A . Furthermore, since $C_2(p_j)$ and $C_1(p_j)$ both intersect τ in exactly two points, A is divided into two open disks by the arcs corresponding to $A \cap \tau$. Since each $C(p_i)$ intersects τ exactly twice, there are only finitely many points of intersection in $A \cap \tau$. Thus we can choose points a and b in $A - \cup_{1 \leq i \leq j-1} C(p_i)$ so that a and b lie on the distinct arcs r_1 and r_2 of $A \cap \tau$.

Next we find an arc α from a to b in A_1 and an arc β from a to b in A_2 such that

$\alpha \cap (\cup_{1 \leq i \leq j-1} C(p_i))$ and $\beta \cap (\cup_{1 \leq i \leq j-1} C(p_i))$ are finite. Then $C(p_j) = \alpha \cup \beta \cup \{a, b\}$ satisfies our requirements.

The remainder of the proof follows exactly as in the proof of Proposition 3.7 ■

Now we are ready to prove Lemma 3.3.

Proof An edge e is an open arc in Σ . Let the ends of e be u and v . Choose arbitrary points u' and v' on e so that in order from u to v we have $u < u' < v' < v$. Now we refer to the closed subarc of e between u' and v' as $[u', v']$. By Proposition 3.8 we have a finite open cover of $[u', v']$, $\{D_i : 1 \leq i \leq n\}$, such that $\text{Bd}(D_i)$ is a simple closed curve for each i , $\text{Bd}(D_i) \cap \text{Bd}(D_j)$ is finite whenever $i \neq j$, and $\text{Bd}(D_i)$ intersects e in exactly two points for each i . Now we proceed to alter this cover so that $n = 1$.

Let d_i be the simple closed curve $d_i = \text{Bd}(D_i)$. For each d_i let $d_i \cap e = \{a_i, b_i\}$ where $a_i < b_i$. Now if we have d_i and d_j such that $a_i \leq a_j < b_j \leq b_i$ then we remove D_j from our cover. The remaining D_i still form a cover with the desired properties.

Now order the D_i according to the left to right ordering of the a_i . We reduce the size of our cover by 1, by amalgamating D_1 and D_2 . We have that d_1 and d_2 intersect finitely and that $a_1 < a_2 < b_1 < b_2$. We also have homeomorphisms $h_i : \text{Cl}(D_i) \rightarrow \text{Cl}(B(0, 1))$. In the plane $h_2(d_2 \cup (\text{Cl}(D_2) \cap (d_1 \cup e)))$ forms a 2-connected plane graph H . Consider the subgraph of H consisting of all vertices and edges of H that lie inside $h_2(\text{Cl}(D_2))$. We call this graph H' and note that H' is also 2-connected.

In H' the outer cycle $h_2(d_2)$ contains all of the vertices of H' except for $h_2(b_1)$ which is incident exactly with the two edges corresponding to e and the two edges corresponding to the arc of $h_2(\text{Cl}(D_2) \cap d_1)$ that intersects $h_2(\text{Cl}(D_2) \cap e)$ inside $h_2(D_2)$. Let f_1 denote the edge of H' corresponding to the subinterval (a_2, b_1) of e , f_2 the edge corresponding to the subinterval (b_1, b_2) of e , and f_3, f_4 the two edges corresponding to subarcs of $h_2(\text{Cl}(D_2) \cap d_1)$. Furthermore, $h_2(b_1)$ is incident with exactly four faces of H' , F_1, F_2, F_3 and F_4 , where $\text{Bd}(F_1)$ contains f_1, f_3 ; $\text{Bd}(F_2)$ contains f_3, f_2 ; $\text{Bd}(F_3)$ contains f_1, f_4 , and $\text{Bd}(F_4)$ contains f_2, f_4 . Now, since we have a plane drawing of H' we know that each F_i is an open disk in the plane. Thus we have two open disks $B_1 = F_1 \cup f_3 \cup F_2$, and $B_2 = F_3 \cup f_4 \cup F_4$.

For $i = 1, 2$, $\text{Bd}(B_i)$ is a simple closed curve in the plane, and

$$\text{Bd}(B_i) \cap h_2(e) = h_2([a_2, b_2]).$$

Furthermore, $\text{Bd}(B_1)$ contains a second edge g_1 of H incident with $h_2(a_2)$, and a second edge g_2 of H incident with $h_2(b_2)$ (i.e. $g_1 \neq f_1$ and $g_2 \neq f_2$). Now we can choose points x_1, x_2 in the interior of g_1 and y_1, y_2 in the interior of g_2 such that in clockwise order on $h_2(d_2)$ we have

$$h_2(a_2) < x_1 < x_2 < y_2 < y_1 < h_2(b_2).$$

Note that the subarcs $[h_2(a_2), x_2]$ and $[y_2, h_2(b_2)]$ of $h_2(d_2)$ contain no points of $h_2(\text{Cl}(D_2) \cap d_1)$.

Since B_1 is a connected subset of $h_2(\text{Cl}(D_2))$, we have an arc τ_1 from x_1 to y_1 in $B_1 \cup \{x_1, y_1\}$, and we have an arc τ_2 from x_2 to y_2 in $(B_1 - \{\tau_1\}) \cup \{x_2, y_2\}$. If either τ_1 or τ_2 intersect $\cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i)$ finitely then we take either τ_1 or τ_2 to be τ . If neither τ_1 nor τ_2 intersect $\cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i)$ finitely, then we find an arc τ from x to y so that $x_1 < x < x_2$, $y_2 < y < y_1$, $\tau \cap \tau_i = \emptyset$ for $i = 1, 2$, and τ intersects $\cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i)$ finitely.

Let A be the open disk bounded by $\tau_1, \tau_2, [x_1, x_2]$ and $[y_2, y_1]$. We call an arc τ in A admissible if $\tau \cap (\cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i))$ is finite. Consider a point $z \in A$. If $z \notin \cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i)$ then z has a small disk neighbourhood in A disjoint from each $\text{Cl}(D_2) \cap d_i$. Any two points in this neighbourhood can be connected by an admissible arc. If $z \in \cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i)$ then we have two possibilities, either z lies on a unique $\text{Cl}(D_2) \cap d_i$ or z is a crossing point of some of the $\text{Cl}(D_2) \cap d_i$. In the first case we can apply Proposition 3.6 to find a disk neighbourhood of z in A so that its intersection with $\text{Cl}(D_2) \cap d_i$ is a single arc. Any two points in this neighbourhood can be connected by an admissible arc. In the second case we can apply Proposition 3.6 a finite number of times to find a disk neighbourhood of z in A so that its intersection with each $\text{Cl}(D_2) \cap d_i$ is a single arc. Any two points in this neighbourhood can be connected by an admissible arc.

Now consider the ‘‘admissible arc’’-wise connected components of A . Since for each $z \in A$ there is an open neighbourhood of z contained in a single ‘‘admissible arc’’-wise connected component, each of these components is both open and closed in A . Thus since A is connected any two points in A can be connected by an

admissible arc. Let x and y be arbitrary points $x \in (x_1, x_2)$, $y \in (y_2, y_1)$ so that

$$x \cap (\text{Cl}(D_2) \cap d_i) = y \cap (\text{Cl}(D_2) \cap d_i) = \emptyset$$

whenever $i \neq 2$. Then we have an arc τ in A with endpoints x and y so that τ is admissible. Thus τ intersects $\cup_{1 \leq i \leq n} (\text{Cl}(D_2) \cap d_i)$ finitely, as required. We repeat this argument on B_2 to obtain an arc τ' from x' to y' that intersects the d_i finitely.

Consider the simple closed curve $h_2^{-1}(\tau \cup [y, y'] \cup \tau' \cup [x', x])$. This curve intersects e at a_2 and b_2 , intersects d_1 exactly twice, encloses an open disk, and contains the subarc $[a_2, b_2]$ of e . We replace D_2 with the open disk enclosed by this simple closed curve. Now $\text{Cl}(D_1) \cap \text{Cl}(D_2)$ is a closed disk. Thus $D = D_1 \cup D_2$ is homeomorphic to an open disk, $\text{Bd}(D)$ is a simple closed curve that intersects e exactly at a_1 and b_2 , and intersects the other d_i finitely. Furthermore $\text{Cl}(D_1) \subset \text{Cl}(D)$. Therefore we replace D_1 and D_2 by D to obtain a cover with the desired properties of size $n - 1$.

By induction, we obtain an open disk B_0 such that $\text{Bd}(B_0)$ is a simple closed curve that intersects e exactly at $\{a, b\}$ where $a < u' < v' < b$, and B_0 contains $[u', v']$.

Now consider the points u_1 and v_1 on e such that u_1 lies halfway between u and u' and v_1 lies halfway between v' and v . The subarc $[u_1, v_1]$ of e is a compact subset of Σ so we can enact the exact same process on $[u_1, v_1]$ in order to obtain an open disk B_1 such that $\text{Bd}(B_1)$ is a simple closed curve that intersects e exactly at $\{a, b\}$ where $a < u_1 < v_1 < b$, and B_1 contains $[u_1, v_1]$. However, we need to be more careful. Instead of enacting our process verbatim, we alter the process slightly so that $\text{Cl}(B_0) \subset \text{Cl}(B_1)$. Recall that by starting with D_1 we were able to ensure that $\text{Cl}(D_1) \subset \text{Cl}(B_0)$. We use this fact in order to choose B_1 with the desired properties.

Let x be an arbitrary fixed point in $[u', v']$. Since Σ is Hausdorff, for any point $p \in \Sigma$ there is an open disk neighbourhood of p that does not contain x .

Consider a point $p \in [u_1, u'] \cup (v', v_1]$. Use Proposition 3.6 to choose an open disk neighbourhood N_p of p such that $N_p \cap K$ is a subarc of e . Furthermore, we can choose N_p so that $x \notin N_p$. Now continue as before. In order to construct an open cover of $[u_1, v_1]$ we take the sets N_p for each $p \in [u_1, u'] \cup (v', v_1]$ together with B_0 . Thus when we choose a finite cover, B_0 is guaranteed to be included since it is the only set containing the point $x \in [u_1, v_1]$. Further, after we have

selected a finite subcover and eliminated nested sets B_0 still remains since none of the N_p contains B_0 . Suppose our finite cover is

$$\{U_i : 1 \leq i \leq n\} \cup \{B_0\} \cup \{V_i : 1 \leq i \leq m\}$$

where each U_i corresponds to a point in $[u_1, u')$ and each V_i corresponds to a point in $(v', v_1]$.

We are now free to enact the remaining part of our process as before. Consider only the set $\{U_i : 1 \leq i \leq n\} \cup \{B_0\}$. We start with B_0 and extend it to an open disk B'_0 such that $\text{Bd}(B'_0)$ is a simple closed curve that intersects e exactly at $\{a, b\}$ where $a < u_1 < v' < b$, B'_0 contains $[u_1, v')$, and $\text{Cl}(B_0) \subset \text{Cl}(B'_0)$. Now consider the set $\{V_i : 1 \leq i \leq m\} \cup \{B'_0\}$. We start with B'_0 and extend it to an open disk B_1 such that $\text{Bd}(B_1)$ is a simple closed curve that intersects e exactly at $\{a, b\}$ where $a < u_1 < v_1 < b$, B_1 contains $[u_1, v_1]$, and $\text{Cl}(B_0) \subset \text{Cl}(B_1)$.

We continue to construct a sequence of open disks $\{B_i\}$ where B_i contains the subarc $[u_i, v_i]$ of e . The corresponding sequences $\{u_i\}$ and $\{v_i\}$ are constructed by setting u_i and v_i to lie on e halfway between u and u_{i-1} and v and v_{i-1} respectively. Further, $\text{Bd}(B_i)$ is a simple closed curve for each i that intersects e in exactly two points. We also have that $\text{Cl}(B_{i-1}) \subset \text{Cl}(B_i)$ for each i .

We have that since $\{u_i\} \rightarrow u$ and $\{v_i\} \rightarrow v$, the set $U_e = \cup_{i=0}^{\infty} B_i$ is an open subset of Σ that contains each point of e . Further, $U_e \cap K = e$. We claim that U_e is homeomorphic to an open disk, and hence satisfies Lemma 3.3. To see this, consider any simple closed curve $\rho \subset U_e$. Since U_e inherits the metric topology on Σ , and ρ is compact, the minimum distance between ρ and $\text{Bd}(U_e)$ is non-zero. Thus there is some $i < \infty$ such that $\rho \subset B_i$. Since B_i is homeomorphic to an open disk, ρ is a contractible curve in B_i , and thus ρ is a contractible curve in U_e . This holds for all simple closed curves $\rho \subset U_e$, so U_e is simply connected. Theorem 16.C.3 in [4] states that any simply connected noncompact 2-manifold is homeomorphic to an open disk. Thus U_e is homeomorphic to an open disk, as required. ■

From Lemma 3.3 we can view each edge e as an open arc in $B(0, 1)$ with two endpoints on \mathbb{S}^1 . This allows us to choose arcs in a face F with very specific relation to e . For instance we can prove the following corollary.

Corollary 3.9

For each edge $e \in E$, e is contained in a closed subset of Σ , D_e , such that D_e is

homeomorphic to a closed disk and $D_e \cap K = \text{Cl}(e)$.

Proof By Lemma 3.3 we know that there is an open disk U_e in Σ such that $U_e \cap K = e$. We have that e is an arc in U_e such that $U_e - \{e\}$ has two connected components, both of which are open disks. Consider the homeomorphism $h : U_e \rightarrow B(0, 1)$, and let D_1 and D_2 be the open disks $h(U_e - \{e\})$. Now in \mathbb{S}^1 there are points u' and v' corresponding to $\text{Bd}(h(e))$. Thus we have simple closed curves τ_1 in D_1 and τ_2 in D_2 from u' to v' . The images $h^{-1}(\tau_1)$ and $h^{-1}(\tau_2)$ are arcs in Σ from u to v that lie in distinct components of $U_e - \{e\}$ (since the arcs τ_1 and τ_2 share the same endpoints as $h(e)$, their images in Σ share the same endpoints as e). Thus the curve

$$\rho = h^{-1}(\tau_1) \cup \{v\} \cup h^{-1}(\tau_2) \cup \{u\}$$

is a simple closed curve in Σ . Further $\rho \cap K = \{u, v\}$ and ρ is the boundary of an open disk D_e that contains e and no other points in K . Thus $\text{Cl}(D_e)$ satisfies the corollary. ■

Not only do we have a closed disk D_e for each edge e so that $D_e \cap K = \text{Cl}(e)$, but we can choose these disks so that they are non-intersecting.

Corollary 3.10

There is a set of closed disks $\{D_e\}$, one for each edge $e \in K$, so that $D_e \cap K = \text{Cl}(e)$ for each e , and $D_e \cap D_{e'}$ is either empty, or a single vertex for each $e \neq e'$.

Proof We begin with the entire edge set E of K . We have by Corollary 2.6 that E is a countable set. We take an ordering of E and consider the first edge. For e_1 we have by Corollary 3.9 that D_{e_1} exists. Now consider $\text{Bd}(D_{e_1})$. We have that $K \cup \text{Bd}(D_{e_1})$ has the same properties as K , except we have added two new edges to K . We continue to apply Corollary 3.9 to each edge e_i of E in turn. Each time we select D_{e_i} we add the boundary of D_{e_i} to K . This ensures that D_i and D_j are disjoint except perhaps at a single vertex for all $i \neq j$. The resulting set of closed disks satisfies our requirements. ■

Note that in the preceding corollary we added finitely many arcs to graph-like space K in order to obtain a new graph-like space. However, at the end of this process we may have added countably many arcs to K , and although

the statement of the corollary holds, we are not guaranteed that if we add the boundary of each closed disk to K the resulting space is graph-like. The problem is that the resulting space may not be compact. We return to this point later.

We now prove a final corollary. We have already shown that each edge either lies entirely in the boundary of F , or is disjoint from the boundary of F . Lemma 3.3 gives us another property of edges in relation to face boundaries.

Corollary 3.11

For each face F and each edge e , $\text{Bd}(F) \cap e$ is either empty, or the entire edge e . Furthermore, each edge e is in the boundary of either one or two faces of K .

Proof For any edge e , we have an open disk U_e . The edge divides U_e into two open disks, D_1 and D_2 . Since D_1 is a connected subset of Σ disjoint from K , $D_1 \subset F$ for some face F . Likewise D_2 is contained in some face F' of K . Thus e is entirely contained in the boundary of F and the boundary of F' , and $\text{Bd}(F'') \cap e = \emptyset$ for all faces $F \neq F'' \neq F'$. Furthermore, e is in the boundary of F and F' , where F and F' are not necessarily distinct. For any edge e , we have that e lies in the boundary of one or two distinct faces of K . ■

3.3 Face Boundaries

In this section we give a procedure for constructing a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ from the homeomorphism $h : B(0, 1) \rightarrow F$. This procedure is straightforward, but lengthy to describe. The basic idea is to define a homeomorphism $h' : B(0, 1) \rightarrow F$, so that boundaries of the disjoint closed disks from Corollary 3.10 are mapped to well-behaved arcs in $B(0, 1)$, specifically, arcs with exactly two endpoints on \mathbb{S}^1 . Then we will use the homeomorphisms between $\text{Cl}(B(0, 1))$ and each closed disk D_e , together with the Jordan-Schönflies Theorem to extend h' to the boundary of $B(0, 1)$.

First note that we may assume that each edge has two distinct endpoints in K . This follows, since if e has only one endpoint, we can choose an arbitrary point $v \in e$ and add v to V . The resulting space K' has two edges corresponding to e , both of which have two distinct ends. Note that regardless of how many edges we bisect in this way, K remains graph-like.

We consider a collection of simple closed curves in the plane that have a common base point. Given a point p in the plane, an n -flower is a finite collection

of simple closed curves τ_1, \dots, τ_n in the plane so that $\tau_i \cap \tau_j = \{p\}$ for all $i \neq j$; and each curve τ_i encloses an open disk D_i so that $D_i \cap \tau_j = \emptyset$ for each j . We call the curves τ_i *petals*.

Proposition 3.12

Suppose f and f' are n -flowers. Let τ_1, \dots, τ_n be a counter-clockwise ordering of the petals of f , and ρ_1, \dots, ρ_n be a counter-clockwise ordering of the petals of f' . There is a homeomorphism $h: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ so that $h(f) = f'$ and $h(\tau_i) = \rho_i$ for each $1 \leq i \leq n$.

Proof Instead of proving the statement of the proposition directly, we prove that given any n -flower, there is such a homeomorphism that maps f onto the standard n -flower. For this proof we refer to the origin of \mathbb{R}^2 as q . We use the standard polar coordinate system in order to describe the standard n -flower.

The standard n -flower is the n -flower centred at q where each petal is defined as follows. For $1 \leq i \leq 3n$ we let l_i be the line segment from q with length 1 at angle $2\pi i/3n$. Let q_i be the other end of l_i . Now for each $1 \leq j \leq n$ we consider the segments l_{3j-1}, l_{3j} and l_{3j+1} where the subscripts are computed modulo $3n$. For each j we let l'_j be the line segment that connects q_{3j-1} to q_{3j+1} in the plane. The curve ρ_j is defined as $l_{3j-1} \cup l'_j \cup l_{3j+1}$. Each ρ_j is a simple closed curve, $\rho_i \cap \rho_j = \{q\}$ for all $i \neq j$ and each ρ_i encloses an open disk D_i so that $D_i \cap \rho_j = \emptyset$ for each j . The petals ρ_i define the standard n -flower.

Consider the n -flower f with petals τ_1, \dots, τ_n in counter-clockwise order. Let q' be the centre of f . For each τ_i choose arbitrary points p_i, p'_i on $\tau_i - \{q'\}$ so that $p_i \neq p'_i$. Now p_i, p'_i define two subarcs of τ_i , one containing q' and one not containing q' . Let the subarc that does not contain q' be α_i . Since $\cup_{i=1}^n \alpha_i$ is a compact subset of the sphere, there is an open disk neighbourhood of q' , D , so that D is disjoint from $\cup \alpha_i$. By applying Proposition 3.6 repeatedly at the point q' we can assume that D is an open disk and that all of the τ_i have exactly two simple arcs inside D . Further we may assume that $\text{Bd}(D \cap \alpha_i) = \emptyset$ for each i .

Now construct a graph G by taking $\text{Bd}(D) \cup (\cup \tau_i)$ to be the entire graph with vertices q' together with the intersection points $\text{Bd}(D) \cap \tau_i$ for each i . We also construct a graph G' by taking a circle C with radius $1/4$ centred at q together with the loops ρ_i . This graph has vertices q together with the intersection points $C \cap \rho_i$ for each i . Note that G and G' are isomorphic 2-connected plane graphs.

If $\phi : G \rightarrow G'$ is an isomorphism then we can create a homeomorphism

$$h : (\cup \tau_i) \cup \text{Bd}(D) \rightarrow (\cup \rho_i) \cup C$$

by mapping $h(v) = \phi(v)$ for each vertex v . For each edge $e \in E(G)$ we have that e is an arc and we let h be the natural homeomorphism between the arc e and the arc $\phi(e)$. Now h is a homeomorphism, and a plane isomorphism, from G to G' . By Theorem 3.5 h extends to a homeomorphism of the entire plane that maps τ_i onto ρ_i for each i .

Thus the statement of the lemma holds for any n -flower f and the standard n -flower. Given two n -flowers f and f' we have homeomorphisms h and h' of the plane that map f and f' to the standard n -flower. Composing h and h' appropriately gives the desired result. ■

We have that K is a connected graph-like space embedded in Σ . We also have that a face F is a connected component of $\Sigma - K$, and we have stipulated that there is a homeomorphism $h : F \rightarrow B(0, 1)$. We have proven that:

1. $\text{Bd}(F)$ is connected,
2. $\forall e \in E(K)$, either $e \cap \text{Bd}(F) = \emptyset$ or $e \cap \text{Bd}(F) = e$,
3. $E(\text{Bd}(F))$ is countable, and
4. \exists a set $\{D_e : e \in E(K)\}$ of subsets of Σ such that there is a homeomorphism $h_e : D_e \rightarrow \text{Cl}(B(0, 1))$ for each e , $D_e \cap K = \text{Cl}(e)$ for each e , and for each $e \neq e'$, $D_e \cap D_{e'} = \text{Cl}(e) \cap \text{Cl}(e')$.

Now for each $e \in E(\text{Bd}(F))$ consider D_e . If F occurs on only one side of e then $\text{Bd}(D_e) \cap F = \tau_e$ is an arc in F . If F occurs on both sides of e then $\text{Bd}(D_e) \cap F$ consists of two arcs, τ_e and τ'_e . Since $E(\text{Bd}(F))$ is countable,

$$A = \{\tau_e : e \in E(\text{Bd}(F))\} \cup \{\tau'_e : e \in E(\text{Bd}(F))\}$$

is countable. Let $\{\tau_i : i \in \mathbb{N}\}$ be an arbitrary ordering of A .

We have that h maps F onto an open disk. While the arcs of A are very well behaved in F (i.e. each τ_i is an arc with two distinct endpoints in the boundary of F) the arcs in $h(A)$ may not share this property. In order to construct

our continuous surjection it will be essential that the arcs $h(A)$ behave in a prescribed way. In order to guarantee this behaviour we first construct a new homeomorphism between F and $B(0, 1)$.

The first step is to compactify $h(F) = B(0, 1)$. We define the quotient map $q : \text{Cl}(B(0, 1)) \rightarrow \mathbb{S}^2$ as,

$$q(x) = \begin{cases} \omega & \text{if } x \in \mathbb{S}^1, \\ \omega' & \text{if } x = (0, 0), \\ (1, \theta, \pi r) & \text{if } x = (r, \theta), \end{cases}$$

where ω is the north pole, ω' is the south pole, and the points are expressed in standard polar and spherical coordinates. Not only is q a quotient map from the closed unit disk to the sphere, but it is a homeomorphism between $B(0, 1)$ and $\mathbb{S}^2 - \{\omega\}$. Consider $q(h(F))$. In the sphere, each τ_i is mapped to a simple closed curve passing through ω . We will be able to apply Proposition 3.12 in order to alter h . Also, we are now able to prove some small facts about the structure of A , $h(A)$ and $q(h(A))$.

Proposition 3.13

$F - \tau_i$ consists of two open disks, one of which contains each arc in $A - \{\tau_i\}$.

Proof Since $q(h(\tau_i)) \cup \{\omega\}$ is a simple closed curve in the sphere through ω , $q(h(\tau_i))$ separates the sphere into two open disks, corresponding to the open disks $F - \tau_i$. Since one of the disks in $F - \tau_i$ contains all other τ_j we have the desired result. ■

Proposition 3.14

For any finite subset $S \subseteq A$, $F - S$ consists of $|S| + 1$ open disks, one of which contains each arc in $A - S$.

Proof This proposition follows directly by applying the previous claim and induction. ■

In order to prove the next proposition we will need $\{q(h(\tau_i)) : i \in \mathbb{N}\} \cup \{\omega\}$ to be a closed subset of the sphere, however this may not be true for arbitrary τ_i . We show that we can choose the disks D_e so that $\{q(h(\tau_i)) : i \in \mathbb{N}\} \cup \{\omega\}$ is closed.

For each τ_i we let $\rho_i = q(h(\tau_i))$. Now for each i , $\rho_i \cup \{\omega\}$ is a simple closed curve that separates the sphere into two open disks, one of which is disjoint from

$q(h(A))$. Note that we can replace ρ_i with any arc ρ in the disk which is disjoint from $q(h(A))$ provided that both ends of ρ are ω . Let C_i be the circle in the sphere centred at ω with radius $1/2^i$. Let D_j be the disk component of $\mathbb{S}^2 - \rho_j$ that is disjoint from $q(h(A))$. Then each C_i separates the disk D_j into open disks, one of which contains ω in its boundary. We replace ρ_j with ρ'_j , an arbitrary arc with both ends tending towards ω so that ρ'_j lies inside D_j and ρ'_j lies entirely inside an open disk component of $D_j - C_j$.

Now when we map the ρ'_j back to F we are effectively re-choosing the disks D_e . We label the disks defined by the ρ'_j as D'_e . The set $\{q(h(\tau'_i)) : i \in \mathbb{N}\} \cup \{\omega\}$ is a closed subset of the sphere. From here on we assume that A is chosen to have this property.

Proposition 3.15

We can choose the disks D_e so that $F - A$ consists of a collection of open disks, one of which, D , has each τ_i in its boundary.

Proof If $E(\text{Bd}(F))$ is finite, then this follows directly from the previous proposition. Assume that $E(\text{Bd}(F))$ is not finite.

Consider F . Each τ_i bounds a subset of F that is an open disk (i.e. τ_i splits F into two open disks, one of which contains all other τ_j). We have that each τ_i contributes an open disk to the set $F - A$. Let B be the collection of open disk components of $F - A$ that correspond to some τ_i . We want to prove that $(F - A) - B$ is an open disk that contains each τ_i in its boundary.

Note that we can replace the disks D_e with new disks, D'_e , so that each D'_e lies inside D_e and $\text{Bd}(D_e) \cap \text{Bd}(D'_e) = \text{Bd}(D_e)$. Now we have a corresponding sets A' and B' and corresponding arcs $\{\tau'_i : i \in \mathbb{N}\}$. By construction we have $D = (F - A') - B' \neq \emptyset$. We show that D is connected by showing that D is arcwise connected.

We have that F is arcwise connected. Thus for any $x, y \in F$, $x, y \notin A' \cup B'$ there is an arc from x to y in F , σ . If $\sigma \cap A = \emptyset$ then σ is an arc from x to y in D and we have the result. If $\sigma \cap A' \neq \emptyset$, then we rechoose σ as follows.

We have an ordering of the arcs in A' , so we consider them in turn. If σ does not intersect τ'_1 then we set $\sigma_1 = \sigma$. Otherwise σ intersects τ'_1 . But then either x and y lie inside $(D_1 \cup \tau_1) - \tau'_1$ in which case we can rechoose σ easily, or σ intersects τ_1 . If σ intersects τ_1 we have two cases. The first case is that $y \notin q(h(D_1))$. In this case we let σ intersect τ_1 first (in the traversal of σ from x

to y) at z_1 and last at z_2 . Now we replace the segment of σ between z_1 and z_2 with the segment of τ_1 from z_1 to z_2 and call the resulting arc σ_1 . The second possibility is that $y \in q(h(D_1))$. In this case we choose an arc from z_1 to y that avoids τ_1' to replace the tail of σ and thus create σ_1 . Now we have an arc σ_1 from x to y that does not intersect τ_1' .

