Moore's Theorem and Zippin's Sphere Characterization

Jack Hackett
Supervisor: Dr. Daniel Meyer

Thesis submitted in accordance with the requirements of the University of Liverpool
for the degree of Master of Philosophy

October 2020

Abstract

Moore's theorem gives sufficient–but not necessary–conditions that a quotient of the 2-sphere S^2 itself be homeomorphic to S^2. In Chapter 1, we give a proof of this classical result by modern means, avoiding explicit use of Moore's original axiomatic approach, which is spread across a number of papers and is rather inaccessible to today's reader. Two applications of Moore's theorem to the study of dynamical systems are also sketched.

Our proof makes essential use of Zippin's characterization of S^2 amongst the class of Peano continua. Much as with Moore's theorem, the original references for this result can be difficult to read from a modern perspective, and so we use Chapter 2 to give a self-contained proof of this theorem.

This self-containedness is contingent upon a variety of results from the theory of continua, and particularly germane is Whyburn's characterization of cyclically connected Peano continua. Appendix A reviews the necessary elements of continuum theory, culminating in a proof of this characterization.

Since our proof of Moore's theorem also relies on a special case of Alexander duality, Appendix B presents a proof of this special case by elementary means.

Acknowledgements

I would like to thank Daniel Meyer for his guidance, support, and seemingly limitless patience throughout the preparation of this thesis. His talent for explaining problems clearly, and his infectious enthusiasm for mathematics, made this project a pleasure to work on, and his help–especially in the socially-distanced climate in which this thesis was written–is greatly appreciated.

I also wish to thank Nicola Pagani, who originally put Daniel and myself in contact, for advice and assistance when it was dearly needed, and for helping to find a solution where I thought one could not be found–and of course, for his help in navigating the bureaucratic necessities.

Finally, I express my gratitude to my friends and family, who have been there for me through thick and thin. Always willing to lend an ear, be that to the latest source of fascination or to whatever obstacle was frustrating me most recently, I could not have done this without them.

Notation and Terminology

For reference, we include below a summary of largely-standard notation and terminology which is not defined explicitly in the main body of text; in particular, we draw the reader's attention to the notational conventions used for disconnected spaces and arcs.

We assume that the standard definitions of topological spaces, quotient spaces, metrizable spaces, covers of a space, and exactness of a sequence are all known.

We also remark that, given sets A and B, we use the notation $A \subset B$ to indicate that every element of A is also an element of B; in particular, this includes the possibility that $A = B$. If we wish to emphasise that A is a strict subset of B, so that the complement $B - A$ is non-empty, we will write $A \subsetneq B$.

Standard Spaces

Throughout, we denote by $\mathbb{N} = \{1, 2, 3, \ldots\}$ the set of natural numbers, and by $\mathbb{Z}$ the set of integers. For each $n \in \mathbb{N}$, we define $\mathbb{R}^n$ to be the real n-space, and unless otherwise specified, we view $\mathbb{R}^n$ as a topological space equipped with the standard (Euclidean) topology. The Euclidean norm on $\mathbb{R}$ is denoted by $|\cdot|$, and the Euclidean norm on $\mathbb{R}^n$ for $n \geq 2$ is denoted by $\|\cdot\|$. For each $n \in \mathbb{N}$, the n-sphere is identified as the subspace

$$S^n = \left\{x \in \mathbb{R}^{n+1} \,\middle|\, \|x\| = 1\right\} \subset \mathbb{R}^{n+1},$$

while the n-cell is identified as the subspace

$$D^n = \{x \in \mathbb{R}^n \,|\, \|x\| \leq 1\} \subset \mathbb{R}^n,$$

and we equip each of these with the standard metric inherited from $\mathbb{R}^n$.

We use $\mathbb{C}$ to denote the complex plane, understood to be topologized in the usual fashion. Given a point $z \in \mathbb{C}$, we denote by $\overline{z}$ its complex conjugate.

Closed intervals in $\mathbb{R}$ are denoted by $[a, b]$, and open intervals by $]a, b[$. Accordingly, half-open intervals are denoted by $[a, b[$ and $]a, b]$.

The empty set is denoted by $\emptyset$.

Topological Concepts

Many of these conventions are broadly similar to those of [Why45] and [Wil49]. Throughout this section, we understand X to denote some topological space.

If $A \subset X$ denotes some subspace, then the *interior of A in X* is denoted by

$$\operatorname{int}_X A = \bigcup \{U \subset X \,|\, U \subset A \text{ and } U \text{ is open}\},$$

and the *closure of A in X* is denoted by

$$\operatorname{cl}_X A = \bigcap \{C \subset X \,|\, A \subset C \text{ and } C \text{ is closed}\}.$$

In turn, the *boundary of* A *in* X is the intersection

$$\operatorname{fr}_X A = \operatorname{cl}_X A \cap \operatorname{cl}_X(X - A).$$

Where there is no confusion as to the choice of ambient space, we omit the subscripts, writing simply $\operatorname{int} A$, $\operatorname{cl} A$ and $\operatorname{fr} A$.

If there exist subspaces $A, B \subset X$ such that $(\operatorname{cl}_X A) \cap B = A \cap (\operatorname{cl}_X B) = \emptyset$, and such that $X = A \cup B$, then we say that X is *disconnected.* In this situation, we call $X = A \mid B$ a *separation (of* X*).*

The connected subspaces of X which are maximal with respect to set-theoretic inclusion are the *(connected) components* of X. If $A \subset X$ denotes some subspace, then the components of the complement $X - A$ are termed the *complementary domains of* A *in* X. If the complement $X - A$ is disconnected, then we say that A *separates* X.

If some open cover of X has a finite subcover, then that cover is termed *essentially finite.* If every open cover of X is essentially finite, then X is *compact.*

If we have some equivalence relation $\sim$ on X, then we denote by $X/\sim$ the associated quotient space. A subspace $A \subset X$ which is a union of equivalence classes of $\sim$ is said to be *saturated (with respect to* $\sim$*).* The notation

$$[x] = \{y \in X \mid x \sim y\} \subset X$$

is used for the equivalence class of a point $x \in X$ with respect to $\sim$.

Given some other topological space Y, a continuous function $f\colon X \to Y$ is said to be *open* (*closed*) if the image under f of every open (closed) subspace of X is open (closed) in Y. If X and Y are homeomorphic to one another, then we indicate this using the shorthand notation $X \cong Y$.

By a *Jordan curve* in X, we mean some homeomorphic image of the circle S^1 embedded in X.

A *path* in X is some continuous function $\gamma\colon [0,1] \to X$. The points $\gamma(0), \gamma(1) \in X$ are called the *end points of* γ, and we say that γ is a path *from* $\gamma(0)$ *to* $\gamma(1)$. If both end points of γ are the same, then γ is said to be a *loop.* If γ is injective, then we call γ (or its image) an *arc from* $\gamma(0)$ *to* $\gamma(1)$.

If $T \subset X$ is some arc, considered as the image of some continuous function $\gamma\colon [0,1] \to X$, then we use the notation $]T[= \gamma(]0,1[) = T - \{\gamma(0), \gamma(1)\}$.

In general, given any pair of points $x, y \in X$, there of course exist many different arcs in X from x to y. Sometimes, however, we are only interested in the existence of some such arc, with specifics not being relevant. In this situation, if γ is some arc from x to y, we may use the notation $[xy] = \gamma([0,1])$. By analogy to our conventions for intervals in $\mathbb{R}$, we further define $]xy[= \gamma(]0,1[)$, and similarly for images $[xy[$ and $]xy]$ of the respective half-open intervals.

If we wish to denote an arc from x to y in X which passes through points $p_1, p_2, \ldots, p_n \in X$ in that order, then we may use the notation $[xp_1p_2 \cdots p_ny] \subset X$.

If $C \subset X$ denotes some compact subspace, an arc $[xy] \subset X$ is said to *span* C if $[xy]$ intersects C only at the points x and y.

If, for each pair of points $x, y \in X$, there exists some path (arc) in X from x to y, then X is said to be *path (arc) connected.* The path (arc) connected subspaces of X which are maximal with respect to set-theoretic inclusion are termed the *path (arc) components* of X.

The space X is said to be *locally (path, arc) connected* if each point of X permits a neighbourhood basis consisting entirely of open (path, arc) connected subspaces.

If X is metrizable, then any metric d which induces the topology of X is said to *topologize* X. If a metric d topologizes X, then the open ε-ball centred at a point $x \in X$ with respect to d is denoted by

$$B_d(x; \varepsilon) = \{y \in X \mid d(x,y) < \varepsilon\},$$

for each $\varepsilon > 0$. If the metric d can be inferred from its context without confusion, then we omit the subscript and write simply $B(x; \varepsilon)$. Given a subspace $A \subset X$, the *diameter of* A *with respect to* d is then defined to be the real number

$$\operatorname{diam}_d A = \sup_{x,y \in A} d(x, y),$$

and again, we write simply $\operatorname{diam} A$ if the choice of metric d is clear from context.

Algebraic Concepts

To indicate that two groups G_1 and G_2 are isomorphic, we use the shorthand $G_1 \cong G_2$.

For a topological space X and any $n \in \mathbb{N} \cup \{0\}$, we denote the n^{th} singular homology group of X by $H_n(X)$.

Suppose that X denotes some topological space with a pair of subspaces $A, B \subset X$ such that $X = \operatorname{int} A \cup \operatorname{int} B$. There then exists a long exact sequence

$$\cdots \longrightarrow H_{n+1}(X) \longrightarrow H_n(A \cap B) \longrightarrow H_n(A) \oplus H_n(B) \longrightarrow H_n(X) \longrightarrow \cdots$$

of singular homology groups, where $\oplus$ denotes the group direct sum. We call this the *Mayer-Vietoris sequence of the triad* (X, A, B), with further details being available in, for example, [Lee00].

Contents

Chapter 1

Moore's Theorem

1.1 Introduction

Our object in this first chapter is a proof of a theorem of R.L. Moore (Theorem 1.1.2.) Before we state this result, however, we pause to define and characterize[1] *upper semicontinuity*, in terms of which Theorem 1.1.2 is most naturally stated.

Definition and Lemma 1.1.1. *Let $\sim$ denote some equivalence relation on a topological space X with the property that, for each point $x \in X$, the equivalence class $[x] \subset X$ is compact.*

The following are equivalent:

- *For each open subspace $U \subset X$, the* saturated interior $U^* = \bigcup\{[x] \mid x \in X \text{ and } [x] \subset U\}$ *is also open;*
- *For each closed subspace $C \subset X$, the* saturated closure $C^\dagger = \bigcup\{[x] \mid x \in C\}$ *is also closed; and*
- *The natural projection $\pi \colon X \to X/\sim$ is a closed map.*

An equivalence relation with the above properties is said to be upper semicontinuous.

Proof. Notice that, for any subspace $A \subset X$, we have the chain of equalities

$$\begin{aligned}(X - A)^* &= \bigcup\{[x] \mid x \in X \text{ and } [x] \cap A = \emptyset\} \\ &= X - \bigcup\{[x] \mid [x] \cap A \neq \emptyset\} \\ &= X - \bigcup\{[x] \mid x \in A\} \\ &= X - A^\dagger.\end{aligned} \tag{1.1}$$

This immediately establishes the equivalence of the first two conditions above. Noticing additionally that

$$A^\dagger = \left(\pi^{-1} \circ \pi\right)(A) \tag{1.2}$$

for any subspace $A \subset X$ shows the final two conditions above to be equivalent as well. This completes the proof. ☐

Heuristically, we can view the upper semicontinuity of an equivalence relation $\sim$ on a topological space X as a guarantee that the quotient $X/\sim$ is, in some sense, 'well-behaved'; for instance, one can show (Lemma 1.4.3) that if X is a separable metrizable space, and if $\sim$ is upper semicontinuous, then $X/\sim$ is also separable and metrizable. A book by Daverman [Dav86] discusses upper semicontinuity in considerable detail.

[1] In the compact metric setting of interest to us here, there are in fact many more equivalent characterizations of upper semicontinuity than the three given in Lemma 1.1.1. A treatment can be found in [PM13].

In the language of Definition 1.1.1, then, the theorem in which we are interested can be stated thus.

Theorem 1.1.2 (Moore's theorem). *Let $\sim$ denote some upper semicontinuous equivalence relation on S^2 with at least two distinct equivalence classes, If, for every point $x \in S^2$,*

- *The equivalence class $[x] \subset S^2$ is connected, and*
- *The complement $S^2 - [x]$ is connected,*

then the quotient $S^2/\sim$ is homeomorphic to S^2 itself. □

In Section 1.2, we loosely describe how this result was first obtained; however, Moore's original proof spans several papers [Moo15, Moo16, Moo25] and is written in language difficult for a modern reader to penetrate. The proof we present here differs substantially from Moore's approach, and is outlined in Section 1.4.1. (Our method of proof adheres to a strategy laid out by Cannon in [Can78].)

Before explicitly proving Moore's theorem, we take some time in Section 1.3 to develop some intuition of why the hypotheses of Theorem 1.1.2 should be what they are, along with an illustration of Moore's theorem in practice. In the same section, we also supply an example of how the converse of Theorem 1.1.2 is not true: the conditions of Moore's theorem are sufficient for a quotient of S^2 to be homeomorphic to S^2, but they are certainly not necessary.

After the aforementioned proof sketch in Section 1.4.1, we use Sections 1.4.2-1.4.4 to make our proof explicit. Finally, we use Section 1.5 to give some indication of how Theorem 1.1.2 is of utility in other fields of mathematics, describing two applications of this result to the study of dynamical systems.

1.2 Historical Background

In [Moo25], Moore proved Theorem 1.1.2 by exploiting an axiomatic characterisation [Moo16] of the Euclidean plane $\mathbb{R}^2$, ultimately derived from earlier work by himself [Moo15] and by Veblen [Veb04], Moore's doctoral supervisor. For their historical interest, we shall briefly review these axioms here, although in an effort to render the source material more readable to a modern audience, we adjust the notation and language of [Moo16] to better reflect the past century's developments. (In fact, we take an additional liberty, and use a reduced set of axioms due to Wilder: the axioms as stated in [Moo16] include another, shown in [Wil27] to be superfluous, which is of a similar flavour to the sixth axiom below.)

Moore begins with the primitive notions of a set Π of points, and some distinguished collection $\mathcal{T}$ of subsets of Π, which he terms *regions*. Then, given any subset $X \subset \Pi$, a point $p \in \Pi$ is classified as a *limit point of* X if, for each region $U \in \mathcal{T}$ which contains p, the intersection $U \cap (X - \{p\})$ is non-empty; in turn, the subset $\operatorname{cl} X \subset \Pi$ is defined as the union of X with its set of limit points. The *boundary* of X, denoted $\operatorname{fr} X$, is the set of limit points of X which do not belong to X itself. Further, the subset X is said to be *connected* if it cannot be expressed as the union of two disjoint subsets, neither of which contains a limit point of the other.

Using these definitions, Moore finds that any pair $(\Pi, \mathcal{T})$ satisfying the following axioms must in fact be topologically equivalent to the plane:

1. There exists some countably-infinite collection $\{B_n\}_{n\in\mathbb{N}} \subset \mathcal{T}$ of regions such that
    - For any $n \in \mathbb{N}$ and for any point $p \in \Pi$, there exists some $m > n$ such that $p \in B_m$; and
    - Given any region $U \in \mathcal{T}$ and any distinct pair of points $p, q \in U$, there exists some $N \in \mathbb{N}$ such that if $n > N$ and $p \in B_n$, then $\operatorname{cl} B_n \subset \Pi - \{q\}$;
2. Every region $U \in \mathcal{T}$ is connected, as is the complement $\Pi - \operatorname{cl} U$;
3. Given any region $U \in \mathcal{T}$, the subset $\operatorname{cl} U$ is compact;
4. There exists an infinite subset of Π with no limit point;

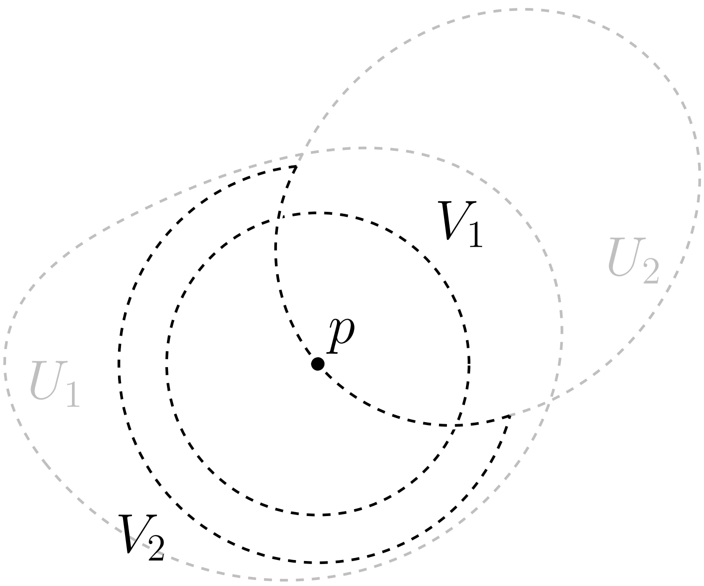

Figure 1.1

5. Every Jordan curve[2] in Π is the boundary of some region;
6. Given regions $U_1, U_2 \subset \mathcal{T}$ and some point $p \in U_1 \cap \operatorname{cl} U_2$, there exists a pair of regions $V_1, V_2 \subset U_1$ such that
   - $p \in V_1$,
   - $V_2 \subset \Pi - \operatorname{cl} U_2$, and
   - $V_1 \cap \operatorname{fr} U_2 \subset \operatorname{fr} V_2$,

   as schematically depicted in Figure 1.1.

The parallels between Moore's regions and the open subspaces defined by a topology are obvious, and it is easy to see how the above notion of limit points and connectedness are direct translations of their modern formulations. It is worth noting, however, that Moore's verbosity in the above axioms is not redundant: at the time, neither abstract topological spaces nor open subspaces (in the modern sense) had a standardised definition. The current definitions of these concepts can be traced back to Sierpiński's work in the 1920s and 1930s [Sie28, Sie34], which in turn drew from earlier work by, most notably, Hausdorff, Tietze and Kuratowski [Hau14, Tie22, Kur22, Tie23].[3]

In 1930, Zippin [Zip30] substantially condensed the above axioms, yielding in particular the following theorem, about which our proof of Theorem 1.1.2 shall revolve. To state the following theorem more concisely, we use the language of Peano continua, which we recall to be precisely the compact, connected, locally connected metrizable spaces; in Appendix A, these spaces are discussed in depth.

Theorem 1.2.1 (Zippin's sphere characterization)**.** *Let X denote some Peano continuum (Definition A.2.1). If X satisfies the Jordan curve theorem in the sense that*

- *X contains a Jordan curve,*
- *Every Jordan curve in X separates X, and*

[2]In [Moo16], Moore defines a Jordan curve as the union of a pair of arcs between a fixed pair of points, disjoint except for at their end points. He uses an intrinsic characterization of arcs similar in essence to Lemma A.4.5.

[3]The development of the modern notion of a topological space over the first half of the twentieth century is recounted in [Moo08].

- *No closed arc in X separates X,*

then X is homeomorphic to S^2. □

This result is itself far from trivial to prove, with such a proof occupying all of Chapter 2 and necessitating the development of a number of continuum-theoretic results in Appendix A. Taking Theorem 1.2.1 as granted for the time being, our proof of Moore's theorem essentially reduces to demonstrating that a quotient $S^2/\sim$ of the kind described in Theorem 1.1.2 satisfies the hypotheses of Zippin's sphere characterization.

1.3 Motivation for the Hypotheses

To gain some idea of why the hypotheses of Moore's theorem are plausible, let us consider some extremely simple quotients of the sphere, determining in each case whether or not they are themselves homeomorphic to the sphere.

Example 1.3.1 (Identification of the closed upper hemisphere to a point)**.** Suppose that we define a quotient map $\pi\colon S^2 \to X$ by letting the only non-degenerate equivalence class be the closed upper hemisphere $H \subset S^2$. What can we say about X?

The dense subspace $\pi(S^2 - H) \subset X$ is of course homeomorphic to the plane, while $\pi(H)$ is, by definition, just a single point. In other words, X is an Alexandroff compactification of the plane–but so too is S^2. Knowing (from, for example, [Wil70]) that the Alexandroff compactification of any Hausdorff space is unique up to homeomorphism, we conclude that $X \cong S^2$. □

Example 1.3.2 (Identification of the open upper hemisphere to a point)**.** Now, let us assume instead that $\pi\colon S^2 \to X$ were defined by taking the open upper hemisphere $U \subset S^2$ as the only non-degenerate equivalence class. The singleton subspace $\pi(U) \subset X$ is necessarily open in X, because $(\pi^{-1} \circ \pi)(U) = U$ is open in S^2; thus, X is not even T_1, much less homeomorphic to S^2.

In particular, we observe that in this case, the quotient map π is not closed, so that the corresponding equivalence relation cannot be upper semicontinuous. □

Example 1.3.3 (Identification of the 'half-open' upper hemisphere to a point)**.** Define the 'half-open' upper hemisphere $A \subset S^2$ to be the union of the open upper hemisphere with some closed arc of the equator. If A is the only non-degenerate equivalence class of the quotient $\pi\colon S^2 \to X$, then the singleton $\pi(A) \subset X$ fails to be closed in X, since the complement $S^2 - (\pi^{-1} \circ \pi)(A) = S^2 - A$ is not open in S^2. Much as in Example 1.3.2, we see that X and S^2 are not homeomorphic. □

Example 1.3.4 (Identification of the north and south poles)**.** Now suppose that the only non-degenerate equivalence class of $\pi\colon S^2 \to X$ is some two-point subspace $\{N, S\} \subset S^2$, where we assume N and S to denote some pair of antipodal points of the sphere, thought of as the north and south poles respectively.

It seems unlikely that $X \cong S^2$ in this case, and we can prove this explicitly. Indeed, the space $X - \{\pi(N)\}$ is homeomorphic to $S^2 - \{N, S\}$, which has the homotopy type of the punctured plane, while $S^2 - \{x\}$ has the homotopy type of the plane, for any point $x \in S^2$. In particular, $X - \{\pi(N)\}$ fails to be contractible, while the complement in S^2 of any point is contractible. □

Example 1.3.5 (Identification of the equator to a point)**.** If we define a quotient map $\pi\colon S^2 \to X$ by taking the equator $E \subset S^2$ as our only non-degenerate equivalence class, then a similar analysis to that of Example 1.3.4 reveals that X and S^2 cannot possibly be homeomorphic: the complement $X - \{\pi(E)\}$ is disconnected, whereas no single point separates S^2. By entirely analogous reasoning, we see that the result of collapsing any subspace of S^2 with disconnected complement cannot be homeomorphic to S^2. □

By comparing Examples 1.3.1-1.3.3 with Definition 1.1.1, we can get some sense of why the statement of Moore's theorem presupposes upper semicontinuity. Examples 1.3.4 and 1.3.5, meanwhile, illustrate some potential obstructions to a homeomorphism between S^2 and some quotient $S^2/\sim$ when the equivalence classes of $\sim$ are either disconnected or separate S^2.