We can carry on this construction for each $i \in \mathbb{N}$. Note that since none of the τ_i' intersect each other we can make the replacements specified in the above construction simultaneously. Thus we can replace σ with an arc σ' from x to y so that σ' lies entirely inside D . Therefore D is arcwise connected and hence connected.

Furthermore, D is simply connected, since given any simple closed curve σ in D , σ is a simple closed curve in F . Further σ bounds a disk in F . This disk contains none of the D_e'' since it contains none of the boundary of F . Thus σ is contractible in F and in D . Therefore D is simply connected. Since A' is a closed set, D is an open set and hence D is an open disk.

It remains to prove that $\tau_i' \subseteq \text{Bd}(D)$ for each i . For each $x \in \tau_i'$, for any open neighbourhood V of x , V contains points of D_e' and of $D_e - D_e'$ by construction. But then V contains points of D_e' and of D because $D_e - D_e' \subset D$. Thus $x \in \text{Bd}(D)$. Therefore τ_i' is in the boundary of D for each $\tau_i' \in A'$. This completes the proof. ■

We henceforth assume that the disks D_e and the arcs τ_i were chosen to comply with the previous proposition. Also note that the arcs $q(h(\tau_i))$ in the sphere all have the property that $\text{Cl}(q(h(\tau_i)))$ is a simple closed curve through ω . Furthermore, $q(h(\tau_i)) \cap q(h(\tau_j)) = \emptyset$. Note that we chose the arcs τ_i so that $q(h(\tau_i))$ is contained in the open disk $B(\omega, 1/2^i)$. Thus

$$K' = (\cup_{i>0} q(h(\tau_i))) \cup \omega$$

is a graph-like space in the sphere, and each edge of K' is a loop. Thus we can apply the result from Section 3.2 to K' . In particular we can apply Corollary 3.10 so there is an open disk containing each $q(h(\tau_i))$ so that any pair of open disks is disjoint. This will be useful later in our construction.

Consider distinct arcs τ_i and τ_j in A so that τ_i corresponds to edge e , τ_j corresponds to edge e' , and e, e' share an endpoint u . Note that e, e' may be the same edge, and in this case the vertex u under consideration is important. For any finite subset S of edges with endpoint u , there is a disk, $N_u(S)$, centred at u

so that, for each edge $f \in S$, $f \cap N_u(S)$ is an arc from u to a point on $\text{Bd}(N_u(S))$. The existence of $N_u(S)$ can be established by repeated application of Proposition 3.6. We say that arcs τ_i and τ_j are *adjacent* if and only if: for any finite subset S of the edges with endpoint u so that $e, e' \in S$, the arcs in $N_u(S)$ corresponding to e, e' are consecutive in the cyclic order, $\tau_i \cap N_u(S), \tau_j \cap N_u(S)$ lie in the same open disk segment of $N_u(S) - \text{Cl}(S)$, and we can rechoose $N_u(S)$ so that only finitely many $x \in V(K)$ lie in the same segment of $N_u(S) - \text{Cl}(S)$ because $\tau_i \cap N_u(S)$ and $\tau_j \cap N_u(S)$.

We define an *infinite flower* to be a flower with countably many petals. If f is an infinite flower centred at ω , and S is any finite subset of petals, we define N_S to be the open disk centred at ω with radius sufficiently small so that each petal in S corresponds to exactly two arcs from ω to $\text{Bd}(N_S)$ in N_S . The disk N_S exists by a repeated application of Proposition 3.6 to the finite flower defined by S . Now, given an infinite flower f centred at ω with petals ρ_i , we say that ρ_i and ρ_j are *adjacent* if and only if: for any finite subset of petals of f , S , so that $\rho_i, \rho_j \in S$, the arcs corresponding to ρ_i and ρ_j in N_S are consecutive in the cyclic order on $\text{Bd}(N_S)$.

Given any face F , the arcs in A are mapped by $h \circ q$ to a flower in the sphere.

Proposition 3.16

Arcs τ_i and τ_j are adjacent in F if and only if $q(h(\tau_i))$ and $q(h(\tau_j))$ are adjacent in the sphere.

Proof Suppose that $q(h(\tau_i))$ and $q(h(\tau_j))$ are not adjacent in the sphere. Then we have petals ρ_1 and ρ_2 so that if $S = \{q(h(\tau_i)), \rho_1, q(h(\tau_j)), \rho_2\}$, then $q(h(\tau_i)), \rho_1, q(h(\tau_j)), \rho_2$ appear in this cyclic order on N_S . Now choose arbitrary $p_1 \in \rho_1$ and $p_2 \in \rho_2$ and connect them by an arc α in D . In the plane $q^{-1}(\rho_1 \cup \rho_2 \cup \alpha)$ partitions $B(0, 1)$ into four open disks, one of which contains $h(\tau_i)$ and another contains $h(\tau_j)$. In Σ , $h^{-1}(q^{-1}(\rho_1 \cup \rho_2 \cup \alpha))$ partitions F into four open disks, one of which contains τ_i and another contains τ_j . But since ρ_1 and ρ_2 correspond to closed disks D_e and D_f we can choose arbitrary points $x \in e$ and $y \in f$ and extend $h^{-1}(q^{-1}(\alpha))$ to an arc β from x to y in F . We achieve this by adding arcs from x to $h^{-1}(q^{-1}(p_1))$ and from $h^{-1}(q^{-1}(p_2))$ to y that lie inside D_e and D_f respectively. Now β divides F into two open disks, one containing τ_i and the other containing τ_j . Thus τ_i and τ_j cannot be adjacent in F , since τ_i and τ_j correspond to edges that do not share a vertex.

Now suppose that τ_i and τ_j are not adjacent in F . We have three possible cases.

Case #1: τ_i corresponds to e , τ_j corresponds to e' and e, e' do not share an endpoint.

Case #2: τ_i, τ_j correspond to $e = uv$ and each of u, v is either the end of multiple edges, or a vertex accumulation point.

Case #3: τ_i corresponds to $e = uv$, τ_j corresponds to $e' = uv'$ and either the arcs in $N_u(S)$ corresponding to e, e' are not consecutive in the cyclic order, $\tau_i \cap N_u(S), \tau_j \cap N_u(S)$ do not lie in the same open disk segment of $N_u(S) - \text{Cl}(S)$, or we cannot rechoose $N_u(S)$ so that only finitely many $x \in V(K)$ lie in the same segment of $N_u(S) - \text{Cl}(S)$ because $\tau_i \cap N_u(S)$ and $\tau_j \cap N_u(S)$.

Now in each case we find an arc that separates $q(h(\tau_i))$ from $q(h(\tau_j))$. The method is the same for each case. Take $p_1 \in \tau_i, p_2 \in \tau_j$ arbitrarily and connect them by an arc α in $F - \{\tau_i, \tau_j\}$. Extend α to an arc β that separates F into two open disks by adding two arcs to α , one from τ_i to an arbitrary interior point of its corresponding edge, and one from τ_j to an arbitrary interior point of its corresponding edge. These disks each contain an edge of K in their boundary. These edges, f, f' , are distinct from e, e' .

We use the same method to construct an edge α' from a point on σ_1 to a point on σ_2 , where σ_1, σ_2 correspond to the boundaries of D_f and $D_{f'}$ respectively. We extend this arc to an arc β' from a point on f to a point on f' as before. Now $F - \beta'$ consists of two open disks, one containing τ_i and the other containing τ_j . The closed curve $q(h(\beta'))$ partitions the sphere into two open disks, but the arcs $q(h(\sigma_1))$ and $q(h(\sigma_2))$ lie between the arcs $q(h(\tau_i))$ and $q(h(\tau_j))$. Thus $q(h(\tau_i))$ and $q(h(\tau_j))$ are not adjacent in the sphere. ■

Now we alter our homeomorphism h so that we can extend it easily to a continuous surjection from \mathbb{S}^1 to $\text{Bd}(F)$. At the end of this process we will have a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ so that $g|_{B(0, 1)}$ is a homeomorphism from $B(0, 1)$ to F . We construct intermediate functions h_1 and h_2 .

We begin by constructing a homeomorphism $h_1 : B(0, 1) \rightarrow F$ so that $h_1^{-1}(\tau_i)$ is well-behaved for each arc τ_i . We continue to make use of the quotient map q from $\text{Cl}(B(0, 1))$ to \mathbb{S}^2 . The construction of h_1 is straightforward, but long to describe. The idea is to repeatedly apply Theorem 3.5 to subsets of arcs $q(h(\tau_i))$ and sets of line segments in the unit disk. In essence we are simply altering the

homeomorphism h . We have that h maps F to $B(0,1)$. For each curve τ_i we identify τ_i with its image $q(h(\tau_i))$ in the sphere.

We have a countable number of arcs τ_i to fix, so we consider them in turn. We specify an ordering of the τ_i so that τ_1 and τ_2 are not adjacent. If this is not possible, then any two arcs τ_i and τ_j are adjacent. In this case we trisect any edge $e \in E(\text{Bd}(F))$. The middle edge e' created by trisecting e has an arc τ that is not adjacent to any arc τ' in the original construction. Thus we may assume that τ_1 and τ_2 are non-adjacent.

We start with C , a copy of \mathbb{S}^1 . Choose points p_1, p_2, p_3 and p_4 on C , so that the p_i are ordered clockwise, and $C - \{p_1, p_2, p_3, p_4\}$ consists of 4 open arcs that are equal in length. Let l_1 be the line segment joining p_1 to p_2 , and let l_2 be the line segment joining p_3 to p_4 . Choose points q_1 and q_2 on l_1 from p_1 to p_2 so that q_1 and q_2 trisect l_1 . Choose points q_3 and q_4 on l_2 from p_3 to p_4 so that q_3 and q_4 trisect l_2 . Now let α_1 be the line segment joining q_1 and q_3 , and let α_2 be the line segment joining q_2 and q_4 . Figure 3.1 roughly depicts our construction.




Figure 3.1: Arcs l_1 and l_2 .

In the sphere, we have a 2-connected graph G formed by $q(l_1 \cup l_2 \cup \alpha_1 \cup \alpha_2) \cup \omega$. We want to map τ_1 to l_1 and τ_2 to l_2 .

Consider τ_1 and τ_2 in the sphere. We can choose arbitrary points $x_1 \in \tau_1$ and

$y_1 \in \tau_2$ and join them by an arc σ_1 in D . The arc σ_1 divides the open disk D into two disks, D' and D'' . Choose arbitrary points $x_2 \in \tau_1$ and $y_2 \in \tau_2$ so that x_2, y_2 both lie in the boundary of D'' . Now join x_2 to y_2 by an arc σ_2 in D'' . We now have a graph G' formed by $q(\tau_1 \cup \tau_2 \cup \sigma_1 \cup \sigma_2) \cup \omega$. Furthermore, there is an homeomorphism between G and G' that maps τ_1 to l_1 and τ_2 to l_2 , and by Theorem 3.5 we can extend this homeomorphism to a homeomorphism of the entire plane.

Note that by construction τ_1 is mapped onto l_1 and the open disk bounded by τ_1 is mapped to the open disk bounded by l_1 . Also, all of the arcs in $A - \{\tau_1, \tau_2\}$ lie either in the open disk bounded by α_1 or the open disk bounded by α_2 . Thus we can proceed by fixing our homeomorphism on $B(0, 1)$ with the exception of the two disks bounded by α_1 and α_2 , where we continue to construct h_1 . This is, in broad terms, our strategy. However, in order for h_1 to be a homeomorphism, we need to map $A - \{\tau_1, \tau_2\}$ more carefully. We begin the process again.

We have the circle C together with line segments l_1, l_2, α_1 and α_2 . Let l'_1 be the line segment joining p_1 and p_4 , and let l'_2 be the line segment joining p_2 and p_3 . Let $x_1 \in l'_1$ bisect l'_1 , $x_2 \in \alpha_1$ bisect α_1 , $x_3 \in \alpha_2$ bisect α_2 and $x_4 \in l'_2$ bisect l'_2 . Let β_1 be the line segment joining x_1 and x_2 , and let β_2 be the line segment joining x_3 and x_4 . Now we have the construction depicted in Figure 3.2.

Now we define the graph G to be

$$q(l_1 \cup l_2 \cup l'_1 \cup l'_2 \cup \alpha_1 \cup \alpha_2 \cup \beta_1 \cup \beta_2) \cup \omega.$$

Note that G is 2-connected.

We have arcs τ_1, τ_2, σ_1 , and σ_2 as before. We note that since we chose the arcs τ_1 so that $K \cup A$ is graph-like, we can apply Lemma 3.3 to obtain open disks B_1 and B_2 so that B_i contains τ_i , and all other τ_j lie outside B_i . Thus, in the sphere we have Figure 3.3, where the dashed circles correspond to the outer boundary of B_1 and B_2 .

Since B_1 is an open disk that contains τ_1 , and no other τ_i , we use B_1 to choose an arc a_1 from ω to x_1 that does not intersect any of the τ_i or σ_i . We have that one face of τ_1 contains each τ_i . The set $B_1 - \tau_1$ consists of two open disks, one of which, B'_1 , lies in the face of τ_1 that contains each τ_i . Furthermore, $B'_1 - \{\sigma_1, \sigma_2\}$ consists of three open disks, one of which, B''_1 , contains both x_1 and ω in its boundary. Thus we can connect x_1 to ω by an arc a_1 that lies entirely inside B''_1 .




Figure 3.2: Arcs l'_1 and l'_2 .




Figure 3.3: Arcs σ_1 and σ_2 .

Also there is an open disk component of $B'_1 - \{\sigma_1, \sigma_2\}$ that contains x_2 and ω in its boundary. Therefore we can join x_2 to ω by an arc a'_1 that does not intersect any of the τ_i . Now we perform the same construction on τ_2 to find an arc a_2 from ω to y_1 and an arc a'_2 from ω to y_2 . We have constructed the object in Figure 3.4.






Figure 3.4: Arcs $a_1, a_2, a'_1,$ and a'_2 .

Now since σ_1 encloses a disk D' that does not contain σ_2 in its boundary, $D' - \{a_1, a_2\}$ consists of three open disks, one of which contains the arcs of $A - \{\tau_1, \tau_2\}$ contained in D' . We can connect an arbitrary interior point z_1 of a_1 to an arbitrary interior point z_2 of a_2 by an arc σ'_1 that does not intersect any of the τ_i , nor σ_1 . Similarly we can choose points z'_1 and z'_2 interior to a'_1 and a'_2 respectively, and join z'_1 to z'_2 by an arc σ'_2 that does not intersect any of the τ_i nor σ_2 . We now have the construction in Figure 3.5.

Finally we let τ'_1 be the simple closed curve consisting of σ'_1 together with the subarcs of a_1 and a_2 that join σ'_1 to ω . We let τ'_2 be the simple closed curve consisting of σ'_2 together with the subarcs of a'_1 and a'_2 that join σ'_2 to ω . We join τ'_1 to σ_1 by an arc b_1 and τ'_2 to σ_2 by an arc b_2 . We have constructed the object depicted in Figure 3.6.

Figure 3.5: Arcs σ_1' and σ_2' .Figure 3.6: Simple closed curves τ_1' and τ_2' .

We define the graph G' to be

$$\tau_1 \cup \tau_2 \cup \tau'_1 \cup \tau'_2 \cup \sigma_1 \cup \sigma_2 \cup b_1 \cup b_2 \cup \omega.$$

Now there is a natural homeomorphism between G and G' that maps l_i to τ_i , l'_i to τ'_i , σ_i to α_i and b_i to β_i . We apply Theorem 3.5 to extend this homeomorphism to a homeomorphism of the entire sphere, $\phi_1 : \mathbb{S}^2 \rightarrow \mathbb{S}^2$.

We take the faces F_1 and F_2 of G' to be the faces bounded by τ'_1 and τ'_2 . Note that these faces are mapped by ϕ_1 to the disks enclosed by l'_1 and l'_2 . By Proposition 3.16 the faces F_1 and F_2 contain all of $A - \{\tau_1, \tau_2\}$. We continue to build h_1 by considering the faces F_1 and F_2 , and the disks $\phi_1(F_1)$ and $\phi_1(F_2)$.

We fix each arc τ_i in turn by a recursive process. For each step we have three cases, depending on how τ_i interacts with the arcs τ_j for $j < i$. The cases are: τ_i is adjacent to two arcs τ_j, τ_k for $j, k < i$; τ_i is adjacent to one arc τ_j for $j < i$; and, τ_i is not adjacent to any τ_j for $j < i$. We describe each of these three cases separately, although they are very similar. At each step we apply a variant of Theorem 3.5. Namely, we can extend a homeomorphism between 2-connected graphs G and G' to the entire sphere.

We consider each of the three cases for τ_3 .

Case #1:

In this case τ_3 is adjacent to both τ_1 and τ_2 . Without loss of generality, τ_3 lies in the disk bounded by τ'_1, F_1 . We would like to map τ_3 to the line segment l'_1 , so instead of continuing with our construction, we take a step back. Instead of constructing τ'_1 , we take $\tau'_1 = \tau_3$, and we construct τ'_2 as before. Now τ_3 is mapped by ϕ_1 to l'_1 , and the open disk bounded by τ_3 is mapped to the open disk bounded by l'_1 . All subsequent arcs τ_i are contained inside τ'_2 , and we map them to the disk bounded by l'_2 .

Case #2:

In this case τ_3 is adjacent to τ_2 , but not to τ_1 . Without loss of generality, τ_3 lies in F_1 . We employ the same methods as in the construction of ϕ_1 to construct a line segment l_3 and map τ_3 to l_3 .

Since $K \cup A \cup \tau'_1$ is a graph-like space, there are open disks containing τ_3 and τ'_1 that contain none of the other τ_i . Thus we can choose arbitrary z_1 on τ_3 and z_2 on τ'_1 . Let a be an arc from z_1 to z_2 in the disk bounded by τ'_1 and each τ_i in the same face of τ'_1 as τ_3 . The arc a defines two new disks. Since τ_2 and τ_3 are

adjacent, one of these disks contains all of the other τ_i in the same face of τ'_1 as τ_3 . Let this disk be D . Let z'_1 be a point on τ_3 in the boundary of D , and let z'_2 be a point on τ'_1 in the boundary of D . Now we use the open disks containing τ_3 and τ'_1 to find arcs a_1 from ω to z'_1 and a_2 from ω to z'_2 as before so that a_1 and a_2 are subsets of D . Choose $x_1 \in a_1$ and $x_2 \in a_2$, and let b be an arc from x_1 to x_2 that does not intersect any τ_i nor a_i . Let τ'_3 be the simple closed curve composed of a together with the subarc of a_1 from ω to x_1 and the subarc of a_2 from ω to x_2 . One face of τ'_3 contains τ_3 , and the other contains all of the τ_i contained F_1 other than τ'_3 . Add an arc a_3 from a point on τ_3 to a point on τ'_3 that does not intersect any τ_i , nor any τ'_i . Add an arc b_3 from a point on a_3 to a point on τ'_1 that does not intersect any τ_i , nor any τ'_i . We have the construction shown in Figure 3.7.




Figure 3.7: Construction of τ'_3 in Case #2.

Let G be the 2-connected graph given by $\tau'_1 \cup \tau_3 \cup \tau'_3 \cup a_3 \cup b_3 \cup \omega$.

Now in the plane we extend our previous construction. Consider the open disk bounded by l'_1 . Let p be the point of C that bisects the subarc of C subtended by l'_1 . Let l'_3 be the line segment from p to p_1 , and l_3 be the line segment from p to p_4 . Let r_1 be the point that bisects l'_3 , r_2 be the point that bisects l_3 , and α_3 be the line segment joining r_1 and r_2 . Let r_3 be the point that bisects α_3 , let r_4 be

the point that bisects l'_1 and β_3 be the line segment joining r_3 and r_4 . We have the picture shown in Figure 3.8.




Figure 3.8: Construction of l'_3 in Case #2.

Let G' be the 2-connected graph given by $q(l'_1 \cup l_3 \cup l'_3 \cup \alpha_3 \cup \beta_3) \cup \omega$. Now we have a natural homeomorphism between G and G' that maps τ'_1 to l'_1 , τ_3 to l_3 , τ'_3 to l'_3 , a_3 to α_3 and b_3 to β_3 . Furthermore we can specify that our homeomorphism $\phi_2 : G \rightarrow G'$ agrees with ϕ_1 on τ'_1 . By Theorem 3.5, ϕ_2 extends to a homeomorphism between disk bounded by τ'_1 in the sphere and the disk bounded by l'_1 . All subsequent arcs τ_i either lie in the disk bounded by τ'_2 or the disk bounded by τ'_3 . We continue to map these arcs to the disks bounded by l'_2 and l'_3 .

Case #3:

In this case τ_3 is adjacent to neither τ_1 nor τ_2 . Without loss of generality, τ_3 lies in F_1 . We employ the same methods as in the construction of ϕ_1 to construct a line segment l_3 and map τ_3 to l_3 .


Since $K \cup A \cup \tau'_1$ is a graph-like space, there are open disks containing τ_3 and τ'_1 that contain none of the other τ_i . Thus we can choose arbitrary z_1 on τ_3 and z_2 on τ'_1 . Let a be an arc from z_1 to z_2 in the disk bounded by τ'_1 and each τ_i in the same face of τ'_1 as τ_3 . The arc a_3 defines two new disks D and D' that contain none of the arcs τ_i . Furthermore, we take D to be the disk containing

the arcs τ_i that lie between τ_1 and τ_3 in the sphere. Since τ_3 is not adjacent to τ_1 or τ_2 , both of these disks contain some of the arcs τ_i in the same face of τ'_1 as τ_3 . Choose y_1 in $\tau_3 - z_1$ in the boundary of D and y_2 in $\tau_3 - z_1$ in the boundary of D' . Let y'_1 be a point in $\tau'_1 - z_2$ in the boundary of D , and y'_2 be a point in $\tau'_1 - z_2$ in the boundary of D' . We have an open disk containing τ_3 that does not contain any other τ_i or τ'_i . We use this disk to find arcs a_1 from ω to y_1 in D and a_2 from ω to y_2 in D' so that each a_i is disjoint from all of the arcs τ_i and τ'_i . We let b_1 be an arc from ω to y'_1 in D and b_2 be an arc from ω to y'_2 so that each b_i is disjoint from all of the arcs τ_i and τ'_i . Now let σ_1 be an arc from an interior point of a_1 to an interior point of b_1 in D , and let σ_2 be an arc from an interior point of a_2 to an interior point of b_2 in D' . We have a simple closed curve τ'_3 consisting of σ_1 together with the subarc of a_1 connecting σ_1 to ω and the subarc of b_1 connecting σ_1 to ω . We also have a simple closed curve τ''_3 consisting of σ_2 together with the subarc of a_2 connecting σ_2 to ω and the subarc of b_2 connecting σ_2 to ω . Let z'_1 be an arbitrary point on τ_3 in the boundary of D' and z'_2 be an arbitrary point on τ'_1 in the boundary of D' . Let a'_3 be an arc from z'_1 to z'_2 in D' so that a'_3 does not intersect τ''_3 . Now let b_3 be an arc connecting an interior point of σ_1 to an interior point of a_3 in D , and let b'_3 be an arc connecting an interior point of σ_2 to an interior point of a'_3 in the open disk component of $D' - a'_3$ that does not contain a_3 in its boundary. This completes the construction at this step. The object we have constructed is depicted in Figure 3.9.

We let G be the 2-connected graph given by

$$\tau'_1 \cup \tau_3 \cup \tau'_3 \cup \tau''_3 \cup a_3 \cup a'_3 \cup b_3 \cup b'_3 \cup \omega.$$

Now in the plane we extend our previous construction. Consider the open disk bounded by l'_1 . Let p and p' be points in the subarc of C subtended by l'_1 so that the points p_4, p, p', p_1 appear in clockwise order, the subarc $[p_4, p]$ has length equal to $1/4$ of the length of $[p_4, p_1]$, the subarc $[p, p']$ has length equal to $1/2$ of the length of $[p_4, p_1]$ and the subarc $[p', p_1]$ has length equal to $1/4$ the length of $[p_4, p_1]$. Let l'_3 be the line segment joining p_1 to p' , l_3 be the line segment joining p' to p and l''_3 be the line segment joining p to p_4 . Let x_1 and x_2 be points that trisect l_3 so that the points p', x_1, x_2, p are in order on l_3 . Let y_1 and y_2 be points that trisect l'_1 so that the points p_1, y_1, y_2, p_4 are in order on l'_1 .

Figure 3.9: Construction of τ'_3 and τ''_3 in Case #3.

Let α_3 be the line segment joining x_1 and y_1 and α'_3 be the line segment joining x_2 and y_2 . Let z_1 bisect l'_3 , z_2 bisect α_3 , z_3 bisect α'_3 and z_4 bisect l''_3 . Let β_3 be the line segment joining z_1 and z_2 and let β'_3 be the line segment joining z_3 and z_4 . We have constructed the object shown in Figure 3.10.

Let G' be the 2-connected graph given by

$$q(l'_1 \cup l_3 \cup l'_3 \cup l''_3 \cup \alpha_3 \cup \alpha'_3 \cup \beta_3 \cup \beta'_3) \cup \omega.$$

Now we have a natural homeomorphism between G and G' that maps τ'_1 to l'_1 , τ_3 to l_3 , τ'_3 to l'_3 , τ''_3 to l''_3 , a_3 to α_3 , a'_3 to α'_3 , b_3 to β_3 , and b'_3 to β'_3 . Furthermore we can specify that our homeomorphism $\phi_2 : G \rightarrow G'$ agrees with ϕ_1 on τ'_1 . By Theorem 3.5, ϕ_2 extends to a homeomorphism between disk bounded by τ'_1 in the sphere and the disk bounded by l'_1 . All subsequent arcs τ_i either lie in the disk bounded by τ'_2 , the disk bounded by τ'_3 , or the disk bounded by τ''_3 . We continue to map these arcs to the disks bounded by l'_2, l'_3 , and l''_3 .

This concludes the description of recursive step. For each subsequent τ_i , we have a disk containing τ_i and we define a homeomorphism ϕ_{i-2} following Case #1, #2 or #3 depending on τ_i . We make one further note on the construction.




Figure 3.10: Construction of l'_3 and l''_3 in Case #3.

Since the curves τ_i were chosen to converge to ω , the curves τ'_i and τ''_i are naturally constrained. In particular, τ_i lies entirely inside the circle of radius $1/2^i$ centred at ω . When considering τ_i we can choose the curves τ'_i and τ''_i so that they lie inside the circle of radius $1/2^i$. We also select the line segments l_i and l'_i carefully.

To start we stipulate that the circular subarcs of C defined by l_1 and l_2 have length $1/4$ of the circumference of C . Now in Case #1, we map τ_3 to l'_1 so there is no choice to make. In Case #2 we chose the segments l_3 and l'_3 so that they both meet the boundary of C at the midpoint of the circular subarc between l_1 and l_2 . In Case #3 we chose the segment l_3 so that the circular subarc defined by l_3 has length equal to half of the length of the circular subarc between l_1 and l_2 , and the subarcs defined by l'_3 and l''_3 have equal length. This guarantees that if there are infinitely many arcs τ_i the area of C not fixed by some ϕ_j approaches zero. Furthermore, the set of endpoints of the segments l_i is a totally disconnected subset of $\text{Bd}(C)$. This follows from the identical argument that appears in Case #3 of the proof of Claim 2.21.

Assume that we have constructed homeomorphisms ϕ_i for each $\tau_i \in A$. We define $h_1 : F \rightarrow B(0, 1)$ as follows:

- If $x \in \tau_1 \cup \tau_2 \cup D_1 \cup D_2$, then $h_1(x) = q^{-1}(\phi_1(q(h(x))))$.
- If $x \in \tau_i \cup D_i$ for $i \notin \{1, 2\}$, then $h_1(x) = q^{-1}(\phi_{i-2}(q(h(x))))$.
- For all other points x , $1/2^j \leq d(q(h(x))) < 1/2^{j-1}$ for some j . We define $h_1(x) = q^{-1}(\phi_k(q(h(x))))$, where k is the least index so that $\phi_k(q(h(x)))$ is defined.

We have that such an index k exists by our previous remarks. Now we show that h_1 is a homeomorphism.