However, it is worth emphasising that, although Moore's theorem gives sufficient conditions for a quotient $S^2/{\sim}$ to be homeomorphic to S^2, these conditions are not necessary. What follows is a counterexample to the converse of Theorem 1.1.2.

Example 1.3.6 (The implication of Theorem 1.1.2 does not reverse)**.** This example is most readily understood by viewing S^2 as the Alexandroff compactification of the complex plane $\mathbb{C}$, with a point at infinity $\infty \in S^2$. Considering S^2 in this way, let us define an equivalence relation $\sim$ on S^2 by declaring that

$$z \sim w \text{ if and only if } z = \pm w. \tag{1.3}$$

Immediately, we see that the equivalence classes of 0 and ∞ are degenerate, and that the equivalence class of any point $z \in S^2 - \{0, \infty\}$ consists precisely two points, namely z and $-z$. Of course, the equivalence relation $\sim$ cannot satisfy the hypotheses of Theorem 1.1.2: almost all of its equivalence classes are disconnected.

Nevertheless, we can show that $S^2/{\sim} \cong S^2$ by supplying an explicit homeomorphism. Indeed, consider the function $S^2 \to S^2/{\sim}$ acting by $z \mapsto [\sqrt{z}]$.[4] This is manifestly a continuous bijection; further, by noticing that the equivalence relation $\sim$ is, according to Definition 1.1.1, upper semicontinuous, we find that this function is also closed. It follows that, as claimed, S^2 and $S^2/{\sim}$ are homeomorphic, despite the fact that $\sim$ does not satisfy the hypotheses of Moore's theorem. □

Our final example in this section shall be an example of how one might use Moore's theorem to identify a quotient space as a sphere, even though this might not be immediately obvious.

Example 1.3.7 (Sierpiński's carpet)**.** Suppose that we subdivide the unit square in $\mathbb{R}^2$ into nine congruent squares of equal area, and delete the interior of the middle square; call the resulting space X_1. Now, let us repeat this process for each of the eight remaining squares, calling the space that results X_2. Continuing in this way, we arrive at a space $X = \bigcap_{n\in\mathbb{N}} X_n \subset \mathbb{R}^2$, which goes by the moniker of *Sierpiński's carpet*. Some stages of this construction are sketched in Figure 1.2.

Viewing X as a subspace of S^2 via stereographic projection, we define an equivalence relation $\sim$ on S^2 which acts to collapse the closure of each complementary domain of X in S^2 to a separate point. Trivially, this equivalence relation satisfies the hypotheses of Moore's theorem, so that $S^2/{\sim} \cong S^2$, but *a priori*, this is not at all obvious. □

Before proceeding, we remark that it is perhaps surprising that the conditions required by Theorem 1.1.2 are relatively weak. In general, quotient spaces can be extremely different from the original spaces; for instance, part of the Hahn-Mazurkiewicz-Sierpiński theorem (Theorem A.6.8) asserts that every compact, connected, locally connected metrizable space is a quotient of the closed unit interval $[0, 1]$.

1.4 Proof of Moore's Theorem

1.4.1 Strategy of Proof

As alluded to in Section 1.2, the engine driving our proof of Moore's theorem is Theorem 1.2.1, with our strategy reducing to proving that the space $S^2/{\sim}$ of Theorem 1.1.2 satisfies the hypotheses of Zippin's sphere characterization. Indeed, Section 1.4.2 is dedicated to proving that $S^2/{\sim}$ is a Peano continuum, while Section 1.4.3 demonstrates that $S^2/{\sim}$ obeys the Jordan curve theorem in a suitable sense.

It is in this latter section that the utility of Zippin's reformulation of Moore's axioms becomes apparent. Theorem 1.2.1 recasts the problem as one of counting path components, and in doing so allows us to argue via singular homology, by inspecting the ranks of the relevant zeroth homology groups. Said homological arguments are enabled by the following fact, itself non-trivial, which we prove in Appendix B.

[4] Given some complex number $z = re^{i\theta} \in \mathbb{C}$, where $r \in [0, \infty[$ and $\theta \in [0, 2\pi[$, we define $\sqrt{z} = \sqrt{r}e^{i\frac{\theta}{2}}$. Additionally, we define $\sqrt{\infty} = \infty$.

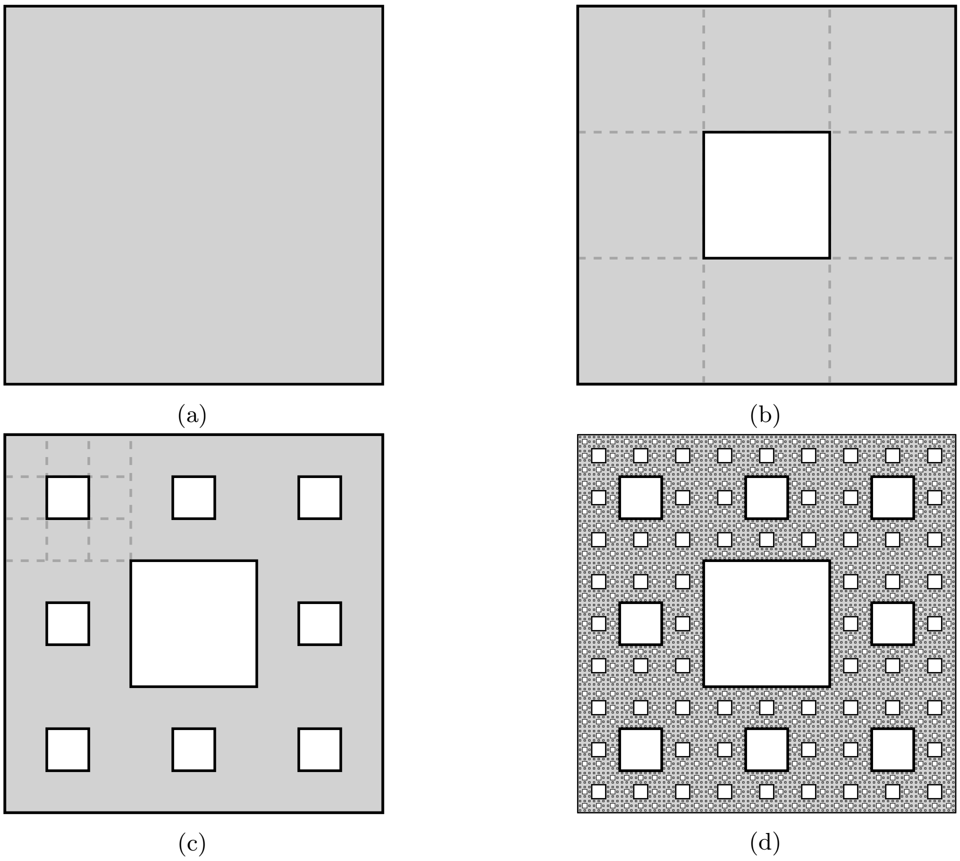

Figure 1.2: The construction of Sierpiński's carpet after (a) zero, (b) one, (c) two and (d) five iterations.

Theorem 1.4.1 (Alexander duality). *Let $C \subset S^2$ denote some compact subspace with $n \in \mathbb{N}$ connected components. Then, there exists an isomorphism $H_1\big(S^2 - C\big) \cong \mathbb{Z}^{n-1}$.* □

1.4.2 Quotients as Peano Continua

Our goal in this section is to demonstrate that the space $S^2/\sim$ of Theorem 1.1.2 is a Peano continuum. In other words, we wish to prove that $S^2/\sim$ is compact, connected, metrizable, and locally connected, and of course the first two properties are trivial to verify: the space $S^2/\sim$ is, by definition, a continuous image of the compact, connected space S^2.

To establish metrizability, we start by recalling *Urysohn's metrization theorem*, although we shall not concern ourselves here with an explicit proof of this classical result. (One implication is essentially immediate, since every metrizable space is regular and every separable space is second countable; the reverse implication may be shown by demonstrating that every regular second countable space can be embedded in the Hilbert cube. An explicit argument can be found in, for instance, [Wil70].)

Using Urysohn's metrization theorem, our proof of the metrizability of $S^2/\sim$ can be reduced to a proof of its second countability.

Theorem 1.4.2 (Urysohn). *For any topological space X, the following are equivalent:*

1. *X is regular and second countable, and*

2. *X is separable and metrizable.* □

Lemma 1.4.3. *If $\sim$ is an equivalence relation on S^2 which satisfies the hypotheses of Theorem 1.1.2, then the quotient $S^2/\sim$ is metrizable.*

Proof. Every compact Hausdorff space is regular [Wil70], so that we need only concern ourselves with the second countability of $S^2/{\sim}$. To this end, let $\mathcal{B}$ denote some countable basis for the topology of S^2, and notice that we lose no generality in assuming that $\mathcal{B}$ is closed under finite unions: were this not the case, then we could simply replace $\mathcal{B}$ by the collection

$$\{B_1 \cup B_2 \cup \cdots \cup B_n \,|\, n \in \mathbb{N} \text{ and } \{B_1, B_2, \ldots, B_n\} \subset \mathcal{B}\}. \tag{1.4}$$

Let $\pi\colon S^2 \to S^2/{\sim}$ denote the natural projection, and consider some point x of any open subspace $U \subset S^2/{\sim}$. For each point a of the fibre $\pi^{-1}(U)$, there of course exists some basis element $B_a \in \mathcal{B}$ such that $a \in B_a \subset \pi^{-1}(U)$; thus, we arrive at an open cover $\{B_a\}_{a\in\pi^{-1}(U)}$ of the subspace $\pi^{-1}(x) \subset \pi^{-1}(U)$. By compactness, this cover must be essentially finite, and so we can exhibit some finite subset $\{a_1, a_2, \ldots, a_n\} \subset \pi^{-1}(U)$ such that

$$\pi^{-1}(x) \subset B_{a_1} \cup B_{a_2} \cup \cdots \cup B_{a_n}. \tag{1.5}$$

Introducing the notation

$$B = B_{a_1} \cup B_{a_2} \cup \cdots \cup B_{a_n}, \tag{1.6}$$

our assumption that $\mathcal{B}$ is closed under finite unions tells us that $B \in \mathcal{B}$.

Now, Lemma 1.1.1 ensures that the saturated interior B^* is open and non-empty; moreover, $\pi(B^*)$ is open in $S^2/{\sim}$ by definition of the quotient topology. In particular, we notice that $x \in B^* \subset U$, so that the collection $\{\pi(B^*)\}_{B\in\mathcal{B}}$ comprises a countable basis for the topology of $S^2/{\sim}$ □

As regards local connectedness, we notice that an immediate corollary of the above is that the quotient $S^2/{\sim}$ is Hausdorff; consequently, we can avail ourselves of Sierpiński's characterization of local connectedness for compact Hausdorff spaces in terms of *Property S*. This notion is discussed in more detail in Section A.6.1, the key statements of which we reproduce below before proving that $S^2/{\sim}$ is indeed locally connected.

Definition A.6.1. *Let X denote some topological space, with the property that every open cover of X permits a refinement by finitely many connected subspaces. Then, the space X is said to have* Property S.