In order to do this, note that our construction gives a partition of two spheres into regions, each of which has an associated homeomorphism ϕ_i . Furthermore, if the regions associated with ϕ_i and ϕ_j meet, ϕ_i and ϕ_j are identical on their shared boundary. Thus the set of functions ϕ_i give us a homeomorphism ϕ from $B(0, 1)$ to itself that maps each τ_i to the segment l_i . Therefore h_1 is the composition of h with ϕ , and is thus a homeomorphism.

We have constructed a homeomorphism $h_1 : B(0, 1) \rightarrow F$ so that the arcs τ_i , in Σ , are mapped onto arcs $h_1^{-1}(\tau_i)$ that have exactly two endpoints on \mathbb{S}^1 . We now construct a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ so that $g|_{B(0, 1)}$ is a homeomorphism. We define g in two steps. First we extend h_1 to a continuous surjection h_2 that is defined on all points of $\text{Cl}(F)$ except for the points in $\text{Bd}(F)$ that are not the end of any τ_i . Then we extend h_2 to g , defined on all of $\text{Cl}(F)$.

Define $A_1 = B(0, 1) \cup X$, where

$$X = \{x \in \mathbb{S}^1 : x \in \text{Bd}(l_i) \text{ for some } i\},$$

and $F_1 = F \cup V$, where

$$V = \{v \in \text{Bd}(F) : v \in \text{Bd}(\tau_i) \text{ for some } i\}.$$

Let γ_i be the subarc of \mathbb{S}^1 defined by the endpoints of l_i so that the simple closed curve $\text{Cl}(l_i \cup \gamma_i)$ encloses an open disk subset of $B(0, 1)$ that contains no other l_j . Define $A_2 = A_1 \cup \{\gamma_i : i \in \mathbb{N}\}$, and $F_2 = F_1 \cup E(\text{Bd}(F))$. We construct $h_2 : A_2 \rightarrow F_2$ as follows.

For each τ_i , $h_1^{-1}(\tau_i) = l_i$, and each l_i is a chord of \mathbb{S}^1 . Consider l_i , and let $e = uv$ be the edge of K corresponding to τ_i . Then we have a homeomorphism $h_e : D_e \rightarrow \text{Cl}(B(0, 1))$. Further, $h_e(e \cup \tau_i)$ bounds an open disk in the plane. There is a

natural homeomorphism between $h_e(e \cup \tau_i)$ and $\text{Cl}(l_i \cup \gamma_i)$ defined by traversing e from u to v then τ_i from v to u while traversing γ_i and l_i correspondingly. Let this homeomorphism be h'_e . By Theorem 3.4, h'_e extends to a homeomorphism from the open disk enclosed by $h_e(e \cup \tau_i)$, contained in $B(0, 1)$, and the open disk enclosed by $\text{Cl}(l_i \cup \gamma_i)$, contained in $B(0, 1)$, D_i .

We define h_2 as,

$$h_2(x) = \begin{cases} h_1(x), & \text{if } x \in D; \\ h_e^{-1}(h_e'^{-1}(x)), & \text{if } x \text{ is in } \text{Cl}(D_i). \end{cases}$$

Note that our definition is slightly ambiguous if $x \in X$. In this case there are two possible edges e and e' that correspond to x . However, it does not matter which edge we choose to map x to, since x is mapped to the shared endpoint of e and e' . Thus we can choose either e or e' and x is mapped to the same point in F_2 . It is important to note that h_2 is still a homeomorphism between $B(0, 1)$ and F . This follows from the same argument that h_1 is a homeomorphism. We have partitioned $B(0, 1)$ into disks D_1, D_2, \dots together with D , and defined a homeomorphism on the closure of each so that they agree on their shared boundaries. Thus $h_2|_{B(0,1)}$ is a homeomorphism, and h_2 is a continuous surjection. We are not guaranteed that h_2 is a homeomorphism, as we may have mapped two of the arcs γ_i to the same edge, and we may have mapped many points in X to the same vertex.

Before we complete the construction we note that we are only missing those points of \mathbb{S}^1 that are not endpoints of any l_i . These correspond to vertices of K in $\text{Bd}(F)$ that are not the endpoint of any edge $e \in E(\text{Bd}(F))$. The definition of g on these points is complicated, but makes the continuity proof easy.

Consider any point $x \in \mathbb{S}^1$. Now consider the concentric open disks $B(x, 1/2^i)$, and set $B_i = B(0, 1) \cap B(x, 1/2^i)$. We have that B_i is an open disk subset of $B(0, 1)$ for each $i \geq 0$, so $h_2(B_i)$ is an open disk subset of F for each $i \geq 0$. Consider the boundary of $h_2(B_i)$.

Proposition 3.17

$\text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$ is a non-empty, closed, connected subset of K for each $i \geq 0$.

Proof We have that $\text{Bd}(B_i) \cap \mathbb{S}^1 \neq \emptyset$ for all $i \geq 0$. In fact, $\text{Bd}(B_i) \cap \mathbb{S}^1$ is a subarc of \mathbb{S}^1 centred at x . Suppose that $\text{Bd}(h_2(B_i)) \cap \text{Bd}(F) = \emptyset$. Then $\text{Bd}(h_2(B_i)) \subset F$

and hence $\text{Bd}(h_2(B_i))$ is a simple closed curve in F . But then $h_2^{-1}(\text{Bd}(h_2(B_i)))$ is a simple closed curve in $B(0, 1)$, a contradiction. Thus $\text{Bd}(h_2(B_i)) \cap \text{Bd}(F) \neq \emptyset$. Since $\text{Bd}(F)$ and $\text{Bd}(h_2(B_i))$ are both closed subsets of Σ , and $\text{Bd}(F) \subseteq K$, it follows that $\text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$ is a closed subset of K .

We need to show that $\text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$ is connected. We use a similar idea as in the proof of Lemma 3.2. Consider the open disks $D_j = B(0, 1 - 1/2^j)$. These are concentric open disks in the plane contained in $B(0, 1)$ for all $j \geq 1$. We note that $D_j \cap B_i$ is either empty or an open disk for all $j \geq 1$. Moreover, $\text{Cl}(B_i) - D_j$ is a closed disk for all $j \geq 1$. Note that $\text{Bd}(B_i) \cap \mathbb{S}^1 = \bigcap_{j \geq 1} (\text{Cl}(B_i) - D_j)$. Now Proposition 2.1 gives us that $\text{Bd}(B_i) \cap \mathbb{S}^1$ is a non-empty, closed connected subset of the plane. Admittedly this is not particularly interesting, but we can use the same strategy in Σ .

We have that $\text{Cl}(h_2(B_i))$ is a closed subset of Σ , and that $\text{Cl}(h_2(B_i)) - h_2(D_j)$ is closed in Σ . This follows from the fact that,

$$\Sigma - (\text{Cl}(h_2(B_i)) - h_2(D_j)) = (\Sigma - \text{Cl}(h_2(B_i))) \cup (h_2(B_i) \cap h_2(D_j))$$

is an open subset of Σ . Further the sets $\text{Cl}(h_2(B_i)) - h_2(D_j)$ are connected, since $h_2(B_i)$ is connected, $h_2(B_i) - h_2(D_j)$ is connected, and

$$\text{Cl}(h_2(B_i)) - h_2(D_j) = \text{Cl}(h_2(B_i) - h_2(D_j)).$$

Thus $\text{Cl}(h_2(B_i)) - h_2(D_j)$ is connected.

Finally, given any $j < j'$

$$\text{Cl}(h_2(B_i)) - h_2(D_j) \supset \text{Cl}(h_2(B_i)) - h_2(D_{j'}).$$

This follows directly from the fact that $B_i - D_j \supset B_i - D_{j'}$. Thus we have by Proposition 2.1 that

$$X_i = \bigcap_{j \geq 1} (\text{Cl}(h_2(B_i)) - h_2(D_j))$$

is a non-empty, closed connected subset of Σ .

But $X_i = \text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$, since if $x \in X_i$ then $x \notin h_2(D_j)$ for any $j \geq 1$. Thus $x \in \text{Bd}(h_2(B_i))$ and $x \in \text{Bd}(F)$. Also, if $x \in \text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$ then $x \in h_2(B_i)$ and $x \notin h_2(D_j)$ for any $j \geq 1$. Therefore we have the desired result. ■

Proposition 3.18

$\cap_{i \geq 0} (\text{Bd}(h_2(B_i)) \cap \text{Bd}(F))$ is a non-empty, closed connected subset of K .

Proof From Proposition 3.17 we have that each $\text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$ is a non-empty, closed connected subset of K . If these sets are nested then we can again apply Proposition 2.1 to derive the result.

We have that if $i > j$ then $B_i \subset B_j$. Also, $\text{Bd}(B_i) \cap \mathbb{S}^1 \subset \text{Bd}(B_j) \cap \mathbb{S}^1$. We prove that these properties also hold in Σ .

We have that if $i > j$ then $h_2(B_i) \subset h_2(B_j)$. Take $y \in \text{Bd}(h_2(B_i)) \cap \text{Bd}(F)$. Each neighbourhood of y contains points of F and $\Sigma - F$, and contains points of $h_2(B_i)$ and $\Sigma - h_2(B_i)$. Thus each neighbourhood of y contains points of F and $\Sigma - F$ and points of $h_2(B_j)$, since $h_2(B_i) \subset h_2(B_j)$, and $\Sigma - h_2(B_j)$, because $h_2(B_j) \subset F$. Therefore

$$\text{Bd}(h_2(B_i)) \cap \text{Bd}(F) \subseteq \text{Bd}(h_2(B_j)) \cap \text{Bd}(F).$$

Now we apply Proposition 2.1 and conclude that $\cap_{i \geq 0} (\text{Bd}(h_2(B_i)) \cap \text{Bd}(F))$ is a non-empty, closed connected subset of K , as required. ■

We need the following theorem from [4]. Given M , a metric space with metric d , we define the distance between S and T for $S, T \subseteq M$ to be $d(S, T) = \inf\{d(s, t) : s \in S, t \in T\}$.

Theorem 3.19 ([4], Thm.3.A.14)

Suppose S and T are disjoint non-empty subsets of a metric space M with metric d . If S is compact and T is closed, then $d(S, T) > 0$. Furthermore, if $d(S, T) = r$ then there is a point $s \in S$ with $d(s, T) = r$.

Proposition 3.20

There is a unique $w \in K$ so that $w = \cap_{i \geq 0} (\text{Bd}(h_2(B_i)) \cap \text{Bd}(F))$.

Proof Assume that

$$X = \cap_{i \geq 0} (\text{Bd}(h_2(B_i)) \cap \text{Bd}(F))$$

contains more than one point. Since X is a connected subset of K , $|X| > 1$ implies that X contains some points in edges of K . Thus there is some $e \in E(\text{Bd}(F))$ so that X contains a subarc of e . Let $[x, y]$ be a subarc of e contained in X . We can add the points x and y to the vertex set of K to obtain a graph-like space that contains $e' = [x, y]$ as an edge. Therefore we can assume that there is an edge $e \in E(K)$ so that $e' \subset X$. Consider the closed disk D_e associated with e as before.

Take two points $a, b \in e$ so that $a \neq b$. Then we have two closed disk neighbourhoods of a and b , U_a and U_b respectively, so that $U_a \cap U_b = \emptyset$. Consider the homeomorphism h_e mapping D_e to a closed disk in the plane. This homeomorphism takes U_a and U_b to disjoint closed disks in the plane. Consider now $V_a = h_2^{-1}(U_a)$ and $V_b = h_2^{-1}(U_b)$. We have two possibilities. Either e appears once or twice in the boundary of F . In the first case, V_a and V_b are closed disk subsets of $\text{Cl}(B(0, 1))$, in the second V_a and V_b are each composed of two closed disk components. In either case, by Theorem 3.19 $d(V_a, V_b) = \epsilon$ for some $\epsilon > 0$, since V_a and V_b are closed. Now take i so that $\epsilon > 1/2^i$, and consider the open disk B_i . We have that for all $s, t \in B_i$, $d(s, t) < 1/2^i < \epsilon$. Thus B_i cannot contain points of both V_a and V_b . Therefore, without loss of generality, $V_a \cap B_i = \emptyset$. But then $a \notin X$, a contradiction. Thus $|X| = 1$, as required. ■

Now we can define $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ as,

$$g(x) = \begin{cases} h_2(x) & \text{if } x \in A_2, \\ w & \text{else.} \end{cases}$$

In this definition w is the unique point from Proposition 3.20.

We now show that g is a continuous surjection.

Lemma 3.21

g is a continuous function.

Proof In order to prove that g is continuous we consider a sequence of points $\{x_i\}$ converging to x in $\text{Cl}(B(0, 1))$. We prove that the sequence $\{g(x_i)\}$ converges to $g(x)$. We consider two cases.

Case #1:

$x \in B(0, 1)$. In this case we note that, for some finite index l , the x_i lie in $B(0, 1)$ for all $i > l$. Thus since g is a homeomorphism on $B(0, 1)$ (as $g|_{B(0,1)} = h_2|_{B(0,1)}$), $\{g(x_i)\}$ converges to $g(x)$, as required.

Case #2:

$x \in \text{Bd}(B(0, 1))$. This case splits naturally into two subcases. Either $g(x) = h_2(x)$, or x is mapped to a vertex of K that is not the end of any edge in $E(\text{Bd}(F))$. Before we address these subcases we make an observation.

Since $\{x_i\}$ converges to x , the distance between x_i and x , d_i , converges to zero. Thus if we consider the open disks $B_i = B(x, d_i) \cap B(0, 1)$, each x_i is in the

boundary of B_i . We show that the same is true of $g(x_i)$ and $g(B_i)$. If $x_i \in B(0, 1)$ this is clear, so assume that $x_i \in \mathbb{S}^1$. If $g(x_i) \notin \text{Bd}(g(B_i))$ then there is a disk D containing $g(x_i)$ so that $(D \cap F) \cap B_i = \emptyset$, but then $g^{-1}(D \cap F) \cap B_i = \emptyset$, which shows that $x_i \notin \text{Bd}(B_i)$. Therefore by Proposition 3.20 $\{g(x_i)\}$ converges to a unique point. It only remains to show that $g(x) = \bigcap_{i \geq 0} \text{Bd}(B_i)$.

In the second subcase, x is mapped to a vertex of K that is not the end of any edge in $E(\text{Bd}(F))$. Now by definition $g(x) = \bigcap_{i \geq 0} \text{Bd}(B_i)$, so $\{g(x_i)\}$ converges to $g(x)$.

In the first subcase, $g(x) = h_2(x)$. We have an arc α in $B(0, 1)$ such that x is one end, and the other end is in $B(0, 1)$. This follows, since if $h_2(x)$ is a vertex, then we take α to be a subarc of the arc τ_i corresponding to an edge incident with $h_2(x)$. If $h_2(x)$ is the interior point of an edge e , then we take α to be an arbitrary arc in the image under h_2^{-1} of the edge disk corresponding to e . Now let $\tau = h_2(\alpha)$. By construction we can choose a sequence $\{y_i\}$ of points on α so $y_i \in B_i - B_{i-1}$. This sequence gives a corresponding sequence $\{g(y_i)\}$. Since α and τ are arcs with endpoint x and $h_2(x)$ respectively, $\{y_i\}$ converges to x , and $\{g(y_i)\}$ converges to $g(x)$. Thus $g(x) = \bigcap_{i \geq 0} \text{Bd}(B_i)$.

We conclude that g is continuous. ■

Before we prove that g is surjective we make the following observation about neighbourhoods of points in $\text{Bd}(F)$.

Proposition 3.22

If $w \in \text{Bd}(F)$ and B is an open disk containing w , then $B \cap F$ is a non-empty collection of open disks, at least one of which contains w in its boundary.

Proof Since B is open, and F is open, $B \cap F$ is open. We also have that each connected component of $B \cap F$ is arcwise connected. Furthermore, consider D , a connected component of $B \cap F$. Suppose we have $\sigma \subset D$, a simple closed curve. Then σ is a simple closed curve in F and a simple closed curve in B . Since B and F are open disks, σ encloses an open disk in both B and F . Further, in the surface Σ , σ encloses an open disk. Thus these disks are all the same and σ encloses an open disk in D , and D is simply connected. Therefore D is an open disk. Furthermore, since every neighbourhood of w contains points of F , $B \cap F \neq \emptyset$. Finally, w is in the boundary of at least one component of $B \cap F$. This follows, since otherwise there is an open neighbourhood of w that is disjoint

from $\text{Bd}(B \cap F)$. However each open neighbourhood of w contains points of F , B and $B \cap F$. Therefore we have the desired result. ■

Lemma 3.23

g is a surjection.

Proof It suffices to show that g is a surjection from \mathbb{S}^1 to $\text{Bd}(F)$. Indeed, it suffices to show that if $x \in \text{Bd}(F) \cap V(K)$ is not in the closure of any edge $e \in E(\text{Bd}(F))$, then there is some $y \in \mathbb{S}^1$ so that $g(y) = x$.

Consider the open disks $B(x, 1/2^i)$ for $i \geq 0$ in Σ . We have by Proposition 3.22 that $B(x, 1/2^i) \cap F$ is a non-empty collection of open disks for each $i \geq 0$. Thus we can choose a point $x_i \in B(x, 1/2^i) \cap F$ for each $i \geq 0$. Now we have a sequence of points $\{x_i\}$ in F that converge to x . Therefore we can consider the set $\{g^{-1}(x_i)\}$. Since this is an infinite set of points in $B(0, 1)$ it has a convergent subsequence $\{y_i\}$ converging to $y \in \text{Cl}(B(0, 1))$ (note that $\{g^{-1}(x_i)\}$ may have many accumulation points, y is one of them). We must have that $y \in \mathbb{S}^1$ since otherwise there is an open disk neighbourhood V of y so that $\text{Bd}(V) \cap \mathbb{S}^1 = \emptyset$. But then there is some finite l so that $B(x, 1/2^i) \cap g(V) = \emptyset$ for all $i > l$. This contradicts the choice of the points x_i . Thus $y \in \mathbb{S}^1$. But now by Lemma 3.21 the sequence $\{y_i\}$ converges to y implies that the sequence $\{g(y_i)\}$ converges to $g(y)$. Therefore $g(y) = x$ and g is a surjection. ■

This completes the construction of g . The material in this section proves the following theorem.

Theorem 3.24

If K is a connected graph-like space embedded in the surface Σ , and F is a face of the embedding, then there is a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ so that $g|_{B(0, 1)}$ is a homeomorphism.

3.4 Connectedness and Embeddings

In the preceding sections we have assumed that G is a connected graph-like space embedded in Σ . Now we consider graph-like spaces that are not connected.

Assume that G is a graph-like space that is not connected. Then by Proposition 2.12 G has finitely many connected components, G_1, \dots, G_n . Now we let the embedding of G in Σ be K , with components K_1, \dots, K_n . We define the faces

of K in the same way as for connected graph-like spaces. However, if K is not connected, then the faces of K are not guaranteed to be open disks. The faces of K are the open, arcwise connected components of $\Sigma - K$.

Following our discussion in Section 3.1 we can assume that if F is a face of K , and $\text{Bd}(F)$ is entirely contained in some K_i , then F is an open disk. Thus a face F of K is an open disk if and only if $\text{Bd}(F)$ is connected. Furthermore, we may assume that if F is not an open disk, then F is an open disk with holes. Therefore F is homeomorphic to $B(0, 1) - \cup_{i=1}^k B_i$ where each B_i is an open disk, $\text{Cl}(B_i) \subset B(0, 1)$, $\text{Cl}(B_i) \cap \text{Cl}(B_j) = \emptyset$ for $i \neq j$, and each $\text{Bd}(B_i)$ lies in a distinct component K_j of K . We refer to these faces as *non-disk faces*.

Note that in Section 3.2 we did not need to assume that our graph-like space was connected. Thus the results from Section 3.2 hold for an embedding K of G . In particular, each edge of G has a neighbourhood homeomorphic to an open disk in Σ . Further, we have by Corollary 3.11 that each edge e appears in the boundary of one or two faces, and $e \cap \text{Bd}(F)$ is either empty or all of e for each face F of K .

In order to prove Theorem 3.24 we needed to assume that F was an open disk. For the faces of K that are open disks, Theorem 3.24 still applies. For the faces of K that are not open disks, we have a similar result. If F is a non-disk face, then for each boundary component of F , C_i , there is a simple closed curve $\sigma \subset F$ so that the component of $\Sigma - \sigma$ containing C_i contains none of the other components C_j . We can create a new surface Σ' by identifying the boundary of a closed disk with σ . Thus in Σ' we have an embedding of some of the components of K , and a new face F' corresponding to F . In Σ' the face F' is an open disk, so we can apply Theorem 3.24 to find a continuous surjection $g_i : \mathbb{S}^1 \rightarrow C_i$. We can use this procedure for each boundary component C_i of F . We have proved the following result.

Theorem 3.25

If K is a graph-like space embedded in the surface Σ , and F is a face of the embedding homeomorphic to an open disk, then there is a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ so that $g|_{B(0, 1)}$ is a homeomorphism. If F is a non-disk face of the embedding then for each component, C_i , of $\text{Bd}(F)$ there is a continuous surjection $g_i : \mathbb{S}^1 \rightarrow C_i$.

Chapter 4

The Thin Cycle Space and the Face Boundary Space

In this chapter we consider the algebraic edge space of a graph-like space. In particular we consider two subspaces of the algebraic edge space, the thin cycle space, and the face boundary space of an embedding. We will use the theory of topological edge spaces developed by Vella and Richter in [19] (see also [3]).

A *topological edge space* (X, E) is a topological space X together with a subset $E \subset X$ of points e so that $\{e\}$ is open, not closed, and has at most two boundary points. The point $e \in E$ is an edge of (X, E) , and any point in $X - E$ is a vertex. The results we wish to apply from [19] concern connected, compact weakly Hausdorff topological edge spaces. A space is *weakly Hausdorff* if for any two points x, y , there are neighbourhoods U_x, U_y of x, y respectively so that $U_x \cap U_y$ is finite.

In Section 4.1 we discuss topological edge spaces and algebraic edge spaces. We will show that there is a natural correspondence between graph-like spaces and topological edge spaces. This will allow us to apply the results of Vella and Richter when considering algebraic edge spaces. In Section 4.3 we will prove that the face boundaries of a graph-like space embedded in a surface generate a subspace of its algebraic edge space. Finally, in Section 4.5 we will consider the cycle space of a graph-like space. We will apply Lemma 2.13 to prove that the face boundary space of a graph-like space embedded in a surface is a subspace of its cycle space.

4.1 Algebraic Edge Spaces and Topological Edge Spaces

Topological edge spaces are not graph-like spaces. In a topological edge space, an edge is a single point rather than an arc, and the vertices do not necessarily form a zero-dimensional set.

For example, if $G = (V, E)$ is a finite connected graph, then we can construct a topological edge space from G as follows. Set $X = V \cup E$. We define the topology on X by specifying the basic open sets. For $e \in E$, $\{e\}$ is a basic open set, and for each $v \in V$,

$$\{v\} \cup \{e \in E : e \text{ is incident with } v\}$$

is a basic open set. These sets generate the topology on X , and (X, E) is a connected, compact weakly Hausdorff topological edge space.

However, we can also define a connected, compact weakly Hausdorff edge space (Y, E) by taking a closed disk for each vertex. We set $Y = E \cup \{B_v : v \in V\}$ where each B_v is a closed disk. Now if $e = uv$, then we specify points x_u, x_v in B_u, B_v respectively so that $\text{Bd}(e) = \{x_u, x_v\}$. The basic open sets are $\{e\}$ for each $e \in E$, together with the open sets for each vertex $v \in B_u$. If v is not in the closure of any edge, then the basic open sets corresponding to v are the neighbourhoods of v in B_u that do not contain the endpoint of any edge. If v is in the closure of some edge, then the basic open sets corresponding to v are of the form $N \cup \{e \in E : v \in \text{Bd}(e)\}$, where N is any neighbourhood of v in B_u that does not contain the endpoint of any edge not incident with v . The resulting topological edge space (Y, E) bears little resemblance to the original graph. However, the main point is that graph-theoretic connection in G corresponds to topological connection in (Y, E) .

Given a connected graph-like space G , we can construct a connected, compact weakly Hausdorff topological edge space $X(G) = (X, E)$ as follows. We take $X = E(G) \cup (G - E(G))$, and $E = E(G)$ (recall that $e \in E(G)$ is an arc in G). We define the topology on X by specifying the basic open sets of X . For each $e \in E$ we take $\{e\}$ to be a basic open set. For each $v \in V$ consider a neighbourhood N of v in G . For each N we take N' to be a basic open set corresponding to v where

$$N' = (V \cap N) \cup \{e \in E : e \cap N \neq \emptyset\}.$$

We have that $X(G)$ is connected, and compact since G is connected and compact. It remains to show that $X(G)$ is weakly Hausdorff. Suppose that $x, y \in X$. If x is an edge we take $U_x = \{x\}$, and U_y to be any neighbourhood of y . Then $U_x \cap U_y$ is finite. Assume that $x, y \in X - E$. Then x and y are vertices of G , and we have neighbourhoods N_x and N_y of x and y respectively so that $N_x \cap N_y = \emptyset$, since G is Hausdorff. Furthermore, since G is metrizable $d(x, y) = \epsilon$ for some $\epsilon > 0$, and we can take $N_x = B(x, \epsilon/4)$ and $N_y = B(y, \epsilon/4)$. Now we claim that $U_x = N'_x$ and $U_y = N'_y$ are neighbourhoods so that $U_x \cap U_y$ is finite. By definition, $U_x \cap U_y$ is the set of edges $e \in E$ so that $e \cap N_x \neq \emptyset$ and $e \cap N_y \neq \emptyset$. Thus by Lemma 2.5, since there are only finitely many edges with diameter greater than $\epsilon/2$, $U_x \cap U_y$ is finite.

We follow the development of the thin cycle space of a topological edge space in [19]. We define the *algebraic edge space* of a topological edge space as the collection of all formal linear combinations $\sum_{e \in E} \alpha_e e$, where each $\alpha_e \in GF(2)$, the finite field with two elements. This collection forms a vector space with addition and scalar-multiplication defined componentwise. From this point on we refer to algebraic edge spaces as *edge spaces*, and if (X, E) is a topological edge space, we refer to the corresponding edge space as 2^E .

If $s \in 2^E$ we define the *support* of $s = \sum_{e \in E} \alpha_e e$ to be the set $\{e \in E : \alpha_e \neq 0\}$. A collection $S \subset 2^E$ is *thin* if for each $e \in E$, there are only finitely many elements $s \in S$ so that e is in the support of s . Given a thin collection S of elements of 2^E we can take the ‘‘symmetric difference’’ of the elements in S . Formally, we can compute $\sum_{s \in S} s = \sum_{e \in E} \beta_e e$ where $\beta_e = 1$ if e is in the support of an odd number of elements of S and $\beta_e = 0$ otherwise. We refer to $\sum_{s \in S} s$ as the *thin sum* of S .

For each subset $S \subseteq 2^E$, S generates three, possibly distinct, subsets of 2^E . The *weak span* of S , $\mathcal{W}(S)$, is the set of all symmetric differences of finite subsets of S . The *algebraic span* of S , $\mathcal{A}(S)$, is the set of all symmetric differences of thin subsets of S . The *strong span* of S , $\mathcal{S}(S)$, is the smallest subset of 2^E that contains S and is closed under *thin summation* (i.e. the thin sum of A is in $\mathcal{S}(S)$ for all thin subsets A of S). A subset $S \subseteq 2^E$ is a *subspace* of 2^E if $S = \mathcal{S}(S)$. Note that in order to prove that S is a subspace of 2^E , it suffices to show that S is closed under thin summation. Given two subspaces S, S' of 2^E , we say that S is a *subspace* of S' if $S \subseteq S'$.

An *edge cycle* is a connected topological edge space (X, E) so that for each $e, f \in E$, $X - e$ is connected, and $X - \{e, f\}$ is not connected. A *cycle* is the edge

set $E(C)$ of any edge cycle C . The *thin cycle space* of (X, E) , $\mathcal{Z}_t(X, E)$, is the strong span of the set of cycles of (X, E) . We have the following results from [19].

Theorem 4.1 ([19], Thm. 14)

Let (X, E) be a compact, weakly Hausdorff topological edge space. Then every element of the thin cycle space is the disjoint union of cycles.

Corollary 4.2 ([19], Cor. 15)

Let (X, E) be a connected, compact, weakly Hausdorff topological edge space. Then the thin cycle space is the algebraic span of the set of cycles of (X, E) .

We are also able to apply the following result from [19].

Theorem 4.3 ([19], Thm. 12)

Given a compact, weakly Hausdorff topological edge space (X, E) and a partition of $X - E$ into two closed sets P, Q , there are only finitely many edges that have one end in P and the other end in Q .

Thus, even though (X, E) may have infinitely many edges, and edge cycles of infinite size, the edge-cuts of X are finite.

Now suppose that G is a graph-like space. For a graph-like space G we define a *cycle* to be connected graph-like subspace C of G so that for all $x, y \in C$, $C - x$ is connected and $C - \{x, y\}$ is not connected. Note that this differs from the definition of an edge cycle above. However, a cycle C in a graph-like space corresponds to an edge cycle in the associated topological edge space. From Theorem 35 in [19] each cycle in G is homeomorphic to \mathbb{S}^1 . Recall that we defined a topological edge space $X(G)$ associated with G . Note that there is a bijection between the cycles of G and the edge cycles of $X(G)$. Thus 2^E is the edge space for both G and $X(G)$, and the thin cycle space $\mathcal{Z}_t(X(G))$ is the strong span of the set of cycles of G . We let $\mathcal{Z}_t(G)$ be the thin cycle space of G , and note that $\mathcal{Z}_t(G) = \mathcal{Z}_t(X(G))$. Furthermore, Theorem 4.1 and Corollary 4.2 hold for $\mathcal{Z}_t(G)$.

4.2 A Property of Surfaces

Before we proceed with the main results of this chapter, we need to prove a property of surfaces. In this section we prove that if Σ is a surface, and $V \subset \Sigma$ is totally disconnected, then $\Sigma - V$ is arcwise connected.

In [12], Richards gives a complete topological classification of non-compact triangulable surfaces. Richards takes *surface* to mean a connected 2-manifold. This differs from our definition in that surfaces can be non-compact, and can have boundaries. The main point of interest to us is the concept of an *ideal boundary*. Richards defines the ideal boundary as follows.

A subset A of a surface Σ is *bounded* if the closure of A is compact in Σ . A *boundary component* of a surface Σ is a nested sequence $P_1 \supset P_2 \supset \dots$ of connected unbounded regions in Σ such that:

1. the boundary of P_n in Σ is compact for all n ;
2. for any bounded subset A of S , $P_n \cap A = \emptyset$ for n sufficiently large.

Two boundary components, $P_1 \supset P_2 \supset \dots$ and $P'_1 \supset P'_2 \supset \dots$ are equivalent if, for any n , there is some N such that $P_N \subset P'_n$ and vice versa. If we p denote the boundary component $P_1 \supset P_2 \dots$, then p^* denotes the equivalence class of boundary components containing p . We call p^* an *ideal boundary point*.

The *ideal boundary* $B(\Sigma)$ of a surface Σ is the topological space consisting of the ideal boundary points p^* of Σ with the following topology. For any set U in Σ whose boundary in Σ is compact, the set U^* is the set of all ideal boundary points p^* so that for $p \in p^*$, $P_N \subset U$ for large enough N . We take the set of all such U^* as a basis for the topology of $B(S)$. Furthermore, if p^* is an ideal boundary point of Σ , then p^* is *planar* if the sets P_n are planar for sufficiently large n . We define p^* to be *orientable* if the sets P_n are orientable for sufficiently large n .

We consider the ideal boundary of Σ to be a nested triple of sets $B \supset B' \supset B''$ where $B = B(\Sigma)$, B' is the part of $B(\Sigma)$ that is not planar and B'' is the part of $B(\Sigma)$ that is not orientable. Richards' main result is the following theorem.

Theorem 4.4 ([12], Thm. 1)

Let Σ and Σ' be two separable surfaces of the same genus and orientability class. Then Σ and Σ' are homeomorphic if and only if their ideal boundaries (considered as triples of spaces) are topologically equivalent.

We use the development of Richards to prove the following topological lemma. We make use of Theorem 4.4 implicitly in the proof.

Lemma 4.5

Let Σ be a compact, connected surface and let V be a totally disconnected compact subset of Σ . Then $\Sigma - V$ is connected.

Proof Let $S = \Sigma - V$. The space S is a non-compact surface. We refer to the connected components of S as the faces of V in Σ .

We employ Richards' classification of non-compact surfaces. By construction, S has finite genus and, therefore, every ideal boundary point is planar. Thus, the ideal boundary of S is the triple B, B', B'' , where $B = V$, and $B' = B'' = \emptyset$.

Each component of S is a non-compact surface. A punctured disk D' is $D - V'$ where D is an open disk and V' is a totally disconnected subset of D . Given a component of S with ideal boundary point p^* , for $p \in p^*$ we have that $P_1 \supset P_2 \supset \dots$ is a nested sequence of punctured disks whose intersection is empty. These punctured disks translate in Σ to disks contained in $F \cup V$, where F is the face of V in question.

If S is not connected, then some point of V is in the boundary of more than one face of V . Suppose that F and F' share a boundary point $p = p'$. Then there are two sets of disks, $P_1 \supset P_2 \supset \dots$ and $P'_1 \supset P'_2 \supset \dots$, in $F \cup V$ and $F' \cup V$ corresponding to the ideal boundary points, p and p' , of F and F' respectively. For each i and each j , $P_i \cap P'_j \subset V$ is a totally disconnected subset of Σ . However, P_i and P_j are both open disks. Thus $P_i \cap P_j$ is an open subset of Σ , and hence each connected component of $P_i \cap P_j$ is open. This is a contradiction since each connected component of $P_i \cap P_j$ is a single point. ■

Lemma 4.5 has a simple corollary that we make use of here and in Chapter 5.

Corollary 4.6

Given K a graph-like space embedded in surface Σ , $\Sigma - V$ is arcwise connected.

Proof The set $V = K - E$ is a totally disconnected, compact subset of Σ . We have from Lemma 4.5 that $\Sigma - V$ is connected. We also have that $\Sigma - V$ is open. Thus by Proposition 3.1, $\Sigma - V$ is arcwise connected. ■

4.3 The Face Boundary Space

In this section we demonstrate that the face boundaries of an embedded graph-like space in a surface form a subspace of the edge space.

Suppose that we have a set $K \subseteq \Sigma$ for a surface Σ such that K is an embedding of a graph-like space G . Then K is compact and every closed connected subset

of K is arcwise connected. We have $V \subseteq K$ such that $K - V$ consists of disjoint open arcs having one or two endpoints in K . We denote $K - V = E$ to be the set of edges of K . Recall that the faces of K are the connected (and hence arcwise connected) components of $\Sigma - K$. Here we stipulate that K be a subset of Σ such that each face is homeomorphic to an open disk. This implies that K is connected.

In Chapter 3 we showed that if each face is an open disk, then for each face F we have a closed curve τ in K such that $\tau = \text{Bd}(F)$. We also have an implicit, arbitrary traversal associated with τ which we call the boundary walk of F in K . Further, if $e \in E(K)$, then either $e \subset \text{Bd}(F)$ or $e \cap \text{Bd}(F) = \emptyset$.

For a face F of K we define $E(F)$ to be the multi-set of edges contained in the boundary of F where $e \in E(F)$ appears as many times in $E(F)$ as it does in the boundary walk of F in K . Recall that a cycle in a graph-like space is a subset $C \subseteq K$ with the property that if $x, y \in C$, then $C - x$ is connected and $C - \{x, y\}$ is not connected. We have that Theorem 4.1 and Corollary 4.2 apply to K , independently of Σ .

Now suppose that F is a face of K . Then we take $\text{Bd}^*(F) \in 2^E$ to be the member of the edge space over $GF(2)$ defined by

$$\text{Bd}^*(F) := \sum_{e \in E(F)} e \pmod{2}.$$

We take $\mathcal{B}_t(K)$ to be the subset of 2^E consisting of all thin sums of faces and we take $\mathcal{Z}_t(K)$ to be the thin cycle space as before. Note that since G and K are homeomorphic, $\mathcal{Z}_t(K)$ and $\mathcal{Z}_t(G)$ are isomorphic. We now work towards the main result of this section.

The set $\mathcal{B}_t(K)$ is defined as the subset of 2^E consisting of all thin sums of faces. Thus the set of faces of K is a generating set for $\mathcal{B}_t(K)$. In order to prove that $\mathcal{B}_t(K)$ is a subspace of 2^E we need to show that $\mathcal{B}_t(K)$ is closed under thin summation.

Let \mathcal{F} denote the set of faces of K . Consider $a_1, a_2, \dots \in \mathcal{B}_t(K)$, a thin collection of elements of $\mathcal{B}_t(K)$. Then for each a_i we have a subset $\mathcal{F}_i \subseteq \mathcal{F}$ of faces such that $a_i = \sum_{F \in \mathcal{F}_i} \text{Bd}^*(F)$ (from here onward we take all sums modulo 2).

Since every edge appears in either one or two face boundaries, each edge appears in either zero or two elements of 2^E that correspond to elements of \mathcal{F} .

Thus $\sum_{F \in \mathcal{F}} \text{Bd}^*(F) = \emptyset$, and

$$a_i = \sum_{F \in \mathcal{F}_i} \text{Bd}^*(F) = \sum_{f \in \mathcal{F} - \mathcal{F}_i} \text{Bd}^*(F).$$

Note that we can view each \mathcal{F}_i as a colouring of \mathcal{F} where the \mathcal{F}_i and the $\mathcal{F} - \mathcal{F}_i$ make up the colour classes. Now the a_i are monochromatic sums over the 2-colourings \mathcal{F}_i . In order to show that $\sum_{i \in \mathbb{N}} a_i \in \mathcal{B}_t(K)$ we use this collection of 2-colourings to define a 2-colouring of \mathcal{F} whose monochromatic faces sum to $\sum a_i$.

We can choose an arbitrary face of K , F_0 , and stipulate that $F_0 \notin \mathcal{F}_i$ for all i . This follows since if $F_0 \in \mathcal{F}_i$ we simply replace \mathcal{F}_i with $\mathcal{F} - \mathcal{F}_i$. Furthermore, if F_0 is the only face of K , then every edge of K appears exactly twice in $\text{Bd}(F_0)$ and $\mathcal{B}_t(K) = \{\emptyset\}$. Now since each $\{a_i\} \subseteq \mathcal{B}_t(K)$ is a thin subset of 2^E , each $e \in E$ appears in a_i for finitely many i . The set \mathcal{F}_i gives us a 2-colouring σ_i of the faces of K , where σ_i is defined by,

$$\sigma_i(F) = \begin{cases} 0 & \text{if } F \in \mathcal{F} - \mathcal{F}_i, \\ 1 & \text{if } F \in \mathcal{F}_i, \end{cases}$$

for each i . Now we show that for each $f \in \mathcal{F}$, $\sigma_i(F) = 1$ for only finitely many i .

Consider an arbitrary face F . Take points $x \in F$ and $y \in F_0$. By Corollary 4.6 there is an arc α from x to y in $\Sigma - V$. We claim that α intersects only finitely many edges of K . Suppose otherwise. Let $\{e_i\}$ be an infinite sequence of edges of K so that each e_i is distinct, $e_i \cap \alpha \neq \emptyset$ and if $i < j$ then α intersects e_i before e_j in the traversal from x to y . Now let $\{z_i\}$ be a sequence of points so that $z_i \in e_i \cap \alpha$ for each i . Then $\{z_i\}$ is an infinite sequence, and has a convergent subsequence $\{z'_i\}$ that converges to z' . Since K and α are both compact, $z' \in K \cap \alpha$. Thus $z' \in e$ for some edge e of K , because $\alpha \subset \Sigma - V$ and $z' \in K \cap \alpha$. Since z' is an interior point of an edge, there is some neighbourhood N of z' so that $N \cap V = \emptyset$. Since the points z'_i converge to z' , N and every neighbourhood of z' contained in N contains infinitely many of the z'_i . Therefore z' has no connected neighbourhood contained in N . This contradicts the local connectedness of K at z' . Thus α intersects only finitely many edges of K . Note that α may intersect K at infinitely many points, and $\alpha \cap e$ may not be totally disconnected for some e .

For faces F and F' of K in Σ , we say that F is *adjacent* to F' if for arbitrary

points $x \in F$ and $y \in F'$ there is an arc α from x to y in $\Sigma - V$ so that $\alpha \cap K \subset e$ for some edge e . Note that the above discussion demonstrates that for any two faces F and F' of K , there is a chain of faces $F = F_1, F_2, \dots, F_n = F'$ so that F_i is adjacent to F_{i+1} for each $i = 1, \dots, n-1$. This follows, since if $x \in F$, $y \in F'$ and α is any arc from x to y in $\Sigma - V$, then α intersects only finitely many edges of K . Now if $E' \subset E$ is the set of edges that α intersects, and \mathcal{F} is the set of faces of K , then we can consider the set $S = (\cup_{F \in \mathcal{F}} F) \cup E'$. There is a connected component C of S that contains α , and all other connected components of S are faces of K . The component C consists of the edges E' together with a finite set of faces of K , \mathcal{F}' . Thus, the faces \mathcal{F}' together with the given definition of adjacency, define a finite connected graph G . There is a path from F to F' in G , $F = F_1, F_2, \dots, F_n = F'$. Therefore we have the desired chain. This point will be useful in Chapter 5, so we note that we have proven the following proposition.

Proposition 4.7

Given a graph-like space K embedded in surface Σ , for any two faces F and F' of K , there is a chain of faces $F = F_1, F_2, \dots, F_n = F'$ so that F_i is adjacent to F_{i+1} for each $i = 1, \dots, n-1$.

Now we prove that for each $F \in \mathcal{F}$, $\sigma_i(F) = 1$ for only finitely many i . For every i , $\sigma_i(F_0) = 0$ by definition. If F' is any face of K , then there is a sequence of faces $F_0, F_1, \dots, F_n = F'$ so that each pair of consecutive faces is adjacent. We prove by induction on j that $\sigma_i(F_j) = 1$ for only finitely many i . This is trivial for $j = 0$, so suppose that $j > 0$, and $\sigma_i(F_{j-1}) = 1$ for only finitely many i . Since F_{j-1} and F_j are adjacent faces, there is an edge $e \in \text{Bd}(F_{j-1}) \cap \text{Bd}(F_j)$. The sets a_i form a thin family, so e appears in only finitely many of the a_i . For each i so that $e \notin a_i$, $\sigma_i(F_{j-1}) = \sigma_i(F_j)$. Thus since $\sigma_i(F_{j-1}) = 1$ for only finitely many i , $\sigma_i(F_j) = 1$ for only finitely many i . Therefore $\sigma_i(F') = 1$ for only finitely many i , for any face F' of K .

Thus we can define a 2-colouring of \mathcal{F} by

$$\sigma(F) = \begin{cases} 1 & \text{if } \sigma_i(F) = 1 \text{ for an odd number of indices } i, \\ 0 & \text{else.} \end{cases}$$

Note that the condition $\sigma_i(F) = 1$ for an odd number of indices i is the same as the condition that $F \in \mathcal{F}_i$ for an odd number of indices i . Also note that $\sigma(F_0) = 0$. We claim that this 2-colouring gives us $\sum a_i = \sum_{\sigma(F)=1} \text{Bd}^*(F)$.

Given $F_i, F_j \in \mathcal{F}$, we define

$$I_{ij} := \{k : F_i \in \mathcal{F}_k, F_j \notin \mathcal{F}_k, \text{ or } F_i \notin \mathcal{F}_k, F_j \in \mathcal{F}_k\}.$$

Note that since both F_i and F_j are in finitely many \mathcal{F}_k , I_{ij} is a finite set. Also note that if $i = j$ then $I_{ij} = \emptyset$. We denote equivalence modulo 2 by $x \equiv_2 y$.

Proposition 4.8

For any $F_1, F_2, F_3 \in \mathcal{F}$ we have $|I_{12}| + |I_{13}| + |I_{23}| \equiv_2 0$.

Proof We have three cases to consider. Either all of the faces are distinct, two faces are the same or all of the faces are the same.

Case #1:

$$F_1 = F_2 = F_3.$$

In this case $I_{12} = I_{13} = I_{23} = \emptyset$ so $|I_{12}| + |I_{13}| + |I_{23}| \equiv_2 0$ trivially.

Case #2:

$$F_1 = F_2 \neq F_3.$$

In this case we have $I_{12} = \emptyset$ and we want to show that $|I_{13}| + |I_{23}| \equiv_2 0$. But since $F_1 = F_2$ we have that $I_{13} = I_{23}$. Thus $|I_{13}| + |I_{23}| \equiv_2 0$ trivially.

Case #3:

The faces F_1, F_2, F_3 are distinct.

Note that

$$I_{12} \cap I_{13} = \{k : F_1 \in \mathcal{F}_k, F_2, F_3 \notin \mathcal{F}_k \text{ or } F_1 \notin \mathcal{F}_k, F_2, F_3 \in \mathcal{F}_k\}.$$

Therefore, $I_{12} \cap I_{13} \cap I_{23} = \emptyset$.

Likewise, $I_{12} \subseteq I_{13} \cup I_{23}$, since $k \in I_{12}$ implies $\sigma_k(F_1) \neq \sigma_k(F_2)$, and thus $\sigma_k(F_3)$ cannot be equal to both $\sigma_k(F_1)$ and $\sigma_k(F_2)$. Since I_{12} is disjoint from $I_{13} \cap I_{23}$, we deduce that $I_{12} \subseteq I_{13} \Delta I_{23}$.

Conversely, suppose $k \in I_{13} \Delta I_{23}$. Then precisely one of $\sigma_k(F_1)$ and $\sigma_k(F_2)$ is equal to $\sigma_k(F_3)$, so obviously $\sigma_k(F_1) \neq \sigma_k(F_2)$, hence $k \in I_{12}$. Thus, $I_{12} = I_{13} \Delta I_{23}$.

It is a standard fact that $|A \Delta B| \equiv_2 |A| + |B|$, so we conclude that

$$\begin{aligned} |I_{12}| &\equiv_2 |I_{13}| + |I_{23}|, \text{ or} \\ 0 &\equiv_2 |I_{12}| + |I_{13}| + |I_{23}|, \end{aligned}$$

as required. ■

We can now prove the main result of this section.

Lemma 4.9

If K is a connected graph-like space embedded in surface Σ , then $\mathcal{B}_t(K)$ is a subspace of 2^E .

Proof We apply Proposition 4.8 to show that σ is a 2-colouring of \mathcal{F} with $\sum a_i = \sum_{\sigma(F)=1} \text{Bd}^*(F)$. Recall that we have fixed $F_0 \in \mathcal{F}$ so that $\sigma_i(F_0) = 0$ for all i .

Suppose that $e \in E$ separates F_j from F_k (we only consider edges that are in the boundaries of two faces, since edges that are in the boundary of only one face vanish in all sums). If $e \in \sum a_i$ then $e \in a_i$ for an odd number of indices i , so $\sigma_i(F_j) \neq \sigma_i(F_k)$ for an odd number of indices i . Now if $j = 0$, then $\sigma(F_j) \neq \sigma(F_k)$ by definition. If $j, k \neq 0$, then $|I_{jk}| \equiv_2 1$. But

$$\begin{aligned} |I_{jk}| &\equiv_2 |I_{0j}| + |I_{0k}|, \quad \text{and} \\ |I_{jk}| &\equiv_2 \sigma(F_j) + \sigma(F_k) \equiv_2 1. \end{aligned}$$

Thus $\sigma(F_j) \not\equiv_2 \sigma(F_k)$ and thus $\sigma(F_j) \neq \sigma(F_k)$ as required.

Now suppose that $e \notin \sum a_i$. Then $e \in a_i$ for an even number of indices i , so $\sigma_i(F_j) = \sigma_i(F_k)$ for an even number of indices i . If $j = 0$, then $\sigma(F_1) = \sigma(F_2)$ by definition. If $j, k \neq 0$, then $|I_{jk}| \equiv_2 0$. But

$$\begin{aligned} |I_{jk}| &\equiv_2 |I_{0j}| + |I_{0k}|, \quad \text{and} \\ |I_{jk}| &\equiv_2 \sigma(F_j) + \sigma(F_k) \equiv_2 0. \end{aligned}$$

Thus $\sigma(F_j) \equiv_2 \sigma(F_k)$ and thus $\sigma(F_j) = \sigma(F_k)$ as required.

Therefore

$$\sum_{\sigma(F)=0} \text{Bd}^*(F) = \sum_{\sigma(F)=1} \text{Bd}^*(F) = \sum a_i.$$

Thus $\sum a_i \in \mathcal{B}_t(K)$ and $\mathcal{B}_t(K)$ is closed under thin summation. Therefore $\mathcal{B}_t(K)$ is a subspace of 2^E . ■

Note that since $\mathcal{B}_t(K)$ is a subspace of 2^E , we can define $\mathcal{B}_t(K)$ as the strong span of the set of face boundaries of K . Lemma 4.9 gives us the first step towards proving that $\mathcal{B}_t(K)$ is a subspace of $\mathcal{Z}_t(K)$.

4.4 The Face Boundaries of a Disconnected Graph-Like Space

In this section we will extend Lemma 4.9 to disconnected graph-like spaces. If G is a graph-like space, then by Proposition 2.12 G has finitely many connected components G_1, \dots, G_n . Furthermore, each G_i is a connected graph-like space. We consider an embedding K of G in surface Σ , and let K_i be the resulting embedding of G_i for each i . Following the development in Section 3.4 we stipulate that every non-disk face of K is homeomorphic to $B(0, 1) - \cup_{i=1}^k B_i$ where each B_i is an open disk, $\text{Cl}(B_i) \subset B(0, 1)$, $\text{Cl}(B_i) \cap \text{Cl}(B_j) = \emptyset$ for $i \neq j$, and each $\text{Bd}(B_i)$ lies in a distinct component K_j of K .

Now by Theorem 3.25 if F is a face of K homeomorphic to an open disk, then there is a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ so that $g|_{B(0, 1)}$ is a homeomorphism. We take $g_1 = g$. If F is a non-disk face of the embedding, then for each component, C_i , of $\text{Bd}(F)$ there is a continuous surjection $g_i : \mathbb{S}^1 \rightarrow C_i$. Thus for any face F we have a set of continuous surjections $\{g_i\}$ from \mathbb{S}^1 to the components of $\text{Bd}(F)$. We also have that if $e \in E(K)$, then either $e \subset \text{Bd}(F)$ or $e \cap \text{Bd}(F) = \emptyset$.

For a face F of K , we define $E(F)$ as follows. Let C_1, \dots, C_k be the components of $\text{Bd}(F)$. Each continuous surjection g_i associated with F gives a traversal of a component C_i of $\text{Bd}(F)$. For each C_i we let E_i be the multi-set of edges in $E(C_i)$ so that each edge $e \in E(C_i)$ appears as many times in E_i as it does in the boundary walk g_i of C_i . Now $E(F)$ is the multi-set $E(F) = \cup_{i=1}^k E_i$. Then we take $\text{Bd}^*(F) \in 2^E$ to be the member of the edge space over $GF(2)$ defined by

$$\text{Bd}^*(F) := \sum_{e \in E(F)} e \pmod{2}.$$

The set $\mathcal{B}_t(K)$ is defined as the subset of 2^E consisting of all thin sums of faces. Thus the set of faces of K is a generating set for $\mathcal{B}_t(K)$.

Furthermore, each K_i is a connected graph-like space embedded in Σ with a finite number of non-disk faces. Thus we can construct an embedding of K_i in surface Σ' by replacing each non-disk face of K_i with a disk. Let K'_i be the resulting embedding. Now the development in Section 4.3 applies to K'_i embedded in Σ' . However, the face boundaries of K'_i in Σ' are equal to the face boundaries of K_i in Σ as elements of 2^E . Thus following the definitions in Section 4.3 we have face boundary spaces $\mathcal{B}_t(K_i)$ for each K_i .

If F is a face of K , then for each K_i , F is contained in a face F_i of K_i . Thus given a face F of K we have a set

$$\begin{aligned} I_F &= \{i : \text{Bd}(F) \cap K_i \neq \emptyset\}, \quad \text{and} \\ F &= \sum_{i \in I_F} F_i \in 2^E. \end{aligned}$$

Lemma 4.10

If K is a graph-like space embedded in surface Σ , then $\mathcal{B}_t(K)$ is a subspace of 2^E .

Proof We prove that $\mathcal{B}_t(K)$ is closed under thin summation. Consider $a_1, a_2, \dots \in \mathcal{B}_t(K)$, a thin collection of elements of $\mathcal{B}_t(K)$. For each a_i we have a subset $\mathcal{F}_i \subseteq \mathcal{F}$ of faces such that $a_i = \sum_{F \in \mathcal{F}_i} \text{Bd}^*(F)$.

Let K_1, \dots, K_n be the connected components of K . Consider $1 \leq j \leq n$. For each face F we have a face F_j of K_j so that $F \subseteq F_j$. For each $i \in \mathbb{N}$ define

$$a_i^j = \sum_{F \in \mathcal{F}_i} \text{Bd}^*(F_j),$$

so $a_i^j \in \mathcal{B}_t(K_j)$. Since $\{a_i\}$ is a thin collection of elements of $2^{E(K)}$, each e appears in finitely many of the a_i . Since $a_i^j \subseteq a_i$, e appears in finitely many of the a_i^j . Therefore, $\{a_i^j\}$ is a thin collection of elements of $\mathcal{B}_t(K_j)$. Since K_j is connected, we have by Lemma 4.9 that

$$a^j = \sum a_i^j \in \mathcal{B}_t(K_j).$$

Thus there is a set \mathcal{F}^j of faces of K_j so that $a^j = \sum_{F \in \mathcal{F}^j} \text{Bd}^*(F)$.

Consider $a = \sum_{j=1}^n a^j$. For each $e \in E(K)$, $e \in a$ if and only if e appears in an odd number of the elements a^j . However, since the subspaces K_j partition $E(K)$, there is exactly one j so that $e \in E(K_j)$. Thus e appears in a if and only if e appears in a^j . Now e appears in a^j if and only if e appears in an odd number of the elements a_i^j . For $e \in E(K_j)$, $e \in a_i^j$ if and only if $e \in a_i$. Thus e lies in an odd number of the a_i^j if and only if e lies in an odd number of the a_i . Therefore, $e \in a$ if and only if $e \in \sum a_i$. Thus

$$a = \sum_{j=1}^n a^j = \sum a_i.$$

In order to show that $a \in \mathcal{B}_t(K)$ we define a set \mathcal{F}^* so that $a = \sum_{F \in \mathcal{F}^*} \text{Bd}^*(F)$. We define \mathcal{F}^* by considering the elements a^j of $\mathcal{B}_t(K_j)$. We have that $a^j = \sum_{F \in \mathcal{F}^j} \text{Bd}^*(F)$. Note that, as before, since the set of all faces of K_j sums to \emptyset , there are two possibilities for \mathcal{F}^j . If we let \mathcal{S}^j be the set of faces of K_j , then

$$a^j = \sum_{F \in \mathcal{F}^j} \text{Bd}^*(F) = \sum_{F \in \mathcal{S}^j - \mathcal{F}^j} \text{Bd}^*(F).$$

For each K_j , each K_i with $i \neq j$ lies entirely in a face of K_j . Thus the faces of K_j partition the remaining K_i . Moreover, for each face F of K_j that is not an open disk, there is a set $I(F, j)$ so that, for each $i \in I(F, j)$, K_i lies in F , and $\text{Bd}(F) \cap K_i \neq \emptyset$. We use these facts to define colourings of the faces of each K_j .

For each $1 \leq j \leq n$ we define a colouring $\sigma_j : \mathcal{S}^j \rightarrow \{0, 1\}$ so that the monochromatic faces of K_j sum to a^j . We define the functions σ_j inductively. Let $\sigma_1(F) = 1$ if $F \in \mathcal{F}^1$ and $\sigma_1(F) = 0$ if $F \notin \mathcal{F}^1$. Now consider a face F of K_1 that is not an open disk. We use σ_1 to define σ_i for each $i \in I(F, 1)$. If $i \in I(F, 1)$, then F contains K_i and $\text{Bd}(F) \cap K_i \neq \emptyset$. Furthermore, F corresponds to a face F_i of K_i that contains K_1 with $\text{Bd}(F_i) \cap K_1 \neq \emptyset$. We let $\sigma_i(F_i) = \sigma_1(F)$. Now define $\sigma_i(F') = \sigma_i(F_i)$ if and only if F' and F_i lie in the same side of the partition $\mathcal{F}^i, \mathcal{S}^i - \mathcal{F}^i$. We repeat this procedure for each $i \in I(F, 1)$ and each non-disk face F of K_1 . This concludes the first inductive step. We now repeat this procedure for each σ_i that has already been defined, considering only those K_j for which σ_j has not been defined. Since there are only finitely many K_j , this process is well-defined and terminates after a finite number of steps.