Lemma A.6.2. *For any compact Hausdorff space X, the following are equivalent:*

- *X is locally connected, and*
- *X has Property S.*

Lemma 1.4.4. *If $\sim$ is an equivalence relation on S^2 which satisfies the hypotheses of Theorem 1.1.2, then the quotient $S^2/{\sim}$ is locally connected.*

Proof. Knowing that $S^2/{\sim}$ is a compact Hausdorff space, it will suffice, according to Lemma A.6.2, to prove that $S^2/{\sim}$ enjoys Property S (Definition A.6.1) in order to conclude that it is locally connected.

To this end, let $\mathcal{U}$ denote some arbitrary open cover of $S^2/{\sim}$, so that the collection of fibres

$$\pi^{-1}(\mathcal{U}) = \{\pi^{-1}(U) \,|\, U \in \mathcal{U}\} \tag{1.7}$$

is an open cover of S^2. Using Lemma A.6.2, we find a refinement $\{V_1, V_2, \ldots, V_n\}$ of $\pi^{-1}(\mathcal{U})$ by finitely many connected subspaces. The surjectivity of π ensures that $\{\pi(V_1), \pi(V_2), \ldots, \pi(V_n)\}$ is a refinement of our original cover $\mathcal{U}$ by finitely many connected subspaces, and so another appeal to Lemma A.6.2 allows us to deduce that $S^2/{\sim}$ is locally connected. □

1.4.3 The Jordan Curve Theorem for Quotients

We mentioned in Section 1.4.1 that the second half of our proof of Moore's theorem–that the space $S^2/{\sim}$ of Theorem 1.1.2 satisfies the Jordan curve theorem in the sense of Theorem 1.2.1–is at its heart a homological argument. That singular homology suffices for our purposes is a consequence of two facts, of which the first is the standard result[5] that the zeroth singular homology group of any space is free on that space's set of path components.

The second is that, in our current setting, we lose nothing by counting path components, rather than connected components, as encapsulated by the following lemma.

Lemma 1.4.5. *If X denotes some locally path connected space, then every connected open subspace of X is itself path connected.*

Proof. Letting $U \subset X$ denote some connected open subspace and distinguishing some point $x \in U$, we denote by A the path component of U which contains x. We know that U inherits the local path connectedness of X, because U is open in X, and so we infer that A is open in U.

Now, suppose towards a contradiction that A is a proper subspace of U, and define $B = U - A$. Selecting any point $y \in B$, the local path connectedness of U demands the existence of some path connected neighbourhood $V \subset U$ of y; moreover, V and A must be disjoint, for otherwise there would exist a path in U connecting x and y. However, this implies that B is open in U, yielding a separation $U = A \mid B$ and contradicting our assumption that U is connected. □

Notice that openness is essential in Lemma 1.4.5, for there exist numerous examples of non-open subspaces of locally path connected spaces which, despite being connected, fail to be path connected. Perhaps the best known of these is the so-called *topologists' sine curve*, defined as the closure in $\mathbb{R}^2$ of the graph of the function $\sin\left(\frac{1}{x}\right)$, defined on the interval $]0, 1]$.[6]

Before tackling any form of Jordan curve theorem, we will need a pair of preliminary results, of which the first is reasonably straightforward, allowing us to 'lift connectedness through quotients', in an appropriate sense. The second is a classical result of plane topology, originally due to Janiszewski [Jan13] and proved independently a few years later by Mullikin [Mul22], which often bears only the name of the former in the literature. Our proof here is in line with the homological theme of this section, using similar ideas to [New85], although avoiding the theory of 'gratings' developed in the cited work.[7]

Lemma 1.4.6. *Consider a quotient map $\pi \colon S^2 \to S^2/{\sim}$, where the equivalence relation $\sim$ satisfies the hypotheses of Theorem 1.1.2. If a subspace $A \subset S^2/{\sim}$ is connected, then so too is its fibre $\pi^{-1}(A) \subset S^2$.*

Proof. We shall argue by contraposition, supposing that the fibre $\pi^{-1}(A)$ permits some separation $\pi^{-1}(A) = U \mid V$.

Given any point $x \in A$, we propose that the fibre $\pi^{-1}(x)$ must be contained entirely within U or V. Indeed, we know by hypothesis that $\pi^{-1}(x)$ must be non-empty and connected, and since

$$\pi^{-1}(x) = \left(\pi^{-1}(x) \cap U\right) \cup \left(\pi^{-1}(x) \cap V\right), \tag{1.8}$$

it follows that precisely one of the intersections $\pi^{-1}(x) \cap U$ and $\pi^{-1}(x) \cap V$ must be empty.

In particular, both U and V are unions of point inverses, so that $\pi(U)$ and $\pi(V)$ are disjoint non-empty open subspaces of A. In other words, $A = \pi(U) \mid \pi(V)$ is a separation, implying that A is disconnected. □

Lemma 1.4.7 (Janiszewski-Mullikin)**.** *Let C_1 and C_2 denote some pair of closed subspaces of the sphere S^2 for which the intersection $C_1 \cap C_2$ is connected. If neither C_1 nor C_2 separates S^2, then nor does the union $C_1 \cup C_2$.*

[5] This is proved in almost any textbook which discusses singular homology; see, for instance, [Lee00].

[6] This space is a classical counterexample in general topology. It, and many other pathological spaces, are discussed in [SS78].

[7] Lemma 1.4.7 appears as Corollary 2 to Theorem 9.1.2 in [New85].

Proof. For the sake of a contradiction, suppose that $C_1 \cup C_2$ separates points x and y in S^2. By hypothesis, both $S^2 - C_1$ and $S^2 - C_2$ are connected, so that we can use Lemma 1.4.5 to deduce the existence of some pair of paths $\gamma_1\colon [0,1] \to S^2 - C_1$ and $\gamma_2\colon [0,1] \to S^2 - C_2$, each of which connects x to y. Of course, the formal difference $\gamma_1 - \gamma_2$ is then a singular 1-chain in $S^2 - C_1 \cap C_2$.

Consider the Mayer-Vietoris sequence of the triad $\big(S^2 - C_1 \cap C_2, S^2 - C_1, S^2 - C_2\big)$. In particular, a fragment of this sequence reads

$$\cdots \longrightarrow H_1(S^2 - C_1 \cap C_2) \xrightarrow{\partial_*} H_0(S^2 - C_1 \cup C_2) \longrightarrow \cdots. \tag{1.9}$$

Now, on the one hand, we are assuming that $C_1 \cap C_2$ is connected, so that we can appeal to Alexander duality (Theorem 1.4.1) to deduce that $H_1\big(S^2 - C_1 \cap C_2\big)$ is the trivial group, so that ∂_* must be the zero homomorphism. On the other, however, we know from the construction of the Mayer-Vietoris sequence [Lee00] that

$$\partial_*([\gamma_1 - \gamma_2]) = [\partial\gamma_1] = [y] - [x], \tag{1.10}$$

where we use square brackets to denote homology classes.

Combining these two observations, we find that $[y] - [x] = 0 \in H_0\big(S^2 - C_1 \cup C_2\big)$. Phrased more transparently, the points x and y must lie in the same path component of $S^2 - C_1 \cup C_2$–but this contradicts our initial assumption that $C_1 \cup C_2$ separates x and y. Thus, $S^2 - C_1 \cup C_2$ must be connected. ☐

With these results in hand, we are able to demonstrate that a quotient $S^2/{\sim}$ of the kind discussed in Theorem 1.1.2 is separated by none of its arcs, and by all of its Jordan curves.

Lemma 1.4.8. *Let $\sim$ denote some equivalence relation on S^2 which satisfies the hypotheses of Theorem 1.1.2. If $T \subset S^2/{\sim}$ denotes some arc, then the complement $\big(S^2/{\sim}\big) - T$ is connected.*

Proof. Let $\pi\colon S^2 \to S^2/{\sim}$ denote the natural projection. Noticing that $\big(S^2/{\sim}\big) - T$ is the continuous image of the complement $S^2 - \pi^{-1}(T)$ under π, it will suffice for us to prove that $\pi^{-1}(T)$ fails to separate S^2. We shall argue by contradiction, supposing that $\pi^{-1}(T)$ separates points $x, y \in S^2$.

Introducing some parametrization $\gamma\colon [0,1] \to T$, let us define

$$T_1 = \gamma\left(\left[0, \frac{1}{2}\right]\right) \text{ and } T_2 = \gamma\left(\left[\frac{1}{2}, 1\right]\right). \tag{1.11}$$

Since $T_1 \cap T_2$ is a singleton, we know by hypothesis that $S^2 - \pi^{-1}(T_1 \cap T_2)$ is connected. Thus, we can invoke (the contrapositive of our statement of) Lemma 1.4.7 in order to deduce that at least one of $\pi^{-1}(T_1)$ and $\pi^{-1}(T_2)$ separates S^2.

Without loss of generality, suppose that $\pi^{-1}(T_1)$ separates S^2. Of course, T_1 is itself an arc in $S^2/{\sim}$, so that we can repeat the above argument to deduce that one of the subarcs

$$\gamma\left(\left[0, \frac{1}{4}\right]\right) \text{ and } \gamma\left(\left[\frac{1}{4}, \frac{1}{2}\right]\right) \tag{1.12}$$

must separate S^2. Continuing in this fashion, we can produce a descending chain

$$[0,1] \supsetneq A_1 \supsetneq A_2 \supsetneq A_3 \supsetneq \cdots \tag{1.13}$$

of closed intervals, with the property that $(\pi^{-1} \circ \gamma)(A_n)$ separates S^2 for each $n \in \mathbb{N}$.

Moreover, $\operatorname{diam} A_n = 2^{-n}$ for each $n \in \mathbb{N}$, so that $\operatorname{diam} A_n \to 0$. It follows that the intersection $\bigcap_{n\in\mathbb{N}} A_n$ is a singleton, containing precisely one point $t \in [0,1]$. Now, our hypotheses ensure that $S^2 - (\pi^{-1} \circ \gamma)(t)$ is connected; in particular, Lemma 1.4.6 allows us to find some path $P \subset S^2 - (\pi^{-1} \circ \gamma)(t)$.

The normality of the sphere allows us to find some neighbourhood $U \subset S^2$ of $(\pi^{-1} \circ \gamma)(t)$ for which $U \cap P = \emptyset$, and we lose no generality in assuming that U is saturated: if necessary, we can just replace U with its saturated interior. Then, the image $\pi(U)$ is a neighbourhood of the point $\gamma(t)$ in $S^2/{\sim}$, so that $\gamma(A_N) \subset \pi(U)$ for sufficiently large $N \in \mathbb{N}$. Thus, $(\pi^{-1} \circ \gamma)(A_N) \subset U$.

However, the disjointness of U and P ensures that U does not separate x and y, and so we arrive at a contradiction. Indeed, although U cannot separate x and y, the subspace $(\pi^{-1} \circ \gamma)(A_N) \subset U$ by construction must. We conclude that $S^2 - \pi^{-1}(T)$, and therefore $\big(S^2/{\sim}\big) - T$, is connected. □

Lemma 1.4.9. *Let $\sim$ denote some equivalence relation on S^2 which satisfies the hypotheses of Theorem 1.1.2. If $J \subset S^2/{\sim}$ denotes some Jordan curve, then the complement $\big(S^2/{\sim}\big) - J$ is disconnected.*

Proof. We can express the Jordan curve J as the union of precisely two proper subarcs $T_1, T_2 \subset J$ which meet precisely at their end points. Introducing the condensed notation

$$T_1' = \pi^{-1}(T_1),\ T_2' = \pi^{-1}(T_2) \text{ and } J' = \pi^{-1}(J), \tag{1.14}$$

where $\pi\colon S^2 \to S^2/{\sim}$ denotes the natural projection, consider the Mayer-Vietoris sequence of the triad $(S^2 - T_1' \cap T_2', S^2 - T_1', S^2 - T_2')$; particularly, the tail of this sequence reads

$$\begin{array}{l} \cdots \longrightarrow H_1(S^2 - T_1') \oplus H_1(S^2 - T_2') \longrightarrow H_1(S^2 - T_1' \cap T_2') \\ \longrightarrow H_0(S^2 - J') \longrightarrow H_0(S^2 - T_1') \oplus H_0(S^2 - T_2') \longrightarrow H_0(S^2 - T_1' \cap T_2') \longrightarrow 0. \end{array} \tag{1.15}$$

We are interested in the rank of the group $H_0(S^2 - J')$, and we propose that we know the ranks of all other groups appearing in (1.15).