Now we redefine \mathcal{F}^j to be $\mathcal{F}^j = \{F : \sigma_j(F) = 1\}$. We still have that $a^j = \sum_{F \in \mathcal{F}^j} \text{Bd}^*(F)$. Note that if F is an open disk face of K , then F is an open disk face of exactly one K_j , and if F is not an open disk face of K , then $\sigma_i(F)$ is constant over all indices i so $\text{Bd}(F) \cap K_i \neq \emptyset$. Thus we can define \mathcal{F}^* using the sets \mathcal{F}^j . Let \mathcal{F} be the set of faces of K . For $F \in \mathcal{F}$, $F \in \mathcal{F}^*$ if and only if $F \in \mathcal{F}^j$ for some $1 \leq j \leq n$. Now we prove that $a = \sum_{F \in \mathcal{F}^*} \text{Bd}^*(F)$.

Suppose that $e \in a$. Then $e \in a^j$ for some j . Thus e bounds two distinct faces of K_j , exactly one of which is in \mathcal{F}^j . Therefore, e bounds two distinct faces of K . Note that each face F of K that is bounded by multiple components K_j corresponds to faces with the same colour in each K_j . Since e bounds exactly one face of K_j that lies in \mathcal{F}^j , e bounds exactly one face of K that lies in \mathcal{F}^* .

Suppose that $e \in \sum_{F \in \mathcal{F}^*} \text{Bd}^*(F)$. Let K_j be the component of K so that $e \in$

$E(K_j)$. Now e bounds two distinct faces F and F' of K and exactly one of F, F' lies in \mathcal{F}^* . Without loss of generality let $F \in \mathcal{F}^*$. The faces F and F' correspond to faces F_j and F'_j of K_j respectively. By definition of \mathcal{F}^* , $F \in \mathcal{F}^*$ if and only if $F_j \in \mathcal{F}^j$. Thus $F_j \in \mathcal{F}^j$ and $F'_j \notin \mathcal{F}^j$. Therefore, $e \in a^j$, and $e \in a$.

Now $\sum a_i = a = \sum_{F \in \mathcal{F}^*} \text{Bd}^*(F)$, and $\sum a_i \in \mathcal{B}_t(K)$. We conclude that $\mathcal{B}_t(K)$ is closed under thin sums, and $\mathcal{B}_t(K)$ is a subspace of 2^E . ■

Note that in proving Lemma 4.10 we have also proved that $\mathcal{B}_t(K)$ is the direct sum of the spaces $\mathcal{B}_t(K_i)$ for $1 \leq i \leq n$.

4.5 The Thin Cycle Space

In this section we will demonstrate that given a graph-like space G that is embedded in a surface Σ , the face boundaries of the embedding generate a subspace of the thin cycle space. Let K be the embedding of G in Σ .

For each face F that is homeomorphic to an open disk we have a continuous surjection $g : \text{Cl}(B(0, 1)) \rightarrow \text{Cl}(F)$ so that g is a homeomorphism between $B(0, 1)$ and F and a continuous surjection between \mathbb{S}^1 and $\text{Bd}(F)$. The function g restricted to the unit circle gives us a traversal of $\text{Bd}(F)$. If F is a non-disk face of the embedding then, for each component C_i of $\text{Bd}(F)$, there is a continuous surjection $g_i : \mathbb{S}^1 \rightarrow C_i$. Each function g_i gives us a traversal of a component of $\text{Bd}(F)$. If F is homeomorphic to an open disk, then we take $g_1 = g$. Thus for any face F we have a set of functions $\{g_i\}$ each of which is a continuous surjection between the circle and a component of $\text{Bd}(F)$.

For a face F , we want to break g_i into homeomorphisms between circles and subsets of $\text{Bd}(F)$ in order to show that for each component C_i of $\text{Bd}(F)$, $C_i \in \mathcal{Z}_t(K)$. Naturally we would like to apply Lemma 2.13. However in order to apply Lemma 2.13 directly we need the points $v \in \text{Bd}(F)$ so that $|g_i^{-1}(v)| > 1$ to be a totally disconnected subset of $\text{Bd}(F)$. If we have an edge $e \in E(\text{Bd}(F))$ so that F occurs on both sides of e then $g_i^{-1}(e)$ consists of two subarcs of \mathbb{S}^1 , so we cannot claim that the points visited multiple times by g_i form a totally disconnected set.

Instead we apply Lemma 2.13 to a graph-like space K' obtained from K . Recall that in Section 3.3 we considered the closed disks D_e for each $e \in E(\text{Bd}(F))$. We needed to ensure that the arcs τ_i arising from the disks D_e formed a closed subset

of Σ . We were able to achieve this by forcing the images of the arcs $q(h(\tau_i))$ to converge to ω in the sphere. At the time we mentioned that we were essentially re-choosing the disks D_e so that their boundaries formed a closed subset of Σ . Our methods give us the following lemma, which shows that for each edge e , we may assume that e lies in the boundary of two distinct faces.

Lemma 4.11

Given a graph-like space K embedded in the surface Σ and a face F of K , we can construct a new graph-like space K' from K by adding two new edges, e', e'' , to K for every edge $e \in E(\text{Bd}(F))$ that appears twice in the boundary of F . Furthermore, e' and e'' define two new faces, F' and F'' so that $e \in \text{Bd}(F') \cap \text{Bd}(F'')$.

Proof Suppose that F is a face of K that is homeomorphic to an open disk. We choose closed disks D_e for each $e \in E(\text{Bd}(F))$. We let τ_i be the arcs formed by the boundaries of these disks. We map the τ_i into the sphere by $h \circ q$. We rechoose the $q(h(\tau_i))$ so that these curves converge to ω as in the proof of Proposition 3.15. Then we re-define the D_e in turn.

The arcs τ_i in Σ are disjoint from each other and $E(K)$. They each have two endpoints in $V(K)$. Define the set of arcs Φ as the subset of $\{\tau_i : i \in \mathbb{N}\}$ so that $\tau_i \in \Phi$ if and only if τ_i corresponds to an edge e that appears twice in the boundary of the face F . The set $K' = K \cup \Phi$ is a closed and hence compact subset of Σ . Finally Proposition 3.15 gives us that $\Sigma - K'$ is a collection of open disks. Thus K' is a graph-like space embedded in Σ containing K .

Now suppose that F is not homeomorphic to an open disk. Then $\text{Bd}(F)$ has finitely many components corresponding to the graph-like spaces K_1, \dots, K_l . For each space K_i , there is a face F_i so that $F \subset F_i$. In the surface Σ_i formed by replacing F_i with an open disk, we can apply the above argument to obtain an embedding of K'_i in Σ so that for every edge $e \in E(\text{Bd}(F_i))$ that appears twice in the boundary of F_i we have new edges e' and e'' in K'_i . We can perform this procedure for each K_i independently. Furthermore, since each $\text{Bd}(F_i)$ is compact, we can choose the new edges so that they only intersect in vertices of K . ■

Note that the edges of K' are the edges of K plus the added arcs τ_i . We have added two arcs for each edge e that appears twice in the boundary of the face F .

Suppose that C is a cycle in K' that passes through an edge $\tau_e \notin E(K)$, and that τ_e corresponds to $e \in E(K)$. If $e \notin C$, then $(C - \tau_e) \cup e$ is a cycle in K' .

Furthermore, if $\tau_e, \tau'_e \in E(K')$ both correspond to $e \in E(K)$, then $\tau_e, \tau'_e \in C$ implies that $E(C) = \{\tau_e, \tau'_e\}$. To see this, let e be the edge between u and v . Now if we have some other edge $e' \in E(C)$, then $C - u$ is connected so there is an arc α from v to e' . But, since $C - \{a, a'\}$ is disconnected for any a, a' in the interior of τ_e, τ'_e , α passes through either a or a' . This is a contradiction. Thus $E(C) = \{\tau_e, \tau'_e\}$. We refer to such cycles in K' as *trivial cycles*.

For non-trivial cycles, C , let $A(C) = E(C) \cap \Phi$, and let $B(C) \subset E(K)$ be the set of edges in K so that $e \in B(C)$ if and only if there is some arc $\tau \in A(C)$ corresponding to e . If C is not trivial, then

$$C' = (C - A(C)) \cup B(C)$$

is a cycle in K . In order to prove this, note that cycles are homeomorphs of circles, and there is a homeomorphism, $\phi : \mathbb{S}^1 \rightarrow C$. Furthermore, for each edge $e \in E(C)$, $\phi^{-1}(e)$ is a subarc of the circle. In order to replace τ_e with e , we simply change ϕ on $\phi^{-1}(\tau_e)$ to the natural homeomorphism onto e so that the ends of e correspond to the ends of τ_e . Now ϕ is still a homeomorphism. Since the edges we are replacing are disjoint, we can perform this operation on each $\tau_e \in A(C)$. Thus we have a homeomorphism between the circle and C' in K , and C' is a cycle.

Let the components of K' be K'_1, \dots, K'_n , where K_1, \dots, K_n are the components of K . For each face F of K that is an open disk, we have some K_i so that $\text{Bd}(F) \subseteq K_i$. Thus there is an open disk, D , defined by Proposition 3.15 (i.e. D is the open disk contained in F so that the arcs τ_i are all in the boundary of D). The disk D is contained in a unique face F' of K' by construction. Furthermore, $\text{Bd}(F') \subseteq K'_i$ for some i . For each face F of K that is not an open disk, the same argument shows that there is a face F' of K' so that F and F' are homeomorphic. Furthermore, $\text{Bd}(F) \cap K_i \neq \emptyset$ if and only if $\text{Bd}(F') \cap K'_i \neq \emptyset$.

Thus we have a set $\{g'_i\}$ where each g'_i is a continuous surjection from \mathbb{S}^1 to a component of $\text{Bd}(F')$. We also have that each edge $e \in E(\text{Bd}(F'))$ has the property that F' appears on exactly one side of e . Therefore the repeated points of $g'^{-1}(\text{Bd}(F'))$ are a subset of $g'^{-1}(V(\text{Bd}(F')))$ which is a totally disconnected subset of Σ . Thus we can apply Lemma 2.13 to each g'_i . The result is a continuous surjection $f : C^* \rightarrow \text{Bd}(F')$ such that $f|_C$ is a continuous injection for each $C \in C^*$.

For $C \in C^*$, we have three possibilities. Firstly, we may have that $f(x) = v$ for all $x \in C$. This case is an artifact of the recursive procedure we used to prove Lemma 2.13 and can be ignored when considering $2^{E(K)}$. Secondly, we may have that $E(f(C))$ contains τ_e and τ'_e where both correspond to the same edge e of K . In this case $f(C)$ is a trivial cycle, and $E(f(C)) = \{\tau_e, \tau'_e\}$. Note that while $f(C)$ does not correspond to a cycle in K , $f(C)$ corresponds to an edge e that appears twice in the traversal of $\text{Bd}(F)$. Thus e does not appear in the element $\text{Bd}^*(F) \in 2^{E(K)}$. Thirdly, we may have $C \in C^*$ so that $f(C)$ is a non-trivial cycle. These cycles in K' correspond to cycles of K by the observation above that we can swap edges of K for their replacement arcs in K' .

We let C_1 be the collection of circles that map to a single vertex, C_2 be the collection of circles that map to trivial cycles, and C_3 be the collection of circles that map to non-trivial cycles. These are disjoint collections of circles, and $C^* = C_1 \cup C_2 \cup C_3$. Let $\text{Bd}^*(F')$ and $\text{Bd}^*(F)$ be the elements of $2^{E'}$ and 2^E corresponding to $E(\text{Bd}(F'))$ and $E(\text{Bd}(F))$ respectively. It is clear that

$$\begin{aligned} \text{Bd}^*(F') &= \sum_{C \in C_1} f(C) + \sum_{C \in C_2} f(C) + \sum_{C \in C_3} f(C) \\ &= \sum_{C \in C_2} f(C) + \sum_{C \in C_3} f(C). \end{aligned}$$

Let C'_3 be the subset of C_3 where $C \in C'_3$ if $E(f(C)) \subset E(K)$, and let $C''_3 = C_3 - C'_3$. For $C \in C''_3$ we let $\rho(f(C))$ be the cycle of K derived from $f(C)$ by replacing the K' edges in $f(C)$ with their corresponding K edges.

Proposition 4.12

In $2^{E(K)}$, $\text{Bd}^*(F) = \sum_{C \in C'_3} f(C) + \sum_{C \in C''_3} \rho(f(C))$.

Proof We have that each edge $e \in \text{Bd}^*(F')$ occurs exactly once in the traversal of $\text{Bd}(F')$. Thus if $e \in \text{Bd}^*(F')$, then e is in exactly one cycle in C_2 , or C_3 . Therefore each edge $e \in E(K)$ appears 0, 1 or 2 times in the summation over C'_3 and C''_3 .

If e does not appear in the summation, then either $e \notin \text{Bd}^*(F')$ and thus $e \notin \text{Bd}^*(F)$, or e corresponds to a trivial cycle, in which case $e \notin \text{Bd}^*(F)$. If e appears once in the summation, then $e \in E(K)$, so $e \in \text{Bd}^*(F)$. If e appears twice in the summation, then $e \in \rho(f(C)), \rho(f(C'))$ for some $C, C' \in C''_3$. Thus e appears twice in the traversal of $\text{Bd}(F)$ and $e \notin \text{Bd}^*(F)$.

Now, if $e \in \text{Bd}^*(F)$, then e appears once in the traversal of $\text{Bd}(F)$, so $e \in \text{Bd}^*(F')$. Thus e appears in exactly one element of C_3 . Therefore e appears once

in the summation, and the equality holds. ■

We have proven the following theorem.

Theorem 4.13

For a graph-like space K embedded in surface Σ , $\mathcal{B}_t(K)$ is generated by the set of cycles of K , and $\mathcal{B}_t(K)$ is a subspace of $\mathcal{Z}_t(K)$.

Chapter 5

MacLane's Theorem

MacLane's Theorem gives an algebraic characterization of finite planar graphs. If G is a finite graph, then we take $\mathcal{C}(G)$ to be the subspace of $2^{E(G)}$ containing all even degree subgraphs of G . We call $\mathcal{C}(G)$ the cycle space of G , as $\mathcal{C}(G)$ is generated by the set of cycles of G . Given any subspace X of $2^{E(G)}$ we call $D \subseteq X$ a 2-basis for X if D generates X , D is independent, and for all $e \in E(G)$ we have that e lies in at most two elements of D .

Theorem 5.1 (MacLane 1937)

A finite graph G is planar if and only if $\mathcal{C}(G)$ has a 2-basis [5].

To prove this theorem one notes that, given a plane graph, the face boundaries generate the cycle space. Then one shows that the Kuratowski graphs K_5 and $K_{3,3}$ do not have 2-bases. Thus MacLane's Theorem is equivalent to Kuratowski's Theorem for finite graphs.

We wish to generalize MacLane's Theorem to include graph-like spaces. In this chapter we will show that MacLane's Theorem holds for 2-connected graph-like spaces. We follow the same strategy as in the proof for finite graphs. In Section 5.1 we show that $\mathcal{B}_t(K) = \mathcal{Z}_t(K)$ for an embedding K of a graph-like space G in the plane. In Section 5.2 we show that $\mathcal{B}_t(K)$ has a 2-basis for any embedding K of G in some surface Σ . This proves the forward direction of MacLane's Theorem. Finally we reduce the backward direction of MacLane's Theorem to Kuratowski's Theorem.

5.1 Face Boundaries in the Plane

In this section we reconsider the face boundary space of a graph-like space G embedded in the plane. Let K be a planar embedding of G . Then by Theorem 4.13 $\mathcal{B}_t(K) \subseteq \mathcal{Z}_t(K)$. In order to prove that $\mathcal{B}_t(K) = \mathcal{Z}_t(K)$ we will apply the Jordan Curve Theorem.

The Jordan Curve Theorem is a standard topological result.

Theorem 5.2 (The Jordan Curve Theorem)

Let C be a simple closed curve in the plane. Then C separates the plane into precisely two components W_1 and W_2 . Each of the sets W_1 and W_2 has C as its boundary [10].

Note that the Jordan-Schönflies Theorem is a stronger version of this result. However, this version will suit our needs in this chapter.

Recall from Chapter 4 that a cycle in a graph-like space G is a homeomorph of the circle. The thin cycle space of G , $\mathcal{Z}_t(G)$ is the strong span of the edge sets of cycles of G . Further, if G is embedded in Σ as K , then $\mathcal{Z}_t(G)$ and $\mathcal{Z}_t(K)$ are isomorphic.

Theorem 5.3

If G is a graph-like space and K is an embedding of G in the plane, then $\mathcal{B}_t(K) = \mathcal{Z}_t(K)$.

Proof First, we show that $\mathcal{B}_t(K) \supseteq \mathcal{Z}_t(K)$. Since $\mathcal{Z}_t(K)$ is the strong span of edge sets of cycles in K and $\mathcal{B}_t(K)$ is closed under thin summation, it suffices to show that if C is a cycle in K , then $E(C) \in \mathcal{B}_t(K)$.

Consider any cycle C in K . We have that C is a homeomorph of a circle in K . Thus C is a simple closed curve in the plane. Now by the Jordan Curve Theorem, C separates the plane into two components, each with C as its boundary. Let these components be W_1 and W_2 .

Now consider the faces of K . Each face is a connected component of $\mathbb{R}^2 - K$. If F is a face, then $F \subset W_i$ for some $i \in \{1, 2\}$. Thus the sets W_1 and W_2 partition the set \mathcal{F} of faces of K . Furthermore, if $e \in E(K)$ then e is in exactly one of W_1, W_2 and $E(C)$.

For a face F of K we let $\text{Bd}^*(F) \in 2^{E(K)}$ denote the set $E(\text{Bd}(F))$. Let \mathcal{F}_1 be the set of faces of K contained in W_1 and \mathcal{F}_2 be the set of faces of K contained in W_2 .

Since $\mathcal{F}_1 \cap \mathcal{F}_2 = \emptyset$ and $\mathcal{F}_1 \cup \mathcal{F}_2 = \mathcal{F}$, $\sum_{F \in \mathcal{F}_1} \text{Bd}^*(F) = \sum_{F \in \mathcal{F}_2} \text{Bd}^*(F)$. We show that $\sum_{F \in \mathcal{F}_1} \text{Bd}^*(F) = E(C)$.

Suppose that $e \in \sum_{F \in \mathcal{F}_1} \text{Bd}^*(F)$. We have that e is contained in W_1, W_2 or $E(C)$. If $e \subset W_2$, then e is not in the boundary of F for any $F \in \mathcal{F}_1$, so e does not appear in the sum. If $e \subset W_1$, then by symmetry, since $\sum_{F \in \mathcal{F}_1} \text{Bd}^*(F) = \sum_{F \in \mathcal{F}_2} \text{Bd}^*(F)$ we can consider the sum to be over the faces of \mathcal{F}_2 . Thus since e is not in the boundary of any $F \in \mathcal{F}_2$, e does not appear in the sum. Therefore, $e \in E(C)$.

Suppose that $e \in E(C)$. Then by Lemma 3.3 we have an open disk U_e so that $U_e \cap K = e$. Now $U_e - K$ consists of two open disks, D_1 and D_2 contained in $\mathbb{R}^2 - K$. Thus $D_1 \subset F_1$ and $D_2 \subset F_2$ for some faces F_1 and F_2 of K . Furthermore, $F_1 \neq F_2$, since C separates D_1 from D_2 , so $F_1 \subset W_1$ and $F_2 \subset W_2$. Thus e appears in exactly one face in \mathcal{F}_1 . Therefore $e \in \sum_{F \in \mathcal{F}_1} \text{Bd}^*(F)$.

Therefore, $\mathcal{B}_t(K) \supseteq \mathcal{Z}_t(K)$. By Theorem 4.13 $\mathcal{B}_t(K) \subseteq \mathcal{Z}_t(K)$. Thus $\mathcal{B}_t(K) = \mathcal{Z}_t(K)$ as required. \blacksquare

Theorem 5.3 relies heavily on the Jordan Curve Theorem, which is a special property of the plane. If Σ is any surface other than the plane, and K is an embedding of G in Σ , then $\mathcal{B}_t(K)$ and $\mathcal{Z}_t(K)$ are not necessarily the same. For example consider the graph-like space G' defined as follows. Let C_i be the circle of radius $1/2^i$ centred at the point v . For C_i let a_i, b_i, c_i, d_i be points on C_i in that clockwise order. Now for each $i \geq 1$ we have edges between $a_i, a_{i+1}, b_i, b_{i+1}, c_i, c_{i+1}$ and d_i, d_{i+1} . Finally we have an edge between a_1 and c_1 and an edge between b_1 and d_1 . The space G' can be embedded in the torus, as shown in Figure 5.1. Note that the cycle through $\{a_i : i \in \mathbb{N}\} \cup \{c_i : i \in \mathbb{N}\}$ is not generated by elements of the face boundary space.

5.2 2-Bases and MacLane's Theorem

Given a subspace \mathcal{A} of the edge space 2^E of a graph-like space, a set $B \subset \mathcal{A}$ is a *generating set* of \mathcal{A} if for each $a \in \mathcal{A}$, a can be expressed as a thin sum over the elements of B . A generating set B of \mathcal{A} is a *basis* of \mathcal{A} if for any thin subset $\emptyset \neq B' \subseteq B$, if $\sum_{b \in B'} \alpha_b b = \emptyset$, then each $\alpha_b = 0$. A basis B of \mathcal{A} is called a *2-basis* if for each $e \in E$, e occurs at most twice in the elements of B .

Now if we consider K a graph-like space embedded in some surface Σ we have by definition that $\mathcal{B}_t(K)$, the face boundary space, is the set of all thin sums




Figure 5.1: Graph-like space G' embedded in the torus.

of face boundaries of K . Furthermore, by Lemma 4.10, $\mathcal{B}_t(K)$ is a subspace of $2^{E(K)}$ and is closed under thin summation. We have by Corollary 3.11 that each edge is either in the boundary of one face, in which case it appears twice in its boundary, or in the boundary of two distinct faces. Thus if we consider the set \mathcal{F} of face boundaries of K , each edge either appears twice or does not appear in the elements of \mathcal{F} . Let $\mathcal{B} = \{\text{Bd}^*(F) : F \in \mathcal{F}\}$. Now \mathcal{B} generates $\mathcal{B}_t(K)$, but is not a basis for $\mathcal{B}_t(K)$. We are summing elements of the edge space of K , so we consider the occurrences of each edge modulo 2. Since each edge appears either 2 times or 0 times, we have $\sum_{F \in \mathcal{F}} \text{Bd}^*(F) = \emptyset$. Thus

$$\text{Bd}^*(F') = \sum_{F \in \mathcal{F} - \{F'\}} \text{Bd}^*(F)$$

for all $F' \in \mathcal{F}$. Therefore if $F \in \mathcal{F}$ is any face of K , then $\mathcal{B} - \{\text{Bd}^*(F)\}$ is a basis for $\mathcal{B}_t(K)$.

Lemma 5.4

Let K be an embedding of a graph-like space in a surface Σ . If $F \in \mathcal{F}$, then $\mathcal{B}' = \mathcal{B} - \{\text{Bd}^*(F)\}$ is a 2-basis for $\mathcal{B}_t(K)$.

Proof We have that \mathcal{B} generates $\mathcal{B}_t(K)$ and \mathcal{B}' generates $\text{Bd}^*(F)$. Since we are working with thin sums as opposed to arbitrary sums, we cannot directly claim that \mathcal{B}' generates $\mathcal{B}_t(K)$. However, suppose for $x \in \mathcal{B}_t(K)$ we have that $x = \sum_{b \in \mathcal{B}} \alpha_b b$ where α_b is a non-negative integer. Let α_F be the coefficient corresponding to $\text{Bd}^*(F)$. If $\alpha_F = 0$, then $x = \sum_{b \in \mathcal{B}'} \alpha_b b$, and if $\alpha_F \neq 0$ then,

$$\begin{aligned} x &= \sum_{b \in \mathcal{B}} \alpha_b b \\ &= \alpha_F \text{Bd}^*(F) + \sum_{b \in \mathcal{B}'} \alpha_b b \\ &= \sum_{b \in \mathcal{B}'} \alpha_F b + \sum_{b \in \mathcal{B}'} \alpha_b b \\ &= \sum_{b \in \mathcal{B}'} (\alpha_F + \alpha_b) b. \end{aligned}$$

But for each b , $\alpha_F + \alpha_b$ is a non-negative integer, and thus can be reduced modulo 2. Thus \mathcal{B}' generates $\mathcal{B}_t(K)$. It remains to show that \mathcal{B}' is independent.

Towards a contradiction, suppose that \mathcal{B}' is not a basis of $\mathcal{B}_t(K)$. Then there is some thin subset $\emptyset \neq B' \subseteq \mathcal{B}'$ such that $\sum_{b \in B'} \alpha_b b = \emptyset$ and not all of the α_b are zero. Let $A = \{b \in B' : \alpha_b = 1\}$. Then we have that $\sum_{a \in A} a = \emptyset$. By definition $\sum_{b \in \mathcal{B}} b = \emptyset$, and $A \subseteq \mathcal{B}' \subset \mathcal{B}$. Thus,

$$\begin{aligned} \emptyset &= \sum_{b \in \mathcal{B}} b, \\ \emptyset &= \sum_{a \in A} a + \sum_{b \in \mathcal{B}-A} b, \\ \emptyset &= \sum_{b \in \mathcal{B}-A} b. \end{aligned}$$

Therefore A and $A' = \mathcal{B} - A$ partition \mathcal{B} into “dependent sets.” If we consider A and A' , each $e \in E$ either appears twice or does not appear in the elements of A and the elements of A' .

By Proposition 4.7 we have, for any faces F and F' of K , a finite chain of faces $F = F_1, F_2, \dots, F_n = F'$ so that F_i is adjacent to F_{i+1} for each $i = 1, \dots, n-1$. Thus we can take any face F so that $\text{Bd}^*(F) \in A$, and any face F' so that $\text{Bd}^*(F') \in A'$, and consider the chain $F = F_1, F_2, \dots, F_n = F'$. There is a least index i so that $\text{Bd}^*(F_i) \in A'$, and so $\text{Bd}^*(F_{i-1}) \in A$. However, F_{i-1} and F_i are adjacent, so there is an edge $e \in E(K)$ so that $e \in \text{Bd}(F) \cap \text{Bd}(F')$. Therefore, $e \in \sum_{a \in A} a$, and

$e \in \sum_{a' \in A'} a'$. This contradicts $\sum_{a \in A} a = \sum_{a' \in A'} a' = \emptyset$. Thus \mathcal{B}' is a 2-basis for $\mathcal{B}_t(K)$. ■

Lemma 5.4 has the following corollary.

Corollary 5.5

If K is a graph-like space embedded in the plane, then $\mathcal{Z}_t(K)$ has a 2-basis.

Proof If K is a graph-like space embedded in the plane then by Theorem 5.3 $\mathcal{Z}_t(K) = \mathcal{B}_t(K)$. Thus a 2-basis of $\mathcal{B}_t(K)$ is a 2-basis of $\mathcal{Z}_t(K)$. By Lemma 5.4, $\mathcal{B}_t(K)$ has a 2-basis. Therefore we have the desired result. ■

Corollary 5.5 gives us half of MacLane's Theorem. In order to prove the other direction, we need to apply Kuratowski's Theorem.

A topological space X is *2-connected* if, for each $p \in X$, $X - p$ is a connected subspace of X . In [17], Thomassen proves the following version of Kuratowski's Theorem.

Theorem 5.6 ([17], Thm. 4.3)

Let M be a locally connected, 2-connected, compact topological space. Then M is embeddable in the sphere if and only if M is metrizable, and contains neither of the Kuratowski graphs K_5 and $K_{3,3}$.

Note that this differs from Kuratowski's Theorem for finite graphs, in that we do not need to consider minors. The reason is that in the standard proof of Kuratowski's Theorem one uses the fact that if a graph G contains a K_5 or $K_{3,3}$ minor, then it contains a subdivision of either K_5 or $K_{3,3}$. In topological spaces we are only concerned with points, and arcs between those points. If we consider a graph as a topological object, we can ignore vertices of degree 2.

Suppose that H is a finite graph. Recall that we have a natural graph-like space X associated with H . If G is a graph-like space, then we say that H is a *subgraph* of G if there is a subspace $K \subseteq G$ so that X is homeomorphic to K . Now we argue that if G is a graph-like space and $\mathcal{Z}_t(G)$ has a 2-basis, then so do all subgraphs of G . We follow the proof of Lemma 17 given by Bruhn and Stein in [2]. For a finite graph H we use the standard theory of cycle spaces, and denote the cycle space of H by $\mathcal{C}(H)$.

Lemma 5.7

Let G be a graph-like space such that $\mathcal{Z}_t(G)$ has a 2-basis, and let H be a subgraph of G . Then $\mathcal{C}(H)$ has a 2-basis.

Proof We let X be the graph-like space derived from H , and we let K be a homeomorphic copy of X in G . Note that the homeomorphism between X and K gives us a bijection between the edges of H and the arcs in G that form the edges of X . Thus if we identify the arc α in X with the edge set $E(\alpha)$, then $\mathcal{C}(H)$ is isomorphic to $\mathcal{Z}_t(K)$.

We may assume that $\mathcal{Z}_t(K) \neq \{\emptyset\}$, as then $\mathcal{Z}_t(K)$ trivially has a 2-basis. For each $Z \in \mathcal{Z}_t(K)$ let $\mathcal{B}_Z \subseteq \mathcal{B}$ be the set so that $Z = \sum_{b \in \mathcal{B}_Z} b$. Since $\mathcal{Z}_t(K)$ is finite there are $Z \in \mathcal{Z}_t(K)$ so that \mathcal{B}_Z is inclusion-wise minimal. Let us denote these by Z_1, \dots, Z_k .

Consider a $D \in \mathcal{Z}_t(K)$ with $\mathcal{B}_D \cap \mathcal{B}_{Z_i} \neq \emptyset$ for some i . We claim that $\mathcal{B}_{Z_i} \subseteq \mathcal{B}_D$. First, note that

$$C := \sum_{B \in \mathcal{B}_D \cap \mathcal{B}_{Z_i}} B \in \mathcal{Z}_t(K).$$

Indeed, consider an edge $e \notin E(K)$. Since Z_i and D are subsets of K and since \mathcal{B} is a 2-basis, e either lies in exactly two or in none of the cycles of \mathcal{B}_{Z_i} , and the same holds for \mathcal{B}_D . Furthermore, if e lies in two cycles of \mathcal{B}_{Z_i} and in two of \mathcal{B}_D , these must be the same, so $e \notin E(C)$.

Therefore, $C \in \mathcal{Z}_t(K)$. Since $\mathcal{B}_C \subseteq \mathcal{B}_{Z_i}$ we obtain, by the minimality of \mathcal{B}_{Z_i} , that $C = Z_i$. Consequently, $\mathcal{B}_{Z_i} = \mathcal{B}_C \subseteq \mathcal{B}_D$, as claimed.

This result also implies $\mathcal{B}_{Z_i} \cap \mathcal{B}_{Z_j} = \emptyset$ for all $1 \leq i < j \leq k$. Thus, every edge of K appears in at most two of the Z_i . Below we prove that $\{Z_1, \dots, Z_k\}$ is a generating set for $\mathcal{Z}_t(K)$. Then $\{Z_1, \dots, Z_k\}$ contains a 2-basis of $\mathcal{C}(H)$, and we are done. This follows, since $\mathcal{Z}_t(K)$ is isomorphic to $\mathcal{C}(H)$, a finite dimensional vector space. Thus every generating set contains a basis.

To show that $\{Z_1, \dots, Z_k\}$ generates $\mathcal{Z}_t(K)$, consider $D \in \mathcal{Z}_t(K)$. Let I denote the set of those indices i with $\mathcal{B}_{Z_i} \cap \mathcal{B}_D \neq \emptyset$. We may assume $I = \{1, \dots, k'\}$ for some $k' \leq k$. Then, since $\mathcal{B}_{Z_i} \subseteq \mathcal{B}_D$ and $\mathcal{B}_{Z_i} \cap \mathcal{B}_{Z_j} = \emptyset$ for $i, j \in I$, it follows that \mathcal{B}_D is the disjoint union of the sets $\mathcal{B}_{Z_1}, \mathcal{B}_{Z_2}, \dots, \mathcal{B}_{Z_{k'}}$, and of $\mathcal{B}' := \mathcal{B}_D - \cup_{i=1}^{k'} \mathcal{B}_{Z_i}$. Consequently,

$$\sum_{B \in \mathcal{B}'} B = \sum_{B \in \mathcal{B}_D} B + \sum_{B \in \mathcal{B}_{Z_1}} B + \dots + \sum_{B \in \mathcal{B}_{Z_{k'}}} B = D + Z_1 + \dots + Z_{k'} \subseteq \mathcal{C}(H)$$

since all the summands lie in $\mathcal{C}(H)$. Now if $\mathcal{B}' \neq \emptyset$ then there is a $Z \in \mathcal{Z}_t(K)$ with a non-empty and minimal $\mathcal{B}_Z \subseteq \mathcal{B}'$ which then must be one of the Z_i , a contradiction. Thus, \mathcal{B}' is empty and we have $D = \sum_{i=1}^{k'} Z_i$. ■

Lemma 5.7 gives us the last piece of the proof of MacLane's Theorem for 2-connected graph-like spaces.

Theorem 5.8

If G is a 2-connected graph-like space, then G can be embedded in the plane if and only if $\mathcal{Z}_t(G)$ has a 2-basis.

Proof If G is a 2-connected graph-like space, and K is an embedding of G in the plane, then by Corollary 5.5, $\mathcal{Z}_t(G)$ has a 2-basis.

If $\mathcal{Z}_t(G)$ has a 2-basis, and H is a subgraph of G , then by Lemma 5.7, $\mathcal{C}(H)$ has a 2-basis. It is a standard result in graph theory that $\mathcal{C}(K_5)$ and $\mathcal{C}(K_{3,3})$ do not have 2-bases (see [5] for proof). Therefore G does not have a K_5 or $K_{3,3}$ subgraph. Thus G is a locally connected, 2-connected, compact metrizable space with no subspace homeomorphic to K_5 or $K_{3,3}$. Therefore by Theorem 5.6, G has an embedding in the plane. ■

Note that Theorem 5.8 only applies to 2-connected graph-like spaces. However, the only part of our proof that requires 2-connectedness is the application of Theorem 5.6. In Chapter 6 we prove a more general version of Kuratowski's Theorem for graph-like spaces that will allow us to remove the 2-connectedness condition from Theorem 5.8.

Chapter 6

Thumbtacks and Kuratowski's Theorem

In Chapter 5 we proved that MacLane's Theorem holds for 2-connected graph-like spaces. Our proof relied on the following variant of Kuratowski's Theorem due to Thomassen.

Theorem 6.1 ([17], Thm. 4.3)

Let M be a locally connected, 2-connected, compact topological space. Then M is embeddable in the sphere if and only if M is metrizable, and contains neither of the Kuratowski graphs K_5 and $K_{3,3}$.

In the context of graph-like spaces if M is graph-like, then M is locally connected, metrizable and compact. Thus if G is a 2-connected graph-like space then Theorem 6.1 applies to G . We would like to remove the condition of 2-connectedness from Theorem 6.1. However, Kuratowski's Theorem fails for compact locally-connected metric-spaces that are not 2-connected.

Consider, for instance, the *thumbtack space*. The thumbtack space is the topological space obtained by identifying one end of a simple arc with the centre of a closed disk. This space cannot be embedded in any 2-manifold since the base point of the arc has no neighbourhood homeomorphic to an open disk. The thumbtack space demonstrates the necessity of the 2-connected condition in Theorem 6.1. Although the thumbtack space is not graph-like we define a class of thumbtack-like spaces in Section 6.3 which are obstructions for planarity.

In this chapter we will demonstrate that thumbtack-like spaces are the only additional obstruction to planarity for graph-like spaces. We will extend Theorem 6.1 to all graph-like spaces, and provide the most general statement of Theorem 5.8. In Section 6.1 we prove some additional properties of graph-like spaces, and

some general topological properties of 2-manifolds. Our proof requires analogues of blocks and cut-vertices in graphs; we provide these in Section 6.2. In Section 6.3 we define thumbtack-like spaces and demonstrate that thumbtack-free planar graph-like spaces have embeddings so that every cut-point lies on some face. These embeddings can be combined in order to embed thumbtack-free connected graph-like spaces as we demonstrate in Section 6.4. Finally, Section 6.5 contains the full generalization of Theorems 6.1 and 5.8.

6.1 Topological Properties of Embeddings

In this section we provide a number of technical results on embeddings of graph-like spaces in the plane. These properties will allow us to construct embeddings in Sections 6.3 and 6.4.

Recall from Chapter 3 that if G is a graph-like space embedded in the plane, then Theorem 3.24 gives us a continuous surjection from the closed disk to $C1(F)$ for any face F of the embedding. Furthermore this surjection is a homeomorphism between the open disk and the face F . This result allows us to choose simple arcs and simple closed curves in $C1(F)$ with very specific properties.

Corollary 6.2

Given an embedding Γ of graph-like space G in the plane, let F be a face of Γ . For any point $v \in \text{Bd}(F)$, there is an arc from any $x \in F$ to v having no point other than v in common with Γ . Furthermore for any point $v \in \text{Bd}(F)$, there is a simple closed curve in $C1(F)$ that intersects any given $x \in F$ and intersects $\text{Bd}(F)$ exactly in v .

Corollary 6.2 allows us to construct new graph-like subspaces of the plane from Γ by adding edges to Γ that lie entirely inside a given face. Richter and Thomassen prove the following result in [14].

Proposition 6.3 ([14], Prop. 3)

Let K be a compact 2-connected locally connected subset of the sphere. Then every face of K is bounded by a simple closed curve.

If G is a 2-connected graph-like space, then we can apply Proposition 6.3. The continuous surjection from Theorem 3.24 is actually a homeomorphism between the closed disk and $C1(F)$ when G is 2-connected.

Recall from Chapter 2 that every edge-cut of a graph-like space is finite. Finite edge-connection translates into a more specific property in the plane. Namely, in

any graph-like subspace of the plane, any point v is finitely separated from the points at distance ϵ from v .

Corollary 6.4

Given an embedding Γ of a graph-like space in the sphere, any $p \in \mathbb{S}^2$, and any $0 < \delta < \epsilon$, there are only finitely many disjoint arcs in Γ that connect any point in $\text{Bd}(B(p, \delta))$ to any point in $\text{Bd}(B(p, \epsilon))$.

Proof The closed disks $\text{Cl}(B(p, \delta))$ and $\mathbb{S}^2 - B(p, \epsilon)$ give us two closed subsets A and B of the vertex set of Γ . Now by Theorem 4.A.11 from [4], since V is compact and Hausdorff, and A, B are disjoint closed sets, there is a separation P, Q of V so that $A \subseteq P$ and $B \subseteq Q$. We apply Theorem 4.3 to the sets P and Q . There are only finitely many edges e that have endpoints in both P and Q , so any set of disjoint arcs from P to Q can only be finite in number. Thus any set of disjoint arcs from A to B can only be finite in number. ■

Furthermore, Corollary 6.4 allows us to find more specific separations of an embedded graph-like space Γ . First we have the following result from [17].

Corollary 6.5 ([17], Cor. 4.4)

Let x, y be elements of a locally connected, connected, compact subset M of the sphere. Then precisely one of the two statements below holds:

1. M has a simple closed curve separating x and y ;
2. The sphere contains a simple arc from x to y having only its ends in common with M .

We can apply Corollary 6.5 in order to find a cycle in Γ that separates v from the points of Γ at a given distance.

Proposition 6.6

Let Γ be an embedding of a connected graph-like space in the sphere and let $v \in V$. Then either v is in the boundary of a face of Γ , or for each $\epsilon > 0$ there is a simple closed curve $C_\epsilon \subset \Gamma$ so that C_ϵ separates v from $\text{Bd}(B(v, \epsilon))$ in the sphere.

Proof Since each edge e is in the boundary of at least one face, and the boundary of any face is closed by definition, $\text{Cl}(e)$ is in the boundary of at least one face for each edge e . Thus, any end of an edge in Γ is incident with a face of Γ , so we only consider vertices v that are not in the closure of any edge.

Suppose v is not in the boundary of any face of Γ . Let $\epsilon > 0$ be given and consider $C = \text{Bd}(B(v, \epsilon))$. Then $\Gamma - C$ is a subspace of the sphere that possibly has many components. Let K be the graph-like space obtained from the component of $\Gamma - C$ that contains v as follows. We take the component K' of $\Gamma - C$ containing v and delete all “partial” edges from K' (i.e. arcs of K' corresponding to edges of Γ that have one end in K' and the other in a component of $\Gamma - C$ on the opposite side of C). Now K is the closure of the resulting space. The space K is a graph-like subspace of Γ since it is closed, and hence compact, and connected.

Thus Γ gives us an embedding of the connected graph-like space K in the sphere. Furthermore, K lies entirely inside $\text{Cl}(B(v, \epsilon))$. Consider the faces of K . There is one face F of the embedding that contains $\mathbb{S}^2 - \text{Cl}(B(v, \epsilon))$, together with $C - K$. Fix any point x in the boundary of F . By Corollary 6.5 we either have a simple closed curve $C' \subset K$ that separates x from v , or there is an arc from x to v in the sphere that has only its ends in common with K . In the first case the curve C' separates x from v and hence separates C from v , so $C_\epsilon = C'$ is as desired. In the second case v is on a face of K . We argue that this implies v is on a face of Γ .

If v is in the boundary of a face other than F , say F' , then F' is a face of Γ , since all of $\Gamma - K$ is embedded in F . On the other hand, suppose that v is in the boundary of F . Consider the neighbourhood $B(v, \delta)$ of v for some $0 < \delta < \epsilon$. Since $\Gamma - K'$ is compact, and hence closed, and Γ is Hausdorff, there is some $0 < \delta < \epsilon$ so that $B(v, \delta)$ and $\Gamma - K'$ are disjoint. Now consider $F \cap B(v, \delta)$. Since v is in the boundary of F , v is in the boundary of $F \cap B(v, \delta)$. Also since $F \cap B(v, \delta) \cap \Gamma = \emptyset$, $F \cap B(v, \delta)$ is contained in some face F' of Γ . Finally, since v is not in the boundary of $B(v, \delta)$, v is in the boundary of F' as required. ■

Proposition 6.6 has a more practical corollary.

Corollary 6.7

Suppose Γ is an embedding of a connected graph-like space in the sphere, $v \in V$ is a point with $v \notin \text{Cl}(e)$ for any $e \in E$ and v is not in the boundary of any face of Γ . Then for any simple closed curve $C \subset \Gamma$ so that $v \notin C$, there is a simple closed curve $C' \subset \Gamma$ so that C' separates v from C in the sphere.

Proof Let ϵ be the minimum distance from v to C . Then for any $0 < \epsilon' < \epsilon$ we have by Proposition 6.6 that there is a simple closed curve $C' \subset \Gamma$ so that C' separates v from $\text{Bd}(B(v, \epsilon'))$. Thus any such C' separates v from C , as required. ■

We will apply Corollary 6.7 to obtain a sequence of nested cycles that converge to v in Γ .

In Section 6.4 we will be constructing an embedding of an abstract graph-like space in the plane. We will need to be able to translate the notion of convergence in one metric space to convergence in another. Towards this end we have the following two results.

Proposition 6.8

Let $\{A_i\}$ be a countably infinite collection of subsets of a compact metric space K so that the diameter of A_i approaches zero. Let $a_i \in A_i$ be such that $\lim_{i \rightarrow \infty} a_i = a$. Let $\{b_i\}$ be any sequence in $\cup_{i=1}^{\infty} A_i$ so that for all i , at most finitely many b_j lie in A_i . Then $\lim_{i \rightarrow \infty} b_i = a$.

Proof Let $M(\epsilon)$ be such that, for $i > M(\epsilon)$, $\text{diam}(A_i) < \epsilon$. Consider the open disk $B(a, \epsilon)$ for some $\epsilon > 0$. There is some natural number $N = N(\epsilon)$ so that $a_i \in B(a, \epsilon)$ for all $i > N$. Thus $A_j \subset B(a, 2\epsilon)$ for all A_j with $j > \max\{M(\epsilon), N(\epsilon)\}$. Therefore $B(a, 2\epsilon)$ contains all but finitely many b_i for all $\epsilon > 0$, and $\lim_{i \rightarrow \infty} b_i = a$. ■

Proposition 6.9

Let K, K' be compact metric spaces and let $\{A_i\}$ be a collection of subsets of K whose diameters approach zero. If $f : K \rightarrow K'$ is a continuous function, then the diameters of $\{f(A_i)\}$ approach zero.

Proof Suppose that the diameters of $\{f(A_i)\}$ do not approach zero. Fix some $\epsilon > 0$. Then we have a sequence of pairs of points $\{(a_i, a'_i)\}$ so that each pair $a_i, a'_i \in A_j$ for some j . For each j , A_j contains at most one of these pairs of points. Furthermore, for every i , $d_{K'}(f(a_i), f(a'_i)) \geq \epsilon$.

Now $\{f(a_i)\}$ and $\{f(a'_i)\}$ are both infinite sequences. Thus $\{f(a_i)\}$ has a convergent subsequence $\{f(a_i)\}_{i \in I_2}$. Further $\{f(a'_i)\}_{i \in I_2}$ has a convergent subsequence $\{f(a'_i)\}_{i \in I_3}$. Let $x = \lim_{i \in I_3} f(a_i)$ and $x' = \lim_{i \in I_3} f(a'_i)$. Since $d_{K'}(f(a_i), f(a'_i)) \geq \epsilon$ for each $i \in I_3$, $d_{K'}(x, x') \geq \epsilon$.

We also have that $\{a_i\}_{i \in I_3}$ and $\{a'_i\}_{i \in I_3}$ are infinite sequences. Thus there is a convergent subsequence $\{a_i\}_{i \in I_4}$ so that $\lim_{i \in I_4} a_i = a$. Now by Proposition 6.8 since each A_j contains at most one a'_i , $\lim_{i \in I_4} a'_i = a$. Thus $\lim_{i \in I_4} f(a_i) = f(a) = x$ and $\lim_{i \in I_4} f(a'_i) = f(a) = x'$. Therefore $x = x'$, contradicting $d_{K'}(x, x') \geq \epsilon > 0$. ■

6.2 The 2-Connected Subspaces of a Graph-Like Space

Our main goal in this chapter is to extend Theorem 6.1 to graph-like spaces that are not necessarily 2-connected. In order to accomplish this we will need to apply Theorem 6.1 to the 2-connected subspaces of a graph-like space G . For a finite graph we have the standard notions of a cut-vertex, which is a vertex whose deletion leaves multiple connected components, and a block, which is a maximal 2-connected subgraph. In this section we define analogues of these objects for graph-like spaces.

The blocks B_1, \dots, B_k of a finite graph H partition the edge set of H , and provide a covering of the vertex set of H . Moreover if v is a vertex of H then there is a unique i so that $v \in B_i$ if and only if v is not a cut-vertex of H . The graph-theoretic definition of connectedness allows for a block to be either a single edge, or a 2-connected graph with more than two vertices. However, if we consider H as a graph-like space, then a point $p \in H$ is a cut-point if and only if $H - p$ consists of more than one topologically connected component. Thus the cut-points of the space H are the cut-vertices of the graph H , together with every point $p \in e$ for each cut-edge e of H . Since graph-like spaces may have infinitely many vertices and edges, we will also need to classify the points of H that are not cut-points and are not members of any 2-connected subspace.

Definition 6.10

Given a graph-like space G :

1. a cut-vertex of G is a vertex $v \in V$ so that $G - v$ has more than one component;
2. an edge-block of G is $\text{Cl}(e)$ for any edge $e \in E$ so that if $p \in e$ then $G - p$ has more than one component;
3. a real-block of G is a maximally 2-connected subspace of G that contains more than one vertex;
4. an artificial-block of G is a vertex $v \in V$ so that v is neither a cut-vertex nor in any real-block of G .

Note that real-blocks of a graph-like space exist, since if K is a 2-connected subspace of a G , then by Zorn's Lemma there is a maximal 2-connected subspace of G that contains K . Further we note that any cut-vertex, edge-block, real-block

or artificial-block is graph-like. Every cut-vertex is a single point, and so is a graph-like space. Likewise, every artificial-block is a graph-like space. Every edge-block is a closed connected subspace of G and hence is graph-like. Every real-block is a connected subspace of a graph-like space, so it too is a graph-like space provided that it is closed. To demonstrate that each real-block is closed we prove the following proposition.

Proposition 6.11

The closure of a 2-connected subspace of a Hausdorff topological space is 2-connected.

Proof Let G be a Hausdorff space and let K be a 2-connected subspace of G . Let $x \in \text{Cl}(K)$. If $x \in \text{Cl}(K) - K$, then $\text{Cl}(K) - x$ is connected. This follows, since if $A \subseteq \text{Bd}(K)$ and K is connected, then $K \cup A$ is connected. Thus

$$\text{Cl}(K) - x = K \cup (\text{Bd}(K) - x)$$

is connected.

If $x \in K$, then we claim that

$$\text{Cl}(K) - x = \text{Cl}(K - x) - x.$$

Suppose $y \in \text{Cl}(K - x) - x$, then for each open set U with $y \in U$, $U \cap (K - x) \neq \emptyset$. Thus $U \cap K \neq \emptyset$, and $y \in \text{Cl}(K) - x$.

Now suppose $y \in \text{Cl}(K) - x$. We have that either $y \in K - x$ or $y \in \text{Bd}(K)$. If $y \in K - x$, then $y \in \text{Cl}(K - x) - x$ trivially. If $y \in \text{Bd}(K)$, then for each open set U with $y \in U$, $U \cap K \neq \emptyset$. Since $y \neq x$ and K is Hausdorff, there is a neighbourhood U' of y so that $x \notin U'$. Thus $U \cap (K - x) \neq \emptyset$. Therefore, $y \in \text{Cl}(K - x) - x$, and

$$\text{Cl}(K) - x = \text{Cl}(K - x) - x.$$

Since K is 2-connected, $K - x$ is connected, and $\text{Cl}(K - x) - x$ is connected. Thus if $x \in K$, then $\text{Cl}(K) - x$ is connected. Therefore $\text{Cl}(K)$ is 2-connected. ■

Thus since real-blocks are maximally 2-connected, the real-blocks of G are closed, and hence graph-like.

We also note that in any 2-connected graph-like space, every edge lies in at least one cycle. If an edge e is in real-blocks K_1 and K_2 , then there is a cycle C_1

through e in K_1 and a cycle C_2 through e in K_2 . Now we can extend K_1 and K_2 unless $K_1 = K_2$. Thus the set of edge-blocks together with the set of real-blocks partitions the edge set of a graph-like space. Since the edge set of any graph-like space is countable, there are only countably many edge-blocks and only countably many real-blocks.

To illustrate these concepts, consider the graph-like space T obtained by taking the Freudenthal compactification of the infinite binary tree. We have that the infinite binary tree has countably many vertices and countably many edges. Each edge e of T gives us an edge-block $\text{Cl}(e)$. Each vertex v of T that is also a vertex of the binary tree is a cut-vertex of T . All other vertices of T are ends of the binary tree that are added in the compactification. There are uncountably many such vertices, and each is an artificial-block. However, if T' is the graph-like space obtained by taking the Alexandroff compactification of the binary tree, then T' is 2-connected, and so T' has a single real-block, T' .

Countability is an essential property, since it allows us to perform recursive procedures and inductions. We have that the number of edges of G is countable as is the total number of edge-blocks and real-blocks. Furthermore we demonstrate that while G may have uncountably many vertices, G has only countably many cut-vertices.

Proposition 6.12

Given a graph-like space G , and real-block B , B contains only countably many cut-vertices of G .

Proof Let G be a graph-like space and let B be a real-block of G . For each cut-vertex x of G in B we select an auxiliary point x' in G so that the component of $G - x$ containing x' contains no point of B . Let X be the set of cut-vertices of G that lie in B , and let $X' = \{x' : x \in X\}$. Consider the subset, X'_n , of X' consisting of points at distance at least $1/n$ from B .

Suppose that X'_n is not a finite set. Then there is some point $z \in G$ that is an accumulation point of X'_n . We have that z is at distance at least $1/n$ from B . Consider the neighbourhood $B(z, 1/2n)$. Since G is locally connected, there is a neighbourhood $N \subset B(z, 1/2n)$ of z so that N is connected. Since z is an accumulation point of X'_n we have points $x', y' \in X'_n \cap N$. Thus there is an arc from x' to y' in $\text{Cl}(N)$. If K is the component of $G - x$ containing x' , then $K \cup x$ is closed and connected. Therefore $K \cup x$ is arcwise connected, and there is an

arc from x to x' in $K \cup x$. Likewise there is an arc from y to y' in $K' \cup y$ where K' is the component of $G - y$ containing y' . Therefore there is an arc α from x to y in G that intersects B exactly in $\{x, y\}$. Thus $B \cup \alpha$ is a 2-connected subspace of G , which contradicts the maximality of B . Therefore X'_n has only finitely many elements.

Now consider the sequence Y_0, Y_1, Y_2, \dots , where $Y_0 = X'_1$ and $Y_i = X'_{1/2^i} - X'_{1/2^{i-1}}$ for $i \geq 1$. We have that each auxiliary point x' is at some positive distance from B , and so $x' \in Y_i$ for some i . Furthermore $Y_i \cap Y_j = \emptyset$ whenever $i \neq j$ and each Y_i is finite. Thus the number of cut-vertices of G contained in B is countable. ■

Since G has countably many real-blocks, G has countably many cut-vertices.

6.3 Embedding 2-Connected Graph-Like Spaces in the Plane

In this section we define the class of thumbtack-like spaces. In Section 6.4 we will demonstrate that thumbtack-like spaces are the only additional obstruction to embedding graph-like spaces in the plane. In order to demonstrate this we will exhibit an embedding of a thumbtack-free graph-like space by combining the embeddings of its real-blocks. We spend the bulk of this section proving that every real-block of a thumbtack-free graph-like space has an embedding conducive to this procedure.

We begin by describing the class of thumbtack-like spaces. Recall that cycles in G are homeomorphs of the unit circle. Given a cycle C we define the C -bridges of G to be the components of $G - C$. If K is a C -bridge then the points in $C \cap \text{Cl}(K)$ are the *points of attachment* of K . Two C -bridges K_1 and K_2 *overlap* if either there are points $a, b, c \in C$ so that a, b, c are points of attachment of both K_1 and K_2 or there are points $a, b, c, d \in C$ that appear in this cyclic order so that a, c are points of attachment of K_1 and b, d are points of attachment of K_2 . A *web* centred at v in a graph-like space is a sequence of cycles C_1, C_2, C_3, \dots together with a centre v with the following properties:

1. v is in none of the C_i ;
2. for each i , there is a C_i -bridge containing $\cup_{j=1}^{i-1} C_j$ and a C_i -bridge containing $(\cup_{j>i} C_j) \cup v$;

3. for each i , the C_i -bridge containing C_1, \dots, C_{i-1} and the C_i -bridge containing v overlap.

A *thumbtack-like space* is a web centred at v together with a simple arc that meets the web only at v . Note that $\{C_i : i \in \mathbb{N}\}$ is a collection of edge-disjoint, closed connected subsets of G , so Lemma 2.11 implies that $\{\text{diam}(C_i)\}$ converges to zero. Thus the cycles C_i converge to the centre point v . Thumbtack-like spaces are clearly obstructions to planar embedding.

For example consider the graph-like space W consisting of the circles of radius $1/2^i$ for $i \in \mathbb{N}$, together with the four rays formed by the coordinate axes in the plane pictured in Figure 6.1. The space, W , is 3-connected and thus from [14] we have that W is uniquely embeddable in the plane. Note that the point v corresponding to the origin does not lie on any face of W . Furthermore v is the centre of a web consisting of every circle in W centred at v . Thus if W is a subspace of a graph-like space G where G contains an arc α that meets W only at v , then G cannot be planar, since in every embedding of W the point v does not lie on any face of W , so there is no way to embed α .