Indeed, we know from Lemma 1.4.6 that T_1' and T_2' are both connected, so that an appeal to Alexander duality (Theorem 1.4.1) reveals that

$$H_1(S^2 - T_1') \cong H_1(S^2 - T_2') \cong 0. \tag{1.16}$$

Similarly, if we let $T_1 \cap T_2 = \{x, y\}$, then

$$S^2 - T_1' \cap T_2' = S^2 - \big(\pi^{-1}(x) \cup \pi^{-1}(y)\big). \tag{1.17}$$

The point inverses $\pi^{-1}(x)$ and $\pi^{-1}(y)$ constitute a pair of disjoint, closed, connected subspaces of S^2, so that their union $\pi^{-1}(x) \cup \pi^{-1}(y)$ has precisely two connected components. Another invocation of Theorem 1.4.1 therefore tells us that the group $H_1(S^2 - T_1' \cap T_2')$ is of rank one.

We have already shown with Lemma 1.4.8 that $S^2 - T_1'$ and $S^2 - T_2'$ are both connected, so that

$$H_0(S^2 - T_1') \cong H_0(S^2 - T_2') \cong \mathbb{Z}, \tag{1.18}$$

thanks to Lemma 1.4.5, leaving us with only the group $H_0(S^2 - T_1' \cap T_2')$ to handle.

Observe that

$$S^2 - T_1' \cap T_2' = (S^2 - T_1') \cup (S^2 - T_2') \tag{1.19}$$

is a union of two connected subspaces. Moreover, these subspaces share some common point, for the intersection

$$(S^2 - T_1') \cap (S^2 - T_2') = S^2 - J' \tag{1.20}$$

is non-empty by hypothesis. Thus, $S^2 - T_1' \cap T_2'$ is connected, allowing us to call upon Lemma 1.4.5 a final time to deduce that

$$H_0(S^2 - T_1' \cap T_2') \cong \mathbb{Z}. \tag{1.21}$$

At this point, we recall that the alternating sum of the ranks of Abelian groups forming an exact sequence vanishes. In this context, this means that

$$\operatorname{rank} H_0(S^2 - J') = 1 + 2 - 1 = 2, \tag{1.22}$$

and since the zeroth singular homology group of any space is free on the set of that space's path components, we conclude that $H_0(S^2 - J') \cong \mathbb{Z}^2$. In particular, $S^2 - J'$ is disconnected, so that (the contrapositive of our statement of) Lemma 1.4.6 asserts that $\big(S^2/{\sim}\big) - J$ is disconnected as well. □

1.4.4 Completing the Proof

We now have all of the components necessary for a proof of Moore's theorem; the only thing that remains for us to do is to assemble them.

Proof of Theorem 1.1.2. Let $\sim$ denote some equivalence relation on S^2 which satisfies the hypotheses of Theorem 1.1.2. Trivially, $S^2/{\sim}$ is both compact and connected, while Lemmas 1.4.4 and 1.4.3 respectively establish that $S^2/{\sim}$ is also locally connected and metrizable; taken together, these facts imply that the quotient $S^2/{\sim}$ constitutes a Peano continuum. Meanwhile, it is precisely the content of Lemmas 1.4.8 and 1.4.9 that $S^2/{\sim}$ satisfies the Jordan curve theorem in the sense of Theorem 1.2.1. It follows that $S^2/{\sim}$ has all of the properties required for us to call upon Zippin's sphere characterization, and we conclude that $S^2/{\sim} \cong S^2$. □

1.5 Applications

Before moving on to a study of Zippin's sphere characterization in its own right, we pause to briefly outline two appearances of Moore's theorem 'in the wild'. In Section 1.5.1, we illustrate how Moore's theorem enables a construction known as *polynomial mating*, while in Section 1.5.2, we discuss how the same result facilitates the study of dynamical systems defined on extremely complicated topological spaces.

1.5.1 Polynomial Mating

Recall the *Riemann mapping theorem*, a classical result of complex analysis.[8] (By a *continuum*, we mean a compact, connected, metrizable space. In Appendix A, these spaces are explored in greater depth.)

Theorem 1.5.1 (Riemann mapping theorem)**.** *Call a continuum $K \subset \mathbb{C}$* non-separating *if the complement $\mathbb{C} - K$ is connected, and let*

$$D^2 = \left\{ re^{i\theta} \in \mathbb{C} \,\middle|\, r \in [0,1] \text{ and } \theta \in [0, 2\pi[\right\}$$

denote the closed unit disc.

Then, given any non-separating continuum $K \subset \mathbb{C}$, there exists a conformal isomorphism

$$\varphi\colon \mathbb{C} - D^2 \to \mathbb{C} - K,$$

with the property that

$$\lim_{|z|\to\infty} \frac{\varphi(z)}{z} = 1.$$

We call φ the exterior Riemann map associated to *K.* □

In their study of the Mandelbrot set, Douady and Hubbard introduced the notion of *external rays* [DH82], which we shall need in what follows.

Definition 1.5.2. *Denote by $\varphi\colon \mathbb{C} - D^2 \to \mathbb{C} - K$ the exterior Riemann map associated to some non-separating continuum $K \subset \mathbb{C}$. For each $\theta \in [0, 2\pi[$, the* external ray of angle θ *is defined to be the curve*

$$R_\theta = \left\{ \varphi\big(re^{i\theta}\big) \,\middle|\, r \in]1, \infty[\right\} \subset \mathbb{C} - K.$$

If the limit $\lim_{r\to 1} \varphi\big(re^{i\theta}\big)$ exists for some $\theta \in [0, 2\pi[$, then this limit is known as the landing point *of the external ray R_θ, and the external ray R_θ is said to* land *(on K).*

Of particular importance is the fact that when the continuum K in Definition 1.5.2 is Peano (Definition A.2.1), all external rays land. This is a direct consequence of the work of Carathéodory [Car13] and Torhorst [Tor21], which we summarise without proof in the following theorem.

[8] Typical treatments of the Riemann mapping theorem in fact state a slightly different result to that presented here. Theorem 1.5.1 is a straightforward consequence of the 'usual' statement of the Riemann mapping theorem, of which proofs may be found in, for instance, [SS03].

Theorem 1.5.3 (Carathéodory-Torhorst)**.** *Let φ denote the exterior Riemann map associated to some non-separating continuum $K \subset \mathbb{C}$. The following are then equivalent:*

- *There exists a continuous extension $\widetilde{\varphi}\colon \mathbb{C} - \operatorname{int} D^2 \to \mathbb{C} - K$, expanding the domain of φ to include the unit circle; and*
- *The boundary $\operatorname{fr} K$ is locally connected.*

In particular, if the continuum K is Peano, then there exists such a continuous extension $\widetilde{\varphi}$ of φ. □

Lemma 1.5.4. *If $K \subset \mathbb{C}$ is a Peano continuum with connected complement, then all of the external rays $\{R_\theta\}_{\theta \in [0,2\pi[}$ land on K.*

Proof. If φ denotes the exterior Riemann map associated to K, and $\widetilde{\varphi}$ the continuous extension guaranteed by Theorem 1.5.3, then the point $\widetilde{\varphi}(e^{i\theta}) \in \operatorname{fr} K$ is readily seen to be the landing point of R_θ in K. □

Now, let us consider the implications of the combination of Lemma 1.5.4 with our knowledge that $\lim_{|z|\to\infty} \frac{\varphi(z)}{z} = 1$ for any exterior Riemann map φ. If we take any pair of non-separating Peano continua $K_1, K_2 \subset \mathbb{C}$, then we have an associated pair of exterior Riemann maps

$$\varphi_1 \colon \mathbb{C} - D^2 \to \mathbb{C} - K_1 \text{ and } \varphi_2 \colon \mathbb{C} - D^2 \to \mathbb{C} - K_2, \tag{1.23}$$

along with corresponding families $\{R_{1,\theta}\}_{\theta \in [0,2\pi[}$ and $\{R_{2,\theta}\}_{\theta \in [0,2\pi[}$, all of which land (Figures 1.3a and 1.3b). Crucially, because $\lim_{|z|\to\infty} \frac{\varphi_i(z)}{z} = 1$ for $i \in \{1, 2\}$, we know that, heuristically, the external rays $R_{1,\theta}$ and $R_{2,\theta}$ each 'look the same' sufficiently far from the origin, for every $\theta \in [0, 2\pi[$, which enables the following construction.

Define *gnomonic projections* $\nu_1, \nu_2 \colon \mathbb{C} \to \mathbb{C} \times \mathbb{R}$ by

$$\nu_1(z) = \frac{1}{\sqrt{|z|^2 + 1}}(z, 1) \text{ and } \nu_2(z) = \frac{1}{\sqrt{|z|^2 + 1}}(\overline{z}, -1) \text{ for each } z \in \mathbb{C} \tag{1.24}$$

mapping the complex plane to the open upper and lower hemispheres, respectively, of the 2-sphere embedded in $\mathbb{C} \times \mathbb{R}$.

Letting $S^1 \subset S^2$ denote the equator, we know that, for each $\theta \in [0, 2\pi[$, both $\nu_1(R_{1,\theta})$ and $\nu_2(R_{2,-\theta})$ have a boundary point $\big(e^{2\pi i\theta}, 0\big) \in S^1$, as sketched in Figure 1.3c. Thus, we can define an equivalence relation $\sim$ on S^2, the so-called *ray equivalence relation*, to be that generated by the relation which identifies $\operatorname{cl}_{S^2} \nu_1(R_{1,\theta})$ and $\operatorname{cl}_{S^2} \nu_2(R_{2,-\theta})$ for each $\theta \in [0, 2\pi[$. We introduce the notation $S^2/{\sim} = K_1 \perp\!\!\!\perp K_2$ for the quotient space which results from this construction.

Of course, there is no *a priori* reason that the space $K_1 \perp\!\!\!\perp K_2$ should be either interesting or well-behaved, and indeed, at this level of generality, there is little more to be said about $K_1 \perp\!\!\!\perp K_2$. In the context of holomorphic dynamics, however, constructions of this kind underlie the rich and expressive theory of *polynomial mating.* Recalling some concepts from holomorphic dynamics will furnish us with the vocabulary to discuss this further.

Definition 1.5.5. *Let $f \colon \mathbb{C} \to \mathbb{C}$ denote some holomorphic function, and for each $n \in \mathbb{N}$, denote by $f^{\circ n}$ its n-fold iterate. (That is, let $f^{\circ 1} = f$, $f^{\circ 2} = f \circ f$, and so on.)*

For each point $z \in \mathbb{N}$, the orbit of z (under f) *is the set*

$$O_f(z) = \{f^{\circ n}(z) \mid n \in \mathbb{N}\} \subset \mathbb{C}.$$

A point $c \in \mathbb{C}$ is said to be a critical point of f *if f has vanishing derivative at c, and the union*

$$P_f = \bigcup \{O_f(c) \mid c \text{ is a critical point of } f\} \subset \mathbb{C}$$

is the postcritical set of f*. When P_f has finite cardinality, the function f is termed* postcritically finite.