Figure 6.1: Graph-like space W embedded in \mathbb{S}^2 .

A graph-like space G is *thumbtack-free* if there is no subspace K of G so that K

is a thumbtack-like space. We spend the remainder of this section demonstrating that each real-block of a thumbtack-free graph-like space G has an embedding in the plane so that no cut-vertex of G is the centre of a web. The ideas presented in the proofs are partially due to Thomassen (private communication).

Lemma 6.13

Suppose G is a thumbtack-free graph-like space, and B is a real-block of G so that B contains neither K_5 , nor $K_{3,3}$, $v \in B$ is a cut-vertex of G , and v is not incident with any edge of B . Then, for every $\epsilon \geq 0$, there is a vertex x at distance at most ϵ from v so that adding an arc α from v to x that is otherwise disjoint from B gives a planar graph-like space $B \cup \alpha$.

Proof Since B is a 2-connected graph-like space that contains neither K_5 nor $K_{3,3}$, by Theorem 6.1 B can be embedded in the plane. We fix any embedding Γ of B . If v is on a face F of Γ , then we can join v to any vertex x in the boundary of F with an arc α that lies entirely in F . Thus for any given $\epsilon > 0$ we can choose $x \in \text{Bd}(F)$ so that x is at distance at most ϵ from v in B . Now $B \cup \alpha$ is planar, since we have exhibited a plane embedding of $B \cup \alpha$.

Suppose that v is not on a face of Γ . For any given $\epsilon > 0$ we can use Proposition 6.6 to find a cycle C in Γ that separates v from all points of B at distance at least ϵ from v . Let M be the C -bridge that contains v . Consider the space $\Gamma - M$. Since M is a C -bridge, M is open and hence $\Gamma - M$ is closed. Furthermore $\Gamma - M$ is connected since all other C -bridges are connected, and are connected to each other through C . Thus $\Gamma - M$ is a graph-like space.

We also have that $\Gamma - M$ is 2-connected, since any cut-vertex of $\Gamma - M$ must lie strictly inside a C -bridge. But such a cut-vertex is a cut-vertex of Γ , a contradiction. Thus $\Gamma - M$ is 2-connected and so, by Proposition 6.3, the faces of $\Gamma - M$ are bounded by cycles. In particular there is a cycle C' that bounds the face of $\Gamma - M$ containing M . Now in Γ , M is the only C' -bridge in the same face of C' as v .

We let $C_1 = C'$ and suppose that we have chosen a sequence of disjoint cycles, C_1, C_2, \dots, C_n so that none of the C_i contains v ; for each i , there is a C_i -bridge containing $\cup_{j=1}^{i-1} C_j$ and a C_i -bridge containing $(\cup_{j=i+1}^n C_j) \cup v$; and, the C_i -bridge containing C_1, \dots, C_{i-1} and the C_i -bridge containing v overlap.

Now we can apply the same argument as above to choose a cycle C that separates v from each C_i so that there is only one C -bridge, M , in the face F of C containing v . Let x be a vertex of attachment of the C -bridge that contains C_n .

Since x lies in the same face of C_1 as v , the distance between x and v is less than ϵ . We consider adding an arc α to B that joins x to v . Let B' be the graph-like space $B \cup \alpha$.

Let $M' = M \cup C \cup \alpha$. Then M' is a closed, connected subspace of B' . Note that M' has only one C -bridge. If M' has an embedding Γ' in the plane then, for any embedding Γ' , C bounds a face of Γ' . This follows, since if C does not bound a face, then the faces of C both contain points of $M \cup \alpha$. However this contradicts the fact that M' has only one C -bridge. Thus if Γ' is an embedding of M' in the plane, we construct an embedding Γ'' of B' in the plane as follows. Let $\phi : B \rightarrow \Gamma$ be the embedding of B in the plane, and let $\phi' : M' \rightarrow \Gamma'$ be the embedding of M' in the plane. Let F be the face of $\phi(C)$ containing $\phi(M)$, and let F' be the face of $\phi'(C)$ containing $\phi'(M \cup \alpha)$. Then we have a natural homeomorphism h from $\phi'(C)$ to $\phi(C)$. By the Jordan-Schönflies Theorem, h extends to a homeomorphism of the plane that maps F' to F . Thus we have $\phi'' : B' \rightarrow \Gamma''$ defined as $\phi''(x) = \phi(x)$ if $x \in B - M$, and $\phi''(x) = h(\phi'(x))$ if $x \in M \cup \alpha$. The result is an embedding Γ'' of B' , as required.

If M' has no planar embedding, then by Theorem 6.1 there is a subspace K of M' that is homeomorphic to either K_5 or $K_{3,3}$. Since K is not a subspace of B , α is an edge of K , and $K - \alpha$ is a subspace of $M \cup C$. Consider the graph K_5 . If we delete the edge $e = uv$ from K_5 the result is a planar graph, and in any embedding of $K_5 - e$, the vertices u, v are separated by a cycle R so that the R -bridge containing u and the R -bridge containing v overlap, and u and v lie in different faces of R . The same holds for any edge e and any embedding of $K_{3,3}$. Thus in $M \cup C$ there is a cycle C' so that the C' -bridges containing x and v overlap, and x and v lie in different faces of C' . Furthermore, C' is disjoint from each C_i with $1 \leq i \leq n$, and by our choice of x , the C' -bridge containing v and the C' -bridge containing C_n overlap. Therefore we set $C_{n+1} = C'$, and we have extended our sequence.

We now repeat this argument. At each step i we either find an arc α so that $B \cup \alpha$ is planar, or we have a sequence of cycles C_1, \dots, C_i that satisfy the definition of a web centred at v . Since G contains no thumbtack-like space, G contains no web centred at v . Thus B contains no web centred at v . Therefore after finitely many repetitions of this argument we find the desired arc α . ■

Since $B \cup \alpha$ is planar, there is an embedding Γ of $B \cup \alpha$ in the plane. Since v is incident with the edge α , v lies on a face of Γ . We now extend this idea to

embed a block B so that for each cut vertex v of G in B , v lies on some face of the embedding.

Lemma 6.13 allows us to embed real-block B of G so that a single cut-vertex lies on a face of the embedding. Note that we can repeatedly apply Lemma 6.13 to embed B so that each cut-vertex $v_i \in \{v_1, v_2, \dots, v_n\}$ lies on a face of the embedding. Thus if B contains finitely many cut-vertices of G , then we can embed B so that each cut-vertex of G in B lies on some face. However, we also need to be able to embed real-blocks of G with infinitely many cut-vertices so that each cut-vertex lies on a face of the embedding.

Lemma 6.14

If G is a thumbtack-free graph-like space, and B is a real-block of G so that B contains neither K_5 , nor $K_{3,3}$, then there is an embedding Γ of B in the plane so that for every cut-vertex v of G in B , there is a face F_v of Γ so that $v \in \text{Bd}(F_v)$.

Proof Since B is a 2-connected graph-like space that does not contain K_5 or $K_{3,3}$, by Theorem 6.1 B has an embedding Γ in the plane. By Proposition 6.12 there are only countably many cut-vertices of G in B , and we let v_1, v_2, \dots be an enumeration of the cut-vertices of G in B that are not incident with an edge of B . Note that if each cut-vertex in B is incident with an edge of B , then in any embedding of B every cut-vertex is on a face. We now proceed to consider the cut-vertices v_i in turn and find for each an arc α_i so that $B' = B \cup (\cup_{i=1}^{\infty} \alpha_i)$ is planar. Then B' has an embedding Γ' in the sphere, and hence Γ' is an embedding of B in the sphere. Moreover, since Γ' is an embedding of B' , each α_i lies on a face of Γ' and hence each v_i lies on a face of Γ' as desired.

We accomplish this exactly as in the proof of Lemma 6.13. For v_1 we choose $0 < \epsilon_1 < 1/2$. Now by Lemma 6.13 there is an arc α_1 so that α_1 is disjoint from B , joins v_1 to a vertex $x_1 \in B$ at distance at most ϵ_1 from v_1 and $B_1 = B \cup \alpha_1$ is a planar graph-like space. Note that α_1 may join v_1 to some v_i with $i > 1$. The space B_1 is graph-like since its vertex set V_1 is the same as the vertex set of B , $B_1 - V_1$ is a set of disjoint arcs, and B_1 is compact. This follows since $B_1 = B \cup \alpha_1$ is the union of two compact spaces, and hence is compact.

Now we iterate this process. At each step i we consider the least index j in our enumeration so that cut-vertex v_j is not the end of any added arc. We find an arc α_i for the cut-vertex v_j . For v_j we choose $0 < \epsilon_i < 1/2^i$ so that no point in $\text{Cl}(\cup_{k=1}^{i-1} \alpha_k)$ is within distance ϵ_i of v_j . We apply Lemma 6.13 to

$B_{i-1} = B \cup (\cup_{k=1}^{i-1} \alpha_k)$ in order to select an arc α_i such that α_i connects v_j to x for some vertex x at distance less than ϵ_i from v_j in B_{i-1} , and $B_i = B_{i-1} \cup \alpha_i$ is planar. Thus by repeated application of Lemma 6.13 we choose a collection of disjoint arcs $\{\alpha_i : i \in \mathbb{N}\}$ so that each B_i is planar.

Now we claim that $B' = B \cup (\cup_{i=1}^{\infty} \alpha_i)$ is planar. First we need to show that B' is graph-like. We have that the set $V' = V(B)$ is totally disconnected, $B' - V'$ is a disjoint set of arcs, and B' is compact. This follows, since if U is an open cover of B' , then each $A \in U$ gives an open set $A - \cup_{i=1}^{\infty} \alpha_i$ of B . Thus we have a finite subcover U_1 of B .

Since U_1 is finite, the corresponding sets in B' all have diameter at least ϵ for some $\epsilon > 0$. We claim that U_1 covers all arcs α_i with length less than δ for some $0 < \delta < \epsilon$. This follows, since if we assume otherwise, then there is a sequence of points $\{x_n\}$ so that each x_n lies in some α_i with length less than $1/n$, and each $x_n \notin U_1 = \cup_{A \in U_1} A$. We have that $\{x_n\}$ has a convergent subsequence, which we take to be $\{x_n\}$, and we let $\{x_n\}$ converge to x . Since each α_i is an open set, and at most finitely many of the x_n lie on α_i , $x \in B$. Thus U_1 is an open subset of B' that contains x , and none of the points x_n , a contradiction. Therefore there is some $0 < \delta < \epsilon$ so that U_1 covers all arcs α_i with length less than δ .

Since there are only finitely many arcs α_i with length greater than δ , and each α_i is compact, we can choose finite open subcovers of U for each α_i . Together with U_1 this gives us a finite subcover of B' , and thus B' is compact. We also have that B' and each B_i is 2-connected, since B is 2-connected. Thus B' is graph-like, so we can apply our previous results. By Theorem 6.1 it suffices to check that B' contains no copy of K_5 or $K_{3,3}$.

Suppose B' contains a copy K of K_5 (the $K_{3,3}$ case is similar). Since each B_i is planar, K is not a subset of any B_i . This implies that K contains infinitely many of the α_i . Let v_1, \dots, v_5 be the vertices of K and e_1, \dots, e_{10} be the edges of K . Then we can choose $\epsilon > 0$ so that the neighbourhoods $B(v_i, \epsilon)$ are pairwise disjoint in B' . Since B is locally connected we have for each v_i some connected neighbourhood N_i of v_i in B so that $N_i \subset B(v_i, \epsilon)$. For v_1 we let the edges of K incident with v_1 be the edges e_1, e_2, e_3, e_4 . For $i = 1, 2, 3, 4$ consider the first point x_i of e_i not in N_1 (x_i is well-defined since $B - N_1$ is closed). If $x_i \in B$ then we set $y_i = x_i$, otherwise x_i is a point on some added arc α_j . In that case we set y_i to be the endpoint of α_j in the subarc of e_i from v_1 to x_i . Now since $\text{Cl}(N_1)$ is a connected subset of B there is a path P_i from v_1 to y_i in $\text{Cl}(N_1)$. Then the

subspace $P = \cup_{i=1}^4 P_i$ is a closed connected subspace of B , and hence graph-like. Thus we can take a spanning tree T_1 of P . We repeat this process for each v_i in order to obtain trees T_1, \dots, T_5 .

Now consider the edge e_1 of K . Suppose that e_1 connects v_1 to v_2 . For v_1, v_2 we have neighbourhoods N_1 and N_2 and spanning trees T_1 and T_2 . Let a be the first point of intersection of e_1 with T_1 (as we traverse e_1 from v_2 to v_1). Likewise, let b be the first point of intersection of e_1 with T_2 (as we traverse e_1 from v_1 to v_2). If a is a vertex, then set $x = a$, otherwise a is the interior point of an edge of B' , and we take x to be the end of that edge in the subarc of e_1 from a to b . In the same way we choose a point $y = b$ or y the first vertex on e_1 in the subarc from b to a . Let α be the subarc of e_1 from x to y . If α contains finitely many added arcs α_i then we do not change α .

Otherwise we show how to choose a new arc β from x to y so that β uses only finitely many added arcs. Note that α is a closed subset of B' , as is $K - e_1$. Thus we have some $0 < \epsilon$ so that for all $z \in K - e_1$, $z' \in \alpha$, z and z' are at distance at least ϵ in B' . Let $c, d \in \alpha$ be the endpoints of an added arc $\alpha_i \subset \alpha$ so that α_i has length less than ϵ . Then there is a path P_i from c to d in B of length less than ϵ , so $P \cap (K - e_1) = \emptyset$. We choose a path P_i for each such α_i . Now α' , the space obtained by deleting all of the arcs α_i with length less than ϵ from α and replacing them with the paths P_i , is closed and connected, and hence arcwise connected. Thus there is an arc β from x to y in α' . Since there are only finitely many arcs α_i with length at least ϵ , β contains only finitely many added arcs, as required. Therefore we can replace the edge e_1 with e'_1 an arc from the last vertex of T_1 on e_1 to the last vertex of T_2 on e_1 so that e'_1 contains only finitely many added arcs.

We repeat the above procedure for each e_i in turn. The result is the space

$$K' = (\cup_{i=1}^5 T_i) \cup (\cup_{i=1}^{10} e'_i)$$

which contains a copy of either K_5 or $K_{3,3}$. Furthermore, K' contains only finitely many added arcs, so there is some index i so that every added arc in K' is contained in B_i . But then $K' \subset B_i$ and G is contained in B_i . This contradicts the planarity of B_i . Thus no such subspace K of B' exists. By Theorem 6.1 B' is planar, and we have the desired result. ■

Finally we prove the following technical lemma.

Lemma 6.15

Suppose G is a thumbtack-free graph-like space, and B is a real-block of G so that B contains neither K_5 , nor $K_{3,3}$. If $x, y \in B$ are cut-vertices of G , then either there is a cycle C_B in B so that the C_B -bridge containing x and the C_B -bridge containing y are distinct and overlap, or there is an embedding Γ of B in the plane so that every cut-vertex of G in B is on a face of Γ , and x and y are on the same face of Γ .

Proof By Lemma 6.14, B has an embedding Γ so that for each cut-vertex v of G in B , there is a face F_v of Γ so that $v \in \text{Bd}(F_v)$. Thus there are faces F_x and F_y such that $x \in \text{Bd}(F_x)$ and $y \in \text{Bd}(F_y)$. If $F_x = F_y$, then we have the result. Otherwise, consider the space $B' = B \cup \alpha$ where α is an arc from x to y .

We have that B' is a 2-connected graph-like space, so by Theorem 6.1 either B' is planar or there is a K_5 or $K_{3,3}$ in B' . If B' is planar, then by Lemma 6.14 we can embed B' in the plane so that each cut-vertex v of G in B lies on a face of the embedding. Furthermore, since α is an edge in B' , x and y lie on the same face of the embedding. Thus, when we remove α we obtain an embedding of B with the required properties.

Suppose that B' is not planar, and K is a copy of a forbidden subgraph in B' . Since B is planar, K is not a subspace of B . Thus α is an edge of K . Consider the subspace $K - \alpha$. If $e = uv$ is any edge of K_5 , $K_5 - e$ can be embedded in the plane, and in any embedding, there is a cycle of $K_5 - e$ that separates u from v in the plane, and the bridges of u and v are distinct and overlap. The same holds for $K_{3,3}$. Therefore there is a cycle C_B in $K - \alpha$ so that x and y lie in distinct faces of C_B . Furthermore, the C_B -bridge containing x and the C_B -bridge containing y are distinct and overlap, as required. ■

6.4 Embedding Connected Graph-Like Spaces in the Plane

In this section we give a procedure for embedding a thumbtack-free graph-like space with no copy of K_5 or $K_{3,3}$ in the plane. Before we begin, consider a finite graph H that is connected, but not 2-connected and contains no K_5 nor $K_{3,3}$ minor. If we can prove Kuratowski's Theorem for 2-connected graphs, then we can extend this to connected graphs as follows.

We have that every block of H is 2-connected and contains no K_5 nor $K_{3,3}$ minor, thus we can embed the blocks of H in the plane. Let B_1, \dots, B_k be the blocks

of H . We have for each B_i that B_i can be embedded in the plane, furthermore B_i can be embedded so that every vertex of B_i lies on a face of the embedding (this is trivially true for all embeddings of B_i). Fix an embedding of B_1 in the plane. Let v be a cut-vertex of H in B_1 , and suppose that v is also in B_2 . Then v lies on a face of the embedding, so we can specify a closed disk B in the plane so that $B \cap B_i = \{v\}$. Now we have an embedding of B_2 in the plane so that v lies on a face F of the embedding. There is a natural homeomorphism between $\text{Bd}(F)$ and $\text{Bd}(B)$ that maps v to itself. The Jordan-Schönflies Theorem allows us to extend this homeomorphism to a homeomorphism between $\mathbb{S}^2 - F$ and B . The result is an embedding of $B_1 \cup B_2$ in the plane. We simply repeat this process until we have embedded each B_i . This gives us an embedding of H in the plane. We follow this strategy in order to embed a thumbtack-free graph-like space, G , in the plane.

Suppose that G is a thumbtack-free graph-like space that contains no copy of K_5 , nor $K_{3,3}$. Then the real-blocks and edge-blocks of G can be embedded in the sphere by Theorem 6.1. Since there are only countably many real-blocks and edge-blocks, let B_1, B_2, \dots be a fixed enumeration of the real-blocks and edge-blocks of G . We begin by embedding the blocks of G in sequence. At step 0 of this process we take H_1 to be B_1 and embed H_1 in the sphere as Γ_1 using any embedding of B_1 so that all of the cut-vertices of G in B_1 are on faces of Γ_1 (such an embedding exists by Lemma 6.14).

At step #1 we consider B_2 . Since G is connected, if $x \in H_1$ and $y \in B_2$ are arbitrary vertices, then there is an arc P from y to x in G . Let u be the last point of B_2 on P , let v be the first point of H_1 on P and let the subarc of P from u to v be P (note that v is “after” u in the traversal of P since B_2 is a real-block). The arc P is unique in the following respect. There are at most countably many cut-vertices of G on P , and the set of cut-vertices P' of G on P does not depend on P , nor does their order from u to v on P . Otherwise we have an arc Q from u to v in G such that the set of cut-vertices Q' is different from P' . Consider $a \in Q' - P'$. Now P is an arc from B_2 to H_1 in $G - a$. This is a contradiction since each cut-point on any arc from B_2 to H_1 separates B_2 from H_1 . Thus P determines a unique sequence of cut-vertices and hence a unique sequence of real-blocks and edge-blocks required to connect B_2 to H_1 .

We have an embedding $h_1 : H_1 \rightarrow \Gamma_1$ of H_1 in the sphere, and we wish to extend our embedding to an embedding of some H_2 containing H_1 and B_2 . First

we have that u is in a face of Γ_1 , so by Corollary 6.2 we can choose a sequence of simple closed curves that intersect H_1 only in u , intersect each other only in u , and have diameter approaching zero. Thus by Proposition 6.9 we have a sequence of simple closed curves in the sphere so that: each curve lies in the same face of Γ_1 ; each curve intersects Γ_1 only in u ; the intersection of any two curves is u ; the curves have diameter approaching zero in the sphere; and, the curves are nested (*i.e.* each curve C has two faces, one containing all of the curves with diameter less than $\text{diam}(C)$, and one containing all of the curves with diameter greater than $\text{diam}(C)$). Thus we can choose three curves, C_1, C_2, C_3 in the sphere so that $C_i \cap C_j = u$, $C_i \cap \Gamma_1 = u$ and the diameter of C_i is less than $1/2^i$. We embed B_2 inside of C_3 . We have two cases.

The first case is that $u = v$ and P is a single point. In this case B_2 has an embedding in the sphere so that each cut-vertex of G is on some face of the embedding. If B_2 is an edge-block then we simply embed B_2 as any chord of C_3 connecting u to an arbitrary point $x \in C_3 - u$. If B_2 is a real-block then B_2 is 2-connected. If F is the face of the embedding with $u \in \text{Bd}(F)$ then F is bounded by a simple closed curve C . There is a natural homeomorphism between C and C_3 that maps u to itself. By the Jordan-Schönflies Theorem this homeomorphism can be extended to a homeomorphism between $\mathbb{S}^2 - F$ and the closed disk bounded by C_3 that contains no points of H_1 other than u . This homeomorphism gives us an embedding $h_2 : H_2 \rightarrow \Gamma_2$ where $H_2 = H_1 \cup \{C_1, C_2\}$.

The second case is that $u \neq v$. In this case the arc P gives us a unique sequence of edge-blocks and real-blocks. We choose C_1, C_2, C_3 exactly as above. Now we add a chord α to C_3 between any two points of $C_3 - u$. The chord α breaks the closed disk bounded by C_3 into two closed disks D_1 and D_2 . Without loss of generality we stipulate that u lies in the boundary of D_1 . Again there are two cases, either B_2 is an edge-block or a real-block.

If B_2 is an edge-block then we embed B_2 as any chord of the boundary cycle of D_2 so that v is mapped to an interior point of α and the other end of B_2 is mapped to a point of C_3 . If B_2 is a real-block then we embed B_2 inside D_2 exactly as above, except that in this case we have a boundary cycle C containing v and we choose a natural homeomorphism between C and the boundary cycle of D_2 so that v is mapped to any interior point of α . We embed B_2 inside D_2 using the Jordan-Schönflies Theorem as before. Now we have H_1, C_1, C_2, D_1 and B_2 embedded in the sphere. We can choose any arc β in the interior of D_1 from u to v , and there

is a natural homeomorphism between P and β . Thus $H_1 \cup B_2 \cup \{C_1, C_2\} \cup P$ is a graph-like space, and we have an embedding of this space in the sphere.

Note that each real-block or edge-block B that intersects P in more than a vertex intersects P in a closed arc. Furthermore, if we let E' be the set of open subarcs of P so that, for each $e' \in E'$, $\text{Cl}(e')$ is the intersection of P with some real-block or edge-block B , then $P - \cup_{e' \in E'} e'$ is closed and totally disconnected. Thus by Lemma 4.11 we can choose a set of closed disks $\{D_{e'}\}$ that enclose each $e' \in E'$. We apply this Lemma to either face with β in its boundary. Now we have an embedding of P together with a closed disk for each element of E' so that the disks only meet at vertices of P that are cut-vertices of G .

For each edge-block B traversed by P , we simply embed B as the element of E' corresponding to B . For each real-block B that P passes through we have $e' \in E'$ corresponding to B . We would like to embed B inside $D_{e'}$. This is possible so long as the ends of e' , x and y correspond to cut-vertices of B that can be embedded on the same face. If there is an embedding of B so that each cut-vertex of G lies on a face and x and y lie on the same face, then as before there is a cycle C bounding that face, and a natural homeomorphism between C and $\text{Bd}(D_{e'})$ that we can extend to the interior of $D_{e'}$. The final possibility is that no such embedding of B exists. In this case by Lemma 6.15 we can identify a cycle C_B in B that separates x from y so that the C_B -bridges containing x and y overlap.

Note that only finitely many elements of E' can give us a bad cycle C_B . If there are infinitely many such bad blocks, then we have a monotonic convergent sequence of cut-vertices $\{z_i\}$ on P , where each z_i is an element of $P - \cup_{e' \in E'} e'$. Thus we have a monotonic subsequence of $\{z_i\}$ that converges to a point $z \in P - \cup_{e' \in E'} e'$. Therefore, z is a cut-vertex of G . Now the bad cycles C_{B_1}, C_{B_2}, \dots form a web centred at z . Since z is a cut-vertex of G , G contains a thumbtack-like space, a contradiction. Thus there are only finitely many bad cycles.

We can choose the block B so that B contains a bad cycle, and B is the first such block traversed by P' from u to v . We replace B_2 by B in our initial ordering and repeat this process starting with the embedding of H_1 . Now the arc P from B to H_1 contains no bad blocks. Note that since for B_2 there are only finitely many bad blocks, we only swap at most finitely many elements in our enumeration of the real-blocks of G . This ensures that as we repeat this process, any block B_n in our original enumeration is considered after finitely many steps.

The purpose of C_1 and C_2 is to separate H_1 from B_2 . In order to embed G we

repeat this process for each B_i . After we have embedded all of the B_i we want our constructed space to be homeomorphic to G . The “cuff” formed by $C_1 \cup C_2$ together with the face of $C_1 \cup C_2$ bounded by both C_1 and C_2 stops us from adding a sequence of B_i 's to B_2 that converges to a point of H_1 . If that were to occur, then H_1 and B_2 would be contained in a 2-connected subspace of our embedding.

The general case of this construction is the same as step 1. At step i we have a least index j so that the block B_j in our ordering has not yet been embedded. We have a graph-like space H_i that contains all B_l for $l < j$, together with additional circles $C_1, C_2, \dots, C_{2i-1}, C_{2i}$ so that $\text{diam}(C_i) < 1/2^i$. At the end of this step we have a homeomorphism $h_{i+1} : H_{i+1} \rightarrow \Gamma_{i+1}$, where

$$H_{i+1} = H_i \cup \{C_{2i+1}, C_{2i+2}\} \cup (\cup_{B \in P} B) \cup B_j$$

and P is the arc in G from B_j to H_i . We claim that $\Gamma = \text{Cl}(\cup_{i=1}^{\infty} \Gamma_i)$ is a graph-like space embedded in the sphere containing G as a subspace.

First we prove a property of the circles C_i . For each i we have a pair of circles C_{2i-1}, C_{2i} . Note that the circles C_{2i-1} and C_{2i} separate the sphere into three open disks. Furthermore, one of the open disks D_i contains both C_{2i-1} and C_{2i} in its boundary. We refer to D_i as the *cuff* corresponding to i . We have that each cuff D_i has the property that $D_i \cap \Gamma = \emptyset$. Thus we can think of the cuff D_i as defining a partition of the sphere, and of Γ . We have an open disk A_i bounded by C_{2i-1} , and an open disk A'_i bounded by C_{2i} . Every point $x \in \Gamma$ lies either in $\text{Cl}(A_i) - C_{2i}$, $\text{Cl}(A'_i) - C_{2i-1}$ or $C_{2i-1} \cap C_{2i}$. Furthermore, $C_{2i-1} \cap C_{2i}$ consists of a single point, and is a vertex of H_i .

Suppose that $v \in \Gamma - \cup_{i=1}^{\infty} \Gamma_i$. We define a 01-string α_v as follows. Since v is not a vertex of H_i for any i , $v \notin C_{2i-1} \cap C_{2i}$. Thus v either lies in $\text{Cl}(A_i) - C_{2i}$ or in $\text{Cl}(A'_i) - C_{2i-1}$. We set $\alpha_v[i] = 1$ if $v \in \text{Cl}(A_i) - C_{2i}$, and $\alpha_v[i] = 0$ if $v \in \text{Cl}(A'_i) - C_{2i-1}$.

Proposition 6.16

If $v \in \Gamma - \cup_{i=1}^{\infty} \Gamma_i$, then α_v contains infinitely many 0's, and the set $\{C_{2i} : \alpha_v[i] = 0\}$ converges to v .

Proof If α_v does not contain infinitely many 0's, then there is some index N so that $\alpha_v[i] = 1$ whenever $i > N$. Thus $v \in \text{Cl}(A_i) - C_{2i}$ for all $i > N$. However, at each step i , we only embed elements of $H_i - H_{i-1}$ in the disk A'_i . Therefore, since

$v \notin A'_i$ whenever $i > N$, $v \in H_j$ for some index j . This contradicts our choice of v , so no such index N exists.

Consider the set $\{C_{2^i} : \alpha_v[i] = 0\}$. Recall that we chose the circles C_i so that $\text{diam}(C_i) < 1/2^i$. Since v lies in the open disk A'_i bounded by C_{2^i} for each i with $\alpha_v[i] = 0$, and the diameters of these disks approach zero, $\{C_{2^i} : \alpha_v[i] = 0\}$ converges to v . ■

Now we proceed with the proof that Γ is a graph-like space containing G as a subspace.

Proposition 6.17

Γ is a graph-like space.

Proof Let $V = \cup_{i=1}^{\infty} V(\Gamma_i)$, and let $V' = \Gamma - \cup_{i=1}^{\infty} \Gamma_i$. We claim that the set $V \cup V'$ is a totally disconnected subset of Γ . In order to prove this we consider points $x, y \in V \cup V'$ and find a separation X, Y of $V \cup V'$ so that $x \in X$ and $y \in Y$. We have two cases.

Suppose that x, y both lie in the same real-block or edge-block of Γ_i for some i . Let B be the real-block or edge-block containing x and y . Since B is graph-like, $V(B)$ is totally disconnected, and there is a separation X_i, Y_i of $V(B)$ so that $x \in X_i$ and $y \in Y_i$. Now we use X_i and Y_i to construct X and Y . At each step of the construction we construct Γ_j by adding vertices and edges to Γ_{j-1} . Furthermore, there is a cut-vertex v_{j-1} that connects Γ_{j-1} to $\Gamma_j - \Gamma_{j-1}$. Thus if $v_{j-1} \in X_{j-1}$ then we set

$$X_j = X_{j-1} \cup (V(\Gamma_j) - V(\Gamma_{j-1}))$$

and $Y_j = Y_{j-1}$. Now $x \in X_j, y \in Y_j$ and X_j, Y_j partition $V(\Gamma_j)$. Furthermore, since

$$(V(\Gamma_j) - V(\Gamma_{j-1})) \cup \{v_{j-1}\}$$

is closed, X_j, Y_j is a separation of $V(\Gamma_j)$. Thus for each $j \geq i$ we have a separation of $V(\Gamma_j)$. Take $X = \text{Cl}(\cup_{j=i}^{\infty} X_j)$ and $Y = \text{Cl}(\cup_{j=i}^{\infty} Y_j)$. If $v \in V'$, then we have a string α_v associated with v . Furthermore, by Proposition 6.16, $\{C_{2^j} : \alpha_v[j] = 0\}$ converges to v . However, for all $j > i$ for which $\alpha_v[j] = 0$, we have without loss of generality that the vertices added to Γ_{j-1} all become part of X_j . Thus $v \in \text{Cl}(\cup_{j=i}^{\infty} X_j)$ and $v \notin \text{Cl}(\cup_{j=i}^{\infty} Y_j)$. Therefore X and Y partition $V \cup V'$ into disjoint closed sets so X, Y is a separation of $V \cup V'$ and $x \in X$ and $y \in Y$.

Now suppose that x, y do not lie in the same real-block or edge-block of Γ_i for some i . Then since Γ is a closed connected subset of the sphere, Γ is arcwise connected, and there is an arc τ from x to y in Γ . Furthermore, by construction τ passes through a real-block or edge-block, B of some Γ_i . Let u and v be the cut-points in B that τ passes through in order from x to y . Note that u and v are unique in the same sense as previously described. There is a separation X', Y' of $B \cap (V \cup V')$ so that $u \in X'$ and $v \in Y'$. We let X_i, Y_i be the separation of Γ_i defined as follows. Let K_1, K_2, \dots, K_n be the components of $\Gamma_i - u$ and let $v \in K_1$. Then we set

$$\begin{aligned} X_i &= X' \cup (\cup_{j=2}^n V(K_j)), \quad \text{and} \\ Y_i &= Y' \cup (V(K_1) - V(B)). \end{aligned}$$

Now X_i, Y_i is a separation of $V(\Gamma_i)$, and if $x \in V(\Gamma_i)$ then $x \in X_i$, and if $y \in V(\Gamma_i)$ then $y \in Y_i$. We continue with the previous construction, constructing each X_j, Y_j by augmenting the previous separation. The result is a separation X, Y of $V \cup V'$. Furthermore, from the construction of Γ we must have that $x \in X$ and $y \in Y$. Therefore $V \cup V'$ is totally disconnected.

Furthermore, $\Gamma - (V \cup V')$ consists of a set of disjoint open arcs. This follows, since by construction, the connected components of $\Gamma - (V \cup V')$ are subsets of $\cup_{i=1}^{\infty} E(\Gamma_i)$, which is a set of open arcs. It remains to show that Γ is compact, Hausdorff and metrizable. Since Γ is a subset of \mathbb{S}^2 , Γ is automatically Hausdorff and metrizable. Furthermore, Γ is closed, and hence compact, by definition. Thus Γ is graph-like. ■

Proposition 6.18

G is homeomorphic to the space $\Gamma - \cup_{i=1}^{\infty} C_i$.

Proof We construct a homeomorphism $h : G \rightarrow \Gamma'$ where

$$\Gamma' = \text{Cl}(\cup_{i=1}^{\infty} \Gamma_i - \cup_{i=1}^{\infty} C_i) = \text{Cl}(\Gamma - \cup_{i=1}^{\infty} C_i).$$

First we have homeomorphisms $h_i : H_i \rightarrow \Gamma_i$ from our construction. Also by construction, $h_i(x) = h_j(x)$ for all $j > i$ and $x \in H_i$. Let A be the set of artificial blocks of G . For each $x \in G - A$ we have some i so that x is mapped into H_i by h_j for all $j \geq i$. We set $h(x) = h_i(x)$ in these cases. This leaves us with points $x \in A$.

Fix some vertex y of G so that y is either in an edge-block or in a real-block of G . For each $x \in A$ choose an arbitrary arc P_x from y to x in G . The arc P_x determines a unique sequence of real-blocks and edge-blocks, and thus P_x determines a unique sequence of circles $\{C_{2i}\}_{i \in I}$ that arise from the embeddings of those blocks in the procedure above. There is a unique point v so that $\{C_{2i}\}_{i \in I}$ converges to v in Γ and we set $h(x) = v$.

We have that h is a bijection between $G - A$ and $\cup_{i=1}^{\infty} \Gamma_i$. We also have that h is a bijection between A and V' (which we recall is $\Gamma - \cup_{i=1}^{\infty} \Gamma_i$). This follows, since by Proposition 6.16 if $v \in V'$, there is a 01-string α_v so that v is the limit of $\{C_{2j} : \alpha_v[j] = 0\}$. Thus if $x \neq x'$ then P_x and $P_{x'}$ define distinct sets of circles $\{C_{2i}\}_{i \in I}$ and $\{C_{2i}\}_{i \in I'}$ that converge to distinct $v, v' \in V'$, and $h : A \rightarrow V$ is an injection. Similarly, if $v \in V$, then let I be the set of indices so that $\alpha_v[i] = 0$. Then, v is embedded in the disk bounded by C_{2i} for each $i \in I$. For each i , the circle C_{2i} arises from embedding a real-block or edge-block B_{j_i} for some index j_i . For each $i \in I$ let p_i be an arbitrary vertex in B_{j_i} , and let P_i be an arc from p_{i-1} to p_i in Γ' . Let $P = \cup_{i \in I} P_i$. Now P is a connected subspace of G with a single accumulation point, p . Thus p is an artificial block of G and there is an arc P_p from y to p in $\text{Cl}(P)$. By definition $h(p) = v$, so h is a surjection. Therefore, $h : G \rightarrow \Gamma'$ is a bijection.

For continuity, we show that for each $y \in \Gamma$, and each $\epsilon > 0$, there is a $\delta > 0$ so that if the distance between x and $h^{-1}(y)$ in G is less than δ then the distance between $h(x)$ and y in the sphere is less than ϵ . We consider two cases.

Case #1:

$h^{-1}(y)$ is an artificial block in G .

Consider the open disk $B(y, \epsilon)$. From the above discussion we have that there is a sequence $\{C_{2i}\}_{i \in I}$ of cycles in $\Gamma - \Gamma'$ so that $\{C_{2i}\}_{i \in I}$ converges to y . Thus there is some $N \in \mathbb{N}$ such that $C_{2i} \subset B(y, \epsilon)$ for all $i \in I$ with $i > N$. Take any $i > N$, and consider C_{2i} . In the original construction the circles C_{2i-1} and C_{2i} were embedded so that $C_{2i-1} \cap C_{2i} = u$ and u is a cut-vertex of G . Since u is a cut-vertex, there is a component L of $G - u$ containing $h^{-1}(y)$ and L is an open set. Thus there is a $\delta > 0$ so that $B(h^{-1}(y), \delta) \subset L$. Now for any point $v \in \Gamma'$ outside of C_{2i-1} , $h^{-1}(v) \notin L$ and hence is at distance greater than δ from $h^{-1}(y)$. Therefore for every point x at distance less than δ from $h^{-1}(y)$ in G , $h(x)$ lies inside C_{2i} and hence inside of $B(y, \epsilon)$.

Case #2:

$h^{-1}(y)$ is not an artificial block in G .

By Proposition 6.17, Γ is a graph-like space. Given an $\epsilon > 0$ we choose some $\epsilon > \epsilon' > 0$ arbitrarily and consider the disks $D = B(y, \epsilon)$ and $D' = B(y, \epsilon')$. Consider the components $\{K_i\}$ of $\Gamma - y$. By Corollary 6.4 only finitely many of these components K_1, K_2, \dots, K_n have points outside of D . Otherwise we would be able to find infinitely many disjoint arcs connecting y to points outside of D . These arcs would give us infinitely many disjoint arcs connecting points in $\text{Bd}(D')$ to points in $\text{Bd}(D)$. Choose $y_i \in K_i$ for each i so that y_i is neither a cut-vertex nor an artificial block of G . Then, for each i , we have a real-block or an edge-block B_i of G so that $y_i \in B_i$. Each B_i has a finite place j_i in the ordering of the real-blocks and edge-blocks in our original construction. Let $l = \max\{j_i : 1 \leq i \leq n\}$. Now H_l contains each B_i , and we have a homeomorphism h_l between H_l and Γ_l . Thus there is some $\delta > 0$ so that for each x in H_l at distance less than δ from $h^{-1}(y)$ in H_l , $h_l(x) = h(x)$ is at distance less than ϵ from y in the sphere. If x is any point of $G - H_l$ at distance less than δ from $h^{-1}(y)$ then $h(x)$ lies in some K_i for $i > n$. Thus $K_i \subset D$ and $h(x)$ is at distance less than ϵ from y .

Therefore h is a continuous function. Now since G is compact and Γ' is Hausdorff, h a continuous bijection implies that h is a homeomorphism. Thus G is homeomorphic to Γ' as required. ■

Propositions 6.17 and 6.18 prove the following theorem.

Theorem 6.19

Let G be a connected graph-like space. Then G is embeddable in the sphere if and only if G contains neither of the Kuratowski graphs K_5 and $K_{3,3}$, nor any thumbtack-like space.

6.5 MacLane's Theorem and Kuratowski's Theorem

In this section we present a simple extension of Theorem 6.19 and a simple generalization of Theorem 5.8. First recall that by Proposition 2.12 if G is a graph-like space, then G has finitely many connected components. If G is thumbtack-free and contains no copy of K_5 nor $K_{3,3}$, then its components satisfy the hypotheses of Theorem 6.19. Thus we have the following result.

Theorem 6.20

Let G be a graph-like space. Then G is embeddable in the sphere if and only if G contains none of the Kuratowski graphs K_5 and $K_{3,3}$, nor any thumbtack-like space.

Proof By Proposition 2.12 G has only finitely many connected components. Let these components be X_1, X_2, \dots, X_n . By Theorem 6.19 since each X_i contains none of the graphs K_5 , nor $K_{3,3}$ and contains no thumbtack-like space, each has an embedding in the sphere. We use these embeddings to construct an embedding of G . First we embed X_1 in the sphere. Now X_1 has a face F . We take an arbitrary simple closed curve C contained entirely in F . We choose an arbitrary face F' and simple closed curve C' contained in F' in our embedding of X_2 . By the Jordan-Schönflies Theorem we have a homeomorphism that maps the side of C' containing X_2 to the side of C that does not contain X_1 . Thus we have an embedding of $X_1 \cup X_2$ in the sphere. We continue this process until we have the desired embedding of G . ■

This gives a more general version of Theorem 6.1 for graph-like spaces. Also note that this proof together with the material developed in Section 6.4 can be adapted to compact, locally connected metric spaces.

Recall the results from Chapter 5. Note that we can reformulate Theorem 5.8 and Theorem 6.1 as the following result.

Theorem 6.21

If G is a 2-connected graph-like space, then $\mathcal{Z}_t(G)$ has a 2-basis if and only if G contains no copy of K_5 nor $K_{3,3}$.

We have that if G is a graph-like space, then the cycle space of G , $\mathcal{Z}_t(G)$ is the space generated by the cycles of G . If G is not 2-connected, then the cycles of G are partitioned by the real-blocks of G . Thus when considering the cycle space of G we can ignore the artificial-blocks of G and the edge-blocks of G . Furthermore, $\mathcal{Z}_t(G)$ is the direct sum of the spaces $\mathcal{Z}_t(B)$ where B is a real-block of G . Therefore $\mathcal{Z}_t(G)$ has a 2-basis if and only if $\mathcal{Z}_t(B)$ has a 2-basis for each real-block B of G . Since K_5 and $K_{3,3}$ are both 2-connected, G contains a copy of K_5 or $K_{3,3}$ if and only if there is a real-block B of G that contains a copy of K_5 or $K_{3,3}$. Thus we can remove the 2-connected criterion from the previous theorem.

Theorem 6.22

If G is a graph-like space, then $\mathcal{Z}_t(G)$ has a 2-basis if and only if G contains no copy of K_5 nor $K_{3,3}$.

Finally, note that if G is a graph-like space consisting of a planar web, together with an arc α that meets the web only at its centre, then $G - \alpha$ is planar and 2-connected. Thus $\mathbb{Z}_t(G - \alpha)$ has a 2-basis. However, $\mathcal{Z}_t(G - \alpha) = \mathcal{Z}_t(G)$, so $\mathcal{Z}_t(G)$ has a 2-basis. Therefore we cannot remove the 2-connectedness criterion from Theorem 5.8. But in light of our above discussion we can extend Theorem 5.8 to the following.

Theorem 6.23

If G is a graph-like space, then G is planar if and only if $\mathcal{Z}_t(G)$ has a 2-basis, and G is thumbtack-free.

Chapter 7

Conclusion

This thesis focused on the topological properties of graph-like spaces necessary to develop a theory of cycle spaces and a theory of embeddings. The material was presented in two main parts. In the first, we developed a theory of embeddings of graph-like spaces in surfaces, together with a theory of algebraic edge spaces of graph-like spaces. The second provided applications of this theory to the face boundary space of an embedded graph-like space and the cycle space of a graph-like space. Further we were able to provide a full characterization of the graph-like spaces embeddable in the plane. We conclude this thesis with some ideas for further research into graph-like spaces, focusing on embeddings of graph-like spaces in surfaces.

7.1 Future Research

Graph-like spaces successfully merge the combinatorial properties of infinite graphs with the topological tools needed to prove natural analogues of classical theorems. The theory of graph-like spaces is in its infancy, and the list of questions is seemingly endless. In this section we present three questions, and one conjecture, about graph-like spaces that arise naturally from the preceding chapters.

In Chapter 4 we discussed the algebraic edge space of a graph-like space. These properties of algebraic edge spaces were largely inherited from the edge space theory of Vella and Richter. The main results from that chapter considered embeddings of graph-like spaces in arbitrary surfaces, and applied the topological

lemma from Chapter 2. We can ask whether the same results hold for a more general class of edge spaces embedded in a surface. In particular, we can consider embeddings of edge spaces that are compact, and have the property that connected subsets are arcwise connected. For example consider the space E defined as follows. Take the set of circles $C_i = \text{Bd}(B(0, 1 - 1/2^i))$ for $i \in \mathbb{N}$ and the set of circles $C'_i = \text{Bd}(B(0, 1 + 1/2^i))$ for $i \in \mathbb{N}$ in the plane. Both of these sets of circles converge to the unit circle C centred at the origin. Construct E by taking the union of all of the C_i and C'_i together with C and the subarcs of the x and y axis in the plane between $3/2$ and $-3/2$. The resulting space is depicted in Figure 7.1.




Figure 7.1: Edge space E embedded in \mathbb{R}^2 .

The space E is not graph-like because the circle C consists entirely of vertices. However, since E is a topological edge space, there is a well-defined cycle space $\mathcal{Z}_t(E)$ of E . Note that the bad set of vertices forms a boundary in the plane that partitions the edge set. None of the points in C lie in faces of the embedding, so the face boundaries are all graph-like (in fact they are all finite cycles). Furthermore, $\mathcal{Z}_t(E) \subseteq \mathcal{B}_t(E)$, since each cycle C' in E is a simple closed curve in the plane. Even if C' contains a subarc of C , it still partitions the faces of

E into two parts, each of which sums to the edge set of C' . Furthermore with slight alteration, we can apply the arguments from Chapter 4 to show that $\mathcal{Z}_t(E) \supseteq \mathcal{B}_t(E)$.

There may be a more general class of edge spaces for which MacLane's Theorem applies. The class of "graph-like spaces with boundaries" may be an interesting collection of spaces. This example inspires the following question.

Question 7.1

Is there a more general class of edge spaces for which MacLane's Theorem applies?

A simple consequence of the general version of Kuratowski's Theorem is that any acyclic graph-like space can be embedded in the plane. Acyclic graph-like spaces are a part of a larger class of topological spaces called *dendrites*. In [11], Nadler shows that every dendrite can be embedded in the plane by exhibiting a universal dendrite, due to Wazewski. A universal dendrite is a dendrite D so that if D' is any dendrite, then D contains a subspace homeomorphic to D' . Nadler constructs an explicit embedding of D in the plane, hence demonstrating that all dendrites are planar. This fact is interesting since, given an acyclic graph-like space T , we can use Nadler's construction to specify an explicit embedding of T .

Unfortunately, Kuratowski's Theorem tells us nothing about specific embeddings of planar spaces. In [9] Mohar and Thomassen present a proof of MacLane's Theorem for finite graphs that involves explicit embeddings. They outline a procedure for constructing an embedding a finite graph G given a 2-basis of $\mathcal{C}(G)$. Both of these points suggest that there may be a way to construct explicit embeddings of graph-like spaces in the plane.

Question 7.2

Given a thumbtack-free graph-like space G and any 2-basis \mathcal{B} of $\mathcal{Z}_t(G)$, is there an embedding of G in the plane so that each $B \in \mathcal{B}$ is the edge set of a face boundary?

In Chapter 6 we defined the class of thumbtack-like graph-like spaces. We proved that not only are these spaces obstructions for planar embedding, but they are the only added obstructions for planar embedding given the obstructions for finite graphs. It is clear that thumbtack-like spaces are also obstructions to embedding graph-like spaces in any 2-manifold. For any finite graph G , there is some surface Σ so that G can be embedded in Σ . The genus problem is the problem of finding the smallest genus surface Σ so that the graph G can be embedded in Σ .

Question 7.3

Is the genus problem well-defined for thumbtack-free graph-like spaces?

Finally, we noted in Section 5.1 that if G is a graph-like space embedded in the plane, then $\mathcal{B}_t(G) = \mathcal{Z}_t(G)$. The space G' embedded in the torus shown in Figure 5.1 demonstrates that this does not hold for arbitrary surfaces. In particular, the face boundaries of G' only generate cycles of G' that are contractible curves in the torus. We have that $\mathcal{B}_t(G')$ is a proper subspace of $\mathcal{Z}_t(G')$, so we can consider the quotient of the spaces. We define $\dim(\mathcal{Z}_t(G)/\mathcal{B}_t(G))$ to be the size of the smallest set $B \subset \mathcal{Z}_t(G) - \mathcal{B}_t(G)$, such that $B \cup \mathcal{B}_t(G)$ generates $\mathcal{Z}_t(G)$. From the example G' it seems that we should be able to generate every cycle of G' given the face boundaries of G' together with one non-contractible cycle of G' for each homotopy class of the torus. Given any surface Σ , Σ is either homeomorphic to the orientable surface of genus n , \mathbb{S}_n , or the non-orientable surface of genus n , \mathbb{N}_n . Define $g(\Sigma)$ as $g(\Sigma) = 2n$ if Σ is homeomorphic to \mathbb{S}_n , or $g(\Sigma) = n$ if Σ is homeomorphic to \mathbb{N}_n .

Conjecture 7.4

If G is a connected graph-like space embedded in the surface Σ so that every face of G is homeomorphic to an open disk, then $\dim(\mathcal{Z}_t(G)/\mathcal{B}_t(G)) = g(\Sigma)$.

This conjecture is true for finite graphs (see [13]).

Bibliography

The numbers at the end of each entry list pages where the reference was cited. In the electronic version, they are clickable links to the pages.

- [1] C. Paul Bonnington and R. Bruce Richter. Graphs embedded in the plane with a bounded number of accumulation points. *J. Graph Theory*, 44(2):132–147, 2003. 3
- [2] Henning Bruhn and Maya Stein. MacLane’s planarity criterion for locally finite graphs. *J. Combin. Theory Ser. B*, 96(2):225–239, 2006. 98
- [3] Karel Casteels and R. Bruce Richter. The bond and cycle spaces of an infinite graph. *J. Graph Theory*. To Appear. 3, 73
- [4] Charles O. Christianson and William L. Voxman. *Aspects of Topology*. BCS Associates, Moscow, Idaho, second edition, 1998. 7, 9, 35, 38, 44, 68, 103
- [5] Reinhard Diestel. *Graph Theory*. Springer, Berlin, third edition, 2006. 93, 100
- [6] Reinhard Diestel and Daniela Kühn. On infinite cycles. I. *Combinatorica*, 24(1):69–89, 2004. 2
- [7] Reinhard Diestel and Daniela Kühn. On infinite cycles. II. *Combinatorica*, 24(1):91–116, 2004. 2
- [8] Reinhard Diestel and Daniela Kühn. Topological paths, cycles and spanning trees in infinite graphs. *Europ. J. Combin.*, 25:835–862, 2004. 2
- [9] Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. The Johns Hopkins University Press, Baltimore, 2001. 37, 129

- [10] James R. Munkres. *Topology*. Prentice Hall, second edition, 2000. 6, 7, 24, 94
- [11] Sam B. Nadler. *Continuum Theory*. Marcel Dekker, Inc., New York, 1992. 129
- [12] Ian Richards. On the classification of non-compact surfaces. *Trans. Amer. Math. Soc.*, 106:259–269, 1963. 77
- [13] R. Bruce Richter and Herbert Shank. The cycle space of an embedded graph. *J. Graph Theory*, 8(3):365–369, 1984. 130
- [14] R. Bruce Richter and Carsten Thomassen. 3-connected planar spaces uniquely embed in the sphere. *Trans. Amer. Math. Soc.*, 354(11):4585–4595, 2002. 102, 110
- [15] Lynn Arthur Steen and J. Arthur Seebach. *Counterexamples in Topology*. Springer-Verlag, Berlin, second edition, 1978.
- [16] Carsten Thomassen. The Jordan-Shönflies Theorem and the classification of surfaces. *Amer. Math. Monthly*, 99(2):116–131, 1992. 37, 38
- [17] Carsten Thomassen. The locally connected compact metric spaces embeddable in the plane. *Combinatorica*, 24(4):699–718, 2004. 98, 101, 103
- [18] Carsten Thomassen and Antoine Vella. Graph-like continua, augmenting arcs, and Menger’s Theorem. *Combinatorica*. To Appear. 3, 5, 8, 9, 11, 12, 13
- [19] Antoine Vella and R. Bruce Richter. Cycle spaces in topological spaces. *J. Graph Theory*. To Appear. 2, 3, 73, 75, 76

Index

Page numbers printed like **123** indicate definitions. Other relevant page numbers are printed like 123.

- 2^E , **75**
- $B(0, 1)$, **35**
- $B(x, \epsilon)$, **6**
- $\text{Bd}^*(F)$, **79**, 84
- \mathbb{S}^1 , **35**
- \mathbb{S}^2 , **33**
- $B_t(K)$, **79**, 83, 85, 91, 94, 96, 129, 130
- $B_t(K_i)$, 87
- $C(G)$, **93**
- $Z_t(G)$, **76**, 91, 94, 98–100, 125, 126, 129, 130
- $Z_t(X, E)$, **76**
- 01-string, 120
- 1-point compactification, *see*
 - Alexandroff compactification
- 2-basis, 93, **95**, 96, 98–100, 125, 126, 129
- 2-connected, 98, **98**, 100, 107
- 2-manifold, **33**
- Alexandroff compactification, **11**, 108
- arc, **7**
 - adjacent, **52**
 - endpoint, **7**
- artificial-block, **106**
- basis, **95**
- block, 106
- boundary component, **77**
 - orientable, **77**
 - planar, **77**
- bounded, **6**
- bridge, **109**
- circle, 15, 16, 71, 72, 89, 103, 104, 109
- closed curve, **15**
 - simple, **15**
- compact, **6**
- connected, **6**
 - arcwise, **7**, 34
 - hereditarily locally, **12**
 - locally, **12**
 - locally arcwise, **7**, 34
- cuff, 120, **120**
- cut-vertex, **106**, 108, 111, 113, 116
- cycle, 109, 116, 129
 - graph-like space, **76**
 - topological edge space, **76**
 - trivial, **89**
- cycle space, **93**

- dendrite, 129
- diameter, **6**, 11
- distance, **6**
- edge, **9**, 12, 36, 37, 45, 46, 108
- edge cycle, **76**
- edge space, **75**
 - algebraic, **75**, 128
 - topological, **73**, 74
- edge-block, **106**
- edge-cut, 76, 103
- embedding, **33**
- face, **34**, 36, 46, 79, 102, 113
 - adjacent, **81**
 - boundary, 35, 46, 71, 72
 - non-disk, **72**, 84
 - simply connected, 35
- flower, 52
 - n -flower, 47, **47**
 - infinite, **52**
- Freudenthal compactification, 11, **11**, 108
- generating set, **95**
- graph, **1**
 - natural topology, 7
- graph-like space, 7, 9, **9**, 10, 11, 15, 75, 76, 88, 94, 102, 107, 121, 127
 - acyclic, 129
 - disconnected, 14, 71, 84, 124, 126
 - genus, 130
 - thumbtack-free, 111, **111**, 113, 124, 126
 - vertex set, 78
 - with boundaries, 129
- Hausdorff, **6**
 - weakly, **73**
- ideal boundary, 77, **77**
- ideal boundary point, **77**
- infinite binary tree, 108
- infinite graph, **2**, 10
 - end, **10**
- Jordan Curve Theorem, 94
- Jordan-Schönflies Theorem, 37
- Kuratowski's Theorem, 98, 101, 124, 125
- locally finite, **10**
- MacLane's Theorem, 93, 100, 126, 129
- metric, **6**
- metrizable, **6**
- overlap, **109**
- petal, 47, **47**
 - adjacent, **52**
- plane, 33
- plane isomorphism, **37**
- point of attachment, **109**
- ray, **10**
 - tail, **10**
- real-block, **106**, 107, 108, 111, 113, 116
- separated, **6**
- separation, **6**
- simply connected, 35
- span
 - algebraic, **75**

- strong, **75**
- weak, **75**
- sphere, 33
- subspace, **75**
- support, **75**
- surface, **33**, 77
 - bounded, **77**
 - genus, 130
 - homotopy, 130
 - non-orientable, 130
 - orientable, 130
- thin collection, **75**
- thin cycle space, 76, **76**, 87
- thin sum, **75**
- thin summation, **75**, 83, 85
- thumbtack space, **101**
- thumbtack-like space, **110**, 129
- torus, 95, 130
- totally disconnected, 9, **9**, 78
- trail, **14**
- vertex, **9**
- walk, 14
- web, **109**, 126
- zero-dimensional, **8**, 9