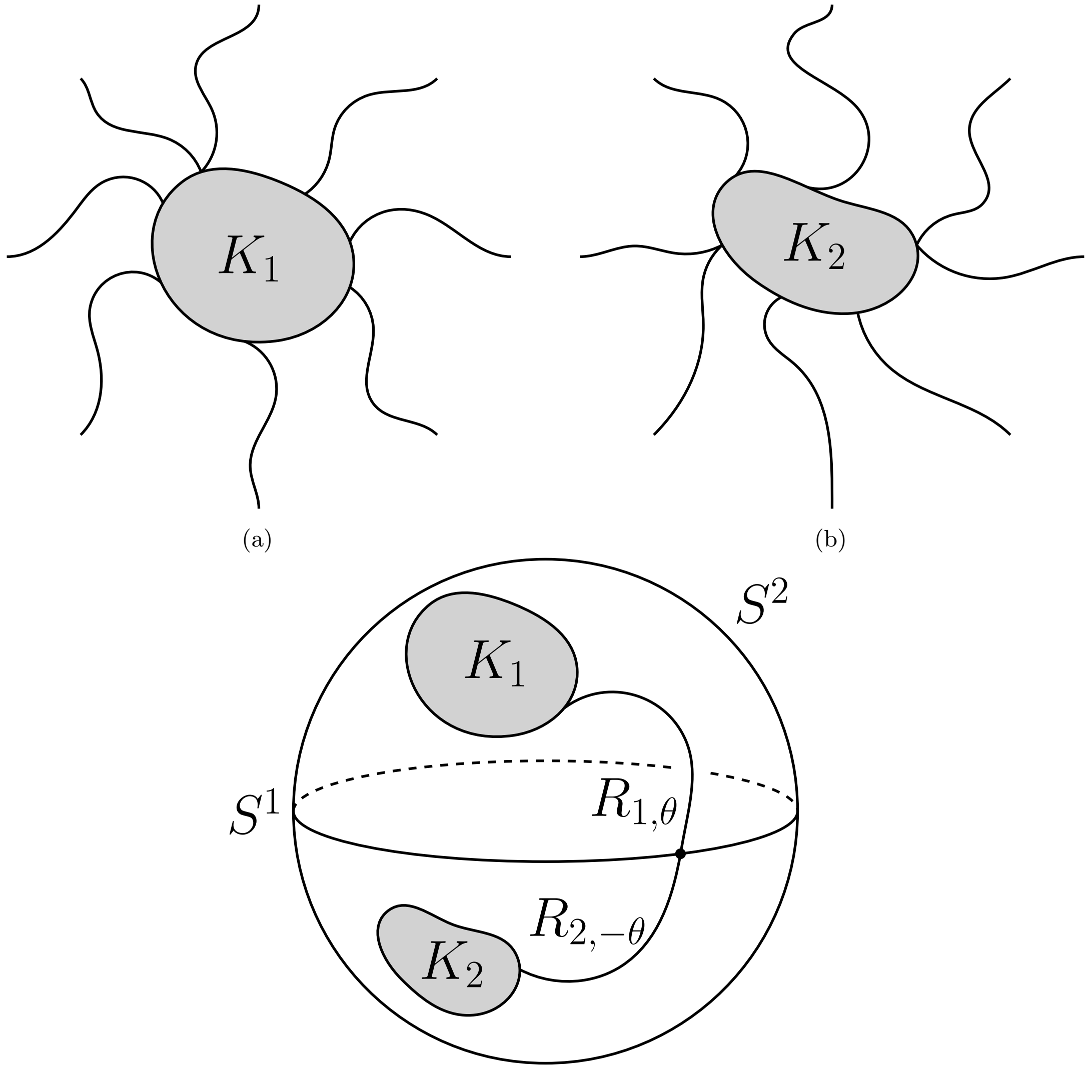

(a) (b)

(c)

If f is a polynomial, then the filled Julia set of f *is defined as the union*

$$K_f = \bigcup \{ O_f(z) \mid O_f(z) \text{ is bounded} \},$$

and we may define the Julia set of f *to be the topological boundary* $J_f = \operatorname{fr} K_f$.

Suppose, then, that $f_1, f_2 \colon \mathbb{C} \to \mathbb{C}$ are some pair of monic polynomials, both of equal degree $d \geq 2$, for which the associated filled Julia sets K_1 and K_2 are Peano continua. In terms of the gnomonic projections ν_1 and ν_2 defined in (1.24), we are able to produce from f_1 and f_2 a smooth function $f_1 \uplus f_2 \colon S^2 \to S^2$ defined by

$$(f_1 \uplus f_2)(z, h) = \begin{cases} (\nu_1 \circ f_1)(z) & \text{if } h > 0, \\ \big(z^d, 0\big) & \text{if } h = 0, \text{ and} \\ (\nu_2 \circ f_2)(z) & \text{if } h < 0, \end{cases} \tag{1.25}$$

where we continue to view S^2 as being embedded in $\mathbb{C} \times \mathbb{R}$. Observing that $f_1 \uplus f_2$ descends to the quotient induced by the ray equivalence relation $\sim$, we arrive at a continuous function $f_1 \perp\!\!\!\perp f_2 \colon K_1 \perp\!\!\!\perp K_2 \to S^2$, which is known as the *mating* of f_1 and f_2.

So far, it does not look as if we have gained much at all: just as the space $K_1 \perp\!\!\!\perp K_2$ may be highly pathological, so too may be the induced function $f_1 \perp\!\!\!\perp f_2$. However, it turns out that in many cases, this construction in fact produces a rational map defined on S^2. Following [PM13], we introduce some jargon which will allow us to discuss this further.

Definition 1.5.6. *Let f_1 and f_2 denote monic complex polynomials, both of equal and at least quadratic degree, for which the corresponding filled Julia sets K_1 and K_2 are Peano continua. If the resulting ray equivalence relation $\sim$ fails to satisfy the hypotheses of Moore's theorem, so that the space $K_1 \perp\!\!\!\perp K_2 \not\cong S^2$, then we classify the mating of f_1 and f_2 as* Moore-obstructed*.*

Suppose that the mating of f_1 and f_2 is not Moore-obstructed, so that $K_1 \perp\!\!\!\perp K_2 \cong S^2$. If the mating $f_1 \perp\!\!\!\perp f_2$ is topologically conjugate to some rational map $F\colon S^2 \to S^2$ (that is, if there exists some homeomorphism $h\colon S^2 \to S^2$ and some rational map $F\colon S^2 \to S^2$ such that $h \circ (f_1 \perp\!\!\!\perp f_2) = F \circ h$), then we say that f_1 and f_2 are topologically mateable*, and $f_1 \perp\!\!\!\perp f_2$ is termed a* topological mating*. If no such conjugacy exists, then the mating of f_1 and f_2 is considered to be* topologically obstructed*.*

If f_1 and f_2 are topologically mateable via some conjugacy h which is holomorphic on the interiors of K_1 and K_2 (when they are non-empty), then f_1 and f_2 are termed geometrically mateable*, and $f_1 \perp\!\!\!\perp f_2$ a* geometric mating*. If this is not the case, then the mating of f_1 and f_2 is referred to as* geometrically obstructed*.*

If h is in fact conformal on the interiors of K_1 and K_2 (again, when non-empty), then f_1 and f_2 are conformally mateable*, and $f_1 \perp\!\!\!\perp f_2$ is a* conformal mating*. Otherwise, the mating of f_1 and f_2 is* conformally obstructed*.*

The first result asserting the existence of matings under reasonably weak hypotheses which we present is due to Yampolsky and Zakeri [YZ00]; for explanations of definitions appearing in this theorem, we refer to [Mil06].

Theorem 1.5.7 (Yampolsky-Zakeri)**.** *Denote by f_1 and f_2 some pair of quadratic complex polynomials which are not anti-holomorphically conjugate to one another, each with a Siegel fixed point of bounded type. There then exists a geometric mating $f_1 \perp\!\!\!\perp f_2$.* □

If we restrict ourselves to the postcritically finite setting, then it is possible to say more about matings, and we give just three examples. The first two are due to Meyer [Mey09, Mey14] and describe classes of rational maps which do or do not arise as matings; the other, thanks to Tan, Rees and Shishikura [Tan92, Ree92, Shi00], asserts the unique existence of certain conformal matings[9].

Theorem 1.5.8 (Meyer)**.** *Suppose that $F\colon S^2 \to S^2$ denotes some postcritically finite rational map, with all of S^2 as its Julia set.*[10] *Then, there exists some $N \in \mathbb{N}$ such that, for every $n \geq N$, the n-fold iterate $F^{\circ n}$ is topologically conjugate to the topological mating of two polynomials.* □

Theorem 1.5.9 (Meyer)**.** *Let $F\colon S^2 \to S^2$ denote some rational map which is not a polynomial, and suppose that the postcritical set of F contains precisely three points. Then, F is not a mating of any pair of polynomials.*

Theorem 1.5.10 (Tan-Rees-Shishikura)**.** *Suppose that the complex polynomials $f_1(z) = z^2 + c_1$ and $f_2(z) = z^2 + c_2$ are postcritically finite. If c_1 and c_2 do not lie in conjugate limbs of the Mandelbrot set, then f_1 and f_2 are conformally mateable, and the result of this mating is unique up to Möbius conjugacy.* □

The uniqueness portion of Theorem 1.5.10 is particularly striking, since extant matings in general fail to be unique. For instance, in the quartic setting, there can be found a topological mating which is conjugate to uncountably many distinct geometric matings [Mil04].

The theory of mating provides a means of constructing new rational maps from appropriate polynomials in a controlled way, and dually, serves as a mechanism for the study of the dynamics of appropriate rational maps by realizing them as matings of polynomials [PM13]. However, many questions remain open in the field. For instance, there is no general theory of the (non-)existence of

[9] The result quoted in Theorem 1.5.10 is in fact weaker than that proved in [Shi00]. There, a more restrictive definition of conformal mating is used, and using this definition, the conditions of Theorem 1.5.10 are not just sufficient, but also necessary.

[10] Although we gave one definition of a Julia set in Definition 1.5.5, we restricted ourselves to the polynomial setting for convenience's sake. More general definitions of Julia sets, appropriate for rational maps, can be found in, for instance, [Mil06].

obstructions to mating, be they Moore, topological, geometric, or conformal.[11] From the perspective of decomposing rational maps, the phenomenon of *shared matings*, whereby particular rational maps may be described as the matings of multiple distinct pairs of polynomials, is well-known [Wit88]; however, no systematic means of enumerating how a given rational map may be realized as a mating is currently available.

1.5.2 Inverse Limits of Tent Maps

Unimodal maps (naïvely, self-maps of some closed interval with a single well defined 'turning point') are of substantial interest in the field of one-dimensional dynamics, with the *kneading theory* of Milnor and Thurston [MT88] enabling their detailed study. An important property of any unimodal map is that its domain of definition can be restricted to some invariant subinterval–its so-called *core*–which contains all non-trivial dynamics.

Before continuing, we make these definitions precise.

Definition 1.5.11. *Let $f\colon [a,b] \to [a,b]$ denote some continuous self-map of a closed interval such that*

- *There exists a point $c \in [a,b]$, the so-called* turning point of f*, such that $f|_{[a,c]}$ is strictly increasing, and $f|_{[c,b]}$ strictly decreasing; and*
- *$x < f(x)$ for each $x \in \,]a,c]$.*

If $f(a) = f(b) = a$, then f is said to be non-core unimodal*; if, instead, $f(c) = b$ and $f(b) = a$, then f is termed* core unimodal.

Given some non-core unimodal map $f\colon [a,b] \to [a,b]$ with turning point c, the core *of f is defined as the interval $C = \left[f^{\circ 2}(c), f(c)\right]$, so that the restriction $f|_C$ is core unimodal.*

Amongst the simplest examples of unimodal maps we find the piecewise linear *tent maps*, which can shed light on the behaviour of more general unimodal maps, in spite of the simplicity of their definition. No small degree of this theoretical significance derives from the (rough) fact that, if f is some unimodal map which 'moves close-together points apart after sufficiently many iterations', then the dynamics of some tent map are, in some sense, contained within those of f. Theorem 1.5.13 [Par66, MT88] couches this in more exact language.

Definition 1.5.12. *Given any $t \in \,]1,2]$, the unimodal map $g_t\colon [0,1] \to [0,1]$ defined by*

$$g_t(x) = t \min\{x, 1-x\}$$

is called the tent map of slope t*. Collectively, the set $\{g_t\}_{t\in]0,2]}$ of all such maps is referred to as the* tent family.

Theorem 1.5.13 (Parry; Milnor-Thurston)**.** *Let $f\colon [a,b] \to [a,b]$ denote some unimodal map with strictly positive* topological entropy[12] $\log s$*. Then, there exists a semiconjugacy $p\colon [a,b] \to [0,1]$ from f to the tent map g_s of slope s. (That is, there exists a continuous surjection $p\colon [a,b] \to [0,1]$ such that $p \circ f = g_s \circ p$.)*

In all cases, the semiconjugacy p can be described by an explicit formula. □

The unimodal maps of Definition 1.5.11 are not homeomorphisms, but there is a well established technique for producing from a non-invertible dynamical system an invertible one, defined on the so-called *inverse limit space*. The book [IM10] discusses the applications of inverse limit spaces to topological dynamics in detail.

[11] It is conjectured that no matings are purely geometrically-obstructed, so that every every topologically mateable pair of polynomials is geometrically mateable, but currently neither a proof nor a counterexample is known.

[12] Heuristically, we can interpret the topological entropy of such a map as a dynamical invariant measuring the degree of divergence of initially close-together points under repeated iterations. We will not need any particular knowledge of the details of topological entropy in what follows; more in-depth treatment of the concept can be found in, for instance, [KH95].

Definition and Lemma 1.5.14. *Let X denote some compact metrizable space, equipped with some continuous surjection $f\colon X \to X$. The* inverse limit of (X, f) *is the subspace*

$$\widehat{X} = \varprojlim(X, f) = \{(x_1, x_2, x_3, \ldots) \mid x_i = f(x_{i+1}) \text{ for each } n \in \mathbb{N}\} \subset X^{\mathbb{N}},$$

where we have endowed $X^{\mathbb{N}}$ with the product topology.

The function $\widehat{f}\colon \widehat{X} \to \widehat{X}$ defined by

$$\widehat{f}(x_1, x_2, x_3, \ldots) = (f(x_1), x_1, x_2, x_3, \ldots)$$

is a homeomorphism, termed the natural extension of f to $\widehat{X}$.

The projection $\pi\colon \widehat{X} \to X$ defined by

$$\pi(x_1, x_2, x_3, \ldots) = x_1$$

is a semiconjugacy from $\widehat{f}$ to f; that is, $\pi \circ \widehat{f} = f \circ \pi$. More generally, any semiconjugacy from any invertible dynamical system to f factors through π. □

A reasonable first step in studying a unimodal dynamical system $([a, b], f)$, then, is to form the inverse limit $\widehat{I} = \varprojlim([a, b], f)$, and then study the dynamics of the invertible natural extension $\widehat{f}$. However, there is some conservation of difficulty in play here: although the homeomorphism $\widehat{f}$ may be better-behaved than our original map f, the topology of inverse limit spaces can, in general, be exceptionally intricate. One canonical example of this complexity is Henderson's construction [Hen64] of the *pseudoarc*[13] as the inverse limit of a particular self-map of $[0, 1]$.

In the context of core tent maps, results of Barge, Brucks and Diamond [BBD96] illustrate just how baroque the topology of the resulting inverse limit spaces can be.

Theorem 1.5.15 (Barge-Brucks-Diamond)**.** *For each $t \in [\sqrt{2}, 2]$, let $\widehat{I}_t$ denote the inverse limit of the tent map $g_t\colon [0, 1] \to [0, 1]$.*

There exists a dense subspace $A \subset [\sqrt{2}, 2]$ such that $\widehat{I}_t$ is, except at finitely many points, locally the product of an arc and a Cantor set, for each $t \in A$.

There also exists a dense full-measure G_δ subspace $B \subset [\sqrt{2}, 2]$ such that, if $s \in B$, then every open subspace of $\widehat{I}_s$ contains a homeomorph of $\widehat{I}_t$ for every $t \in [\sqrt{2}, 2]$. □

Adding more layers of detail still, the spaces $\widehat{I}_s$ and $\widehat{I}_t$, in the notation of Theorem 1.5.15, are conjecturally non-homeomorphic when $s \neq t$. This so-called *Ingram conjecture* was answered in the affirmative for non-core tent maps by Barge, Bruin and Štimac [BBŠ12]; for general core tent maps, it is unknown whether or not the Ingram conjecture holds in full generality, although it has been shown by Anušić, Bruin and Činč [ABČ15] that there are uncountably many homeomorphism classes.[14]

In light of the convoluted topologies of the inverse limit spaces of unimodal maps, it is perhaps surprising that dynamics in such a space are tightly entwined with dynamics on the 2-sphere. The proof of this fact, due to Boyland, de Carvalho and Hall [BdCH17] is highly non-trivial, and shall not be dwelled upon in any great detail; we content ourselves with remarking that the passage from an inverse limit space to the 2-sphere is enabled by Moore's theorem.

Theorem 1.5.16 (Boyland-de Carvalho-Hall)**.** *Let $\{f_t\}_{t \in J}\colon [a, b] \to [a, b]$ denote some family of core unimodal maps, all of which satisfy the conditions of Convention 2.8 in [BdCH17]. where the parameter interval $J \subset \mathbb{R}$ is compact. For each $t \in J$, let $\widehat{I}_t$ denote the inverse limit $\varprojlim([a, b], f_t)$, and $\widehat{f}_t$ the corresponding natural extension.*

[13] A non-degenerate continuum P with the striking property of *hereditary indecomposability*: it is impossible to express P, or any subcontinuum of P, as a union of two proper subcontinua of P.

[14] More precisely, it has been shown that the Ingram conjecture holds for core tent maps when their turning point fails to be *preperiodic*, so that $g_t^{\circ n}(\frac{1}{2}) \neq g_t^{\circ m}(\frac{1}{2})$ when $n \neq m$, and is moreover *non-recurrent*, so that there exists some neighbourhood of $\frac{1}{2}$ containing no iterates of $\frac{1}{2}$ under g_s.

There then exists a continuously varying family $\{\chi_t\}_{t\in J}$ *of self-homeomorphisms of the* 2*-sphere* S^2*, such that there exists, for each* $t \in J$*, a semiconjugacy* $\pi_t \colon \widehat{I_t} \to S^2$ *from* $\widehat{f_t}$ *to* χ_t*.*

Moreover, for each such semiconjugacy π_t*, all but one fibre contains three or fewer points, and only countably many fibres contain precisely three points.* □

Chapter 2

Zippin's Sphere Characterization

2.1 Overview

This chapter is dedicated entirely to a proof of Zippin's sphere characterization (Theorem 1.2.1), which enabled our proof of Moore's theorem in Chapter 1, and our approach shall mirror that of van Kampen in [vK35]. For convenience, we restate Zippin's sphere characterisation below.

Theorem 1.2.1 (Zippin's sphere characterization). *Let X denote some Peano continuum (Definition A.2.1). If X satisfies the Jordan curve theorem in the sense that*

- *X contains a Jordan curve,*
- *Every Jordan curve in X separates X, and*
- *No closed arc in X separates X,*

then X is homeomorphic to S^2. □

Quite aside from its utility in Chapter 1, this result is interesting in its own right: given the intuitively plausible condition of 'satisfying the Jordan curve theorem', we are able to distinguish S^2 among the class of Peano continua, which can in general be extremely pathological.

Before continuing, we would also remark on the similarities between Zippin's sphere characterization and the characterization of S^1 given in Lemma A.4.6. There, the space S^1 is identified as the unique continuum separated by every pair of distinct points; that is, as the unique continuum separated by each embedded copy of the 0-sphere. Meanwhile, Theorem 1.2.1 characterizes S^2 in terms of separation by Jordan curves, which is to say by embedded copies of the 1-sphere, although also requires additional hypotheses to do so.

The similarity between Theorem 1.2.1 and Lemma A.4.6 is further reflected in our methods of proof: in each case, we take a copy of S^0 or S^1 in our space X, and show that its complementary domains in X are (the interiors of) a pair of 1- or 2-cells with the original 0- or 1-sphere as their common boundary. One may well wonder whether a similar construction is possible in higher dimensions, but even for the 3-sphere new techniques are required. Part of the reason for this is the existence in S^3 of pathological objects such as *Alexander's horned sphere* [Ale24], an embedding of S^2 with a complementary domain which fails to be simply connected; a modern treatment can be found in [Hat02].

2.2 Strategy of Proof

To streamline our discussion slightly, we introduce the following terminology. (By a *generalized Peano continuum* is meant a connected, locally connected, locally compact, metrizable space. See

Definition A.2.1, and Appendix A in general.)

Definition 2.2.1. *A generalized Peano continuum X is said to be a* Zippin space *if*

- *X contains a Jordan curve,*
- *Every Jordan curve in X separates X, and*
- *No closed arc in X separates X.*

Now, as a first step towards proving Theorem 1.2.1, we recall a more straightforward characterization of S^2. Namely, we realize the sphere as a CW complex consisting of a pair of 2-cells identified along their boundaries–or, perhaps more evocatively, as two hemispheres glued together. Rather than restricting ourselves to the two-dimensional case, we prove that we can in fact describe the n-sphere in this way for all $n \in \mathbb{N}$; this generalization adds no complexity to the proof, and will be of use in Section A.4.1.

Lemma 2.2.2. *Given some $n \in \mathbb{N}$, explicitly identify the n-cell D^n and the $(n-1)$-sphere S^{n-1} with the subspaces*

$$D^n = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \leq 1\} \text{ and } S^{n-1} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| = 1\},$$

of $\mathbb{R}^n$, and denote by $\iota\colon S^{n-1} \to D^n$ the set-theoretic inclusion. The adjunction space $D^n \cup_\iota D^n$ is homeomorphic to S^n.

(Explicitly, define the disjoint union $D^n \sqcup D^n$ to be the topological product $D^n \times \{1,2\}$, where $\{1,2\}$ is understood to be equipped with the discrete topology. Further, let $\sim'$ denote the relation on the disjoint union $D^n \sqcup D^n$ defined by $(x,i) \sim' (y,j)$ if and only if $x = \iota(y)$, and $\sim$ the equivalence relation on $D^n \sqcup D^n$ generated by $\sim'$. Then, the quotient $D^n \cup_\iota D^n = (D^n \sqcup D^n)/{\sim}$ is homeomorphic to S^n.)

Proof. As suggested in this section's opening paragraph, our strategy shall be to map each copy of the n-cell in $D^n \sqcup D^n$ to a hemisphere of S^n; concretely, we define a surjection $f\colon D^n \sqcup D^n \to S^n$ by

$$f(\mathbf{x}, i) = \begin{cases} \left(x_1, x_2, \ldots, x_n, \sqrt{1 - \sum_{k=1}^n x_k^2}\right) & \text{if } i = 1 \text{ and} \\ \left(x_1, x_2, \ldots, x_n, -\sqrt{1 - \sum_{k=1}^n x_k^2}\right) & \text{if } i = 2, \end{cases} \tag{2.1}$$

where D^n and S^n are defined as in the statement of the lemma.

We immediately notice that this function is closed, as a continuous function with compact domain and Hausdorff codomain, whence it follows that f is a quotient map.[1] Observing additionally that the equivalence kernel[2] of f is precisely $\sim$ allows us to conclude that $D^n \cup_\iota D^n \cong S^n$: any continuous function constant on the equivalence classes of $\sim$ must factor through f, and this is precisely the universal property of the quotient $D^n \cup_\iota D^n = (D^n \sqcup D^n)/{\sim}$. □

Our proof of Zippin's sphere characterization shall ultimately reduce to an application of the above lemma. More concretely, suppose that we have some compact Zippin space X, alongside some arbitrary Jordan curve $J \subset X$. If we can show that

- Each complementary domain of J in X has precisely J as its boundary;
- There are precisely two complementary domains of J in X, say U and V;
- Both $U \cup J$ and $V \cup J$ are homeomorphs of the 2-cell; and
- It is possible to attach $U \cup J$ to $V \cup J$ along J in an appropriate fashion;

[1] We recall that, given topological spaces X and Y, a continuous surjection $f\colon X \to Y$ is termed a *quotient map* if the topology of Y is final with respect to f. That is, the function f is a quotient map if the topology of Y is generated from a subbasis consisting of those subsets $A \subset Y$ for which $f^{-1}(A) \subset X$ is open. This is discussed in any textbook treating elementary topology, such as [Lee00].

[2] That is, the equivalence relation on X which identifies points $x, y \in X$ if and only if $f(x) = f(y)$.

then Lemma 2.2.2 will yield a homeomorphism $X \cong S^2$. Before we embark upon this undertaking, however, it is worth taking some time to consider how we might proceed.

Proving the first two points above amount to a proof of the Jordan curve theorem for compact Zippin spaces. Modern proofs of the Jordan curve theorem in $\mathbb{R}^2$ tend to involve results from algebraic topology, often broadly similar to that used during our proof of Moore's theorem in Chapter 1, but no direct translation of these techniques to the current setting presents itself: we do not have enough information about the homology or homotopy groups of arbitrary Peano continua. It therefore is necessary for us to approach the problem using more elementary technology, and our method shall be essentially that of van Kampen [vK35] and Wilder [Wil49], using results from continuum theory which can be traced back to Whyburn in [Why28, Why31].

As an aside, it is interesting and perhaps surprising to note that this proof will essentially frame the Jordan curve theorem as a consequence of the non-planarity of the complete bipartite graph $K_{3,3}$ (Figure 2.1), as we shall see with Lemma 2.3.6. (This relationship is at the core of a proof [Tho92] of the Jordan curve theorem in the plane due to Thomassen.)

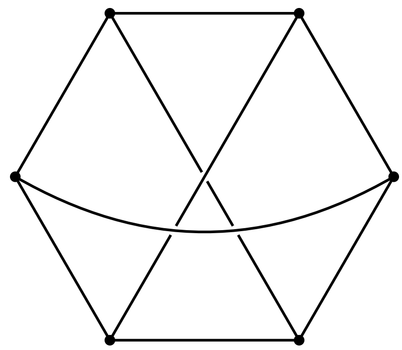

Figure 2.1: The complete bipartite graph $K_{3,3}$. In this depiction, each of the top-left, bottom-left and middle-right vertices is connected to each of the top-right, bottom-right and center-left vertices.

The third point, meanwhile, is a clear analogue of Schoenflies' theorem. In a similar vein to our above comments, proofs of Schoenflies' theorem in the plane often appeal to structure which we do not here have at our disposal–this time, to complex-analytic techniques, with the Riemann mapping theorem and the extension theorem of Carathéodory-Torhorst chief among them. It is again van Kampen [vK35] who supplies us with a way around the issue, making use of the completeness of Peano continua as metric spaces. Roughly speaking, an inductive construction shall equip us with a homeomorphism between a dense subspace of the 2-cell and a dense subspace of (the closure of) a complementary domain of a Jordan curve in a compact Zippin space; then, completeness allows us to extend this to a homeomorphism between the entirety of the spaces.

The remainder of this chapter is given over to addressing each of the four bullet points above. As alluded to, results from the theory of Peano continua–and especially from Whyburn's theory of cyclic connectivity [Why31]–will be crucial throughout this process, but the development of such machinery at this stage would represent a lengthy digression from the task at hand. To streamline our exposition, we therefore defer our study of Peano continua *per se* to Appendix A, freely calling upon results as the need arises.

2.3 The Jordan Curve Theorem

In this section, we wish to show that if J denotes some Jordan curve in a Zippin space X, then the complement $X - J$ has precisely two components, each with J as its boundary. To compartmentalise our discussion, we deal exclusively with the latter condition in Section 2.3.1, leaving the task of enumerating the complementary domains of J in X until Section 2.3.2.

2.3.1 Boundaries of Complementary Domains

We start with the technically simpler portion of the Jordan curve theorem for Zippin spaces: namely, the condition for boundaries. Thus, given a Zippin space X and a Jordan curve $J \subset X$, we wish to show that if A is some component of $X - J$, then $\operatorname{fr}_X A = J$.

The inclusion $\operatorname{fr}_X A \subset J$ in fact holds in any locally connected space, and is a corollary of the following.

Lemma 2.3.1. *Let X denote some locally connected topological space, and consider some subspace $Y \subset X$. If A denotes some connected component of Y, then $\operatorname{fr}_X A \subset \operatorname{fr}_X Y$.*

Proof. We first propose that $(\operatorname{cl}_X A) \cap (\operatorname{int}_X Y) \subset \operatorname{int}_X A$. Heuristically, one could view this as stating that the connected component A is, in some sense, 'as large as it can be': if a point of the closure of A is interior to Y, then that point must in fact be interior to A as well.

Indeed, consider some point $x \in (\operatorname{cl}_X A) \cap (\operatorname{int}_X Y)$. Using local connectedness, we can find some connected neighbourhood $U \subset \operatorname{int}_X Y$ of x, and since we chose A to be a connected component of Y, it follows that $U \subset A$. We find that $x \in \operatorname{int}_X A$, as suggested.

From here, we can reach the desired result by a series of set-theoretic manipulations. Explicitly, we notice that

$$\begin{aligned} \operatorname{fr}_X A &= \operatorname{cl}_X A - \operatorname{int}_X A \\ &\subset \operatorname{cl}_X A - \operatorname{int}_X Y \\ &\subset \operatorname{cl}_X Y - \operatorname{int}_X Y \\ &= \operatorname{fr}_X Y, \end{aligned} \tag{2.2}$$

where we use our above observation to pass from the first line to the second. □

Corollary 2.3.2. *Let X denote some locally connected topological space, and consider some closed subspace $C \subset X$. If A denotes some connected component of $X - C$, then $\operatorname{fr}_X A \subset C$.*

Proof. We know from the preceding lemma that $\operatorname{fr}_X A \subset \operatorname{fr}_X(X - C)$, and of course $\operatorname{fr}_X(X - C)$ is precisely $\operatorname{fr}_X C$. Since we have assumed C to be closed in X, we know that $\operatorname{fr}_X C \subset C$, and we are done. □

To prove the inclusion $J \subset \operatorname{fr}_X A$, we require an intuitively plausible result: if we are given some path which starts in an open subspace $U \subset X$, and which ends in that subspace's complement, then that path must meet the boundary $\operatorname{fr}_X U$.

Lemma 2.3.3. *Let Y denote some subspace of a topological space X. If $T \subset X$ denotes some path with one end point in Y and the other in $X - \operatorname{cl} Y$, then $T \cap \operatorname{fr} Y \neq \emptyset$.*

Proof. We start by proving the result for the special case in which Y is open in X, and we then use this to arrive at the full result.

Suppose, then, that Y is open in X. If one end point of T lies on $\operatorname{fr} Y$, then there is nothing for us to prove, so let us assume that one end point of T lies in Y and the other in $X - \operatorname{cl} Y$. Now, for the sake of a contradiction, suppose that $T \cap \operatorname{fr} Y = \emptyset$.

Using the openness of Y to observe that

$$(X - \operatorname{cl} Y) \cup Y = X - (\operatorname{cl} Y - Y) = X - \operatorname{fr} Y, \tag{2.3}$$

we find that

$$T = (T \cap Y) \cup (T \cap (X - \operatorname{cl} Y)), \tag{2.4}$$

but this is a contradiction. We know that both $T \cap Y$ and $T \cap (X - \operatorname{cl} Y)$ are disjoint non-empty open subspaces of T, but T is connected by definition. This establishes the result when Y is open.

We are now ready to approach the general case. Given some arbitrary subspace $Y \subset X$, consider some path $T \subset X$ with one end point in Y and the other in $X - Y$. Knowing that $Y = \operatorname{int} Y \cup \operatorname{fr} Y$, the result follows immediately, for the end point of T lying in $X - Y$ lies either in $\operatorname{fr} Y$, in which case we are done, or in the open subspace $X - \operatorname{cl} Y$, in which case the result follows from the special case proved above. □

Corollary 2.3.4. *Let U denote some connected open subspace of a path connected topological space X. The boundary $\operatorname{fr} U$ separates U and $X - \operatorname{cl} U$ in X.*

Proof. Were $\operatorname{fr} U$ to fail to separate U from $X - \operatorname{cl} U$, then we would be able to find some path $T \subset X - \operatorname{fr} U$ connecting a point of U to a point of $X - \operatorname{cl} U$–but the preceding lemma tells us that such a path cannot exist. □

All that remains for us to do here is to put everything together.

Lemma 2.3.5. *If J denotes some Jordan curve in a Zippin space X, and A some complementary domain of J in X, then $\operatorname{fr}_X A = J$.*

Proof. The inclusion $\operatorname{fr}_X A \subset J$ was proved in Corollary 2.3.2. To establish the reverse inclusion, we argue by contradiction, supposing that there exists some point $z \in J - \operatorname{fr}_X A$.

Being a Jordan curve, J is of course locally connected, and so we are able to find some connected open neighbourhood $U \subset J - \operatorname{fr}_X A$ of z in J; accordingly, the complement $J - U$ is a closed arc in J which contains $\operatorname{fr}_X A$. Now, select some point $x \in A$, and some point $y \in X$ separated from x by J, noticing that our assumption that X is Zippin guarantees the existence of such a point. This same assumption precludes $J - U$ from separating x and y, but Corollary 2.3.4 states that $\operatorname{fr}_X A$ separates x and y. (Notice that here we have made use of the fact that, since X is locally connected, the complementary domains of J in X are all open.)

This is the contradiction we seek: if $J - U$ does not separate x and y, then certainly no subspace of $J - U$ can separate those points either. □

2.3.2 Counting Complementary Domains

With one part of the Jordan curve theorem proven, it remains for us to show that a Jordan curve in a Zippin space has precisely two complementary domains. We mentioned in Section 2.2 that this result is closely related to the failure of the graph $K_{3,3}$ to be planar, and before addressing this section's proofs in earnest, it is worth thinking about why this might be the case.

The Euclidean space $\mathbb{R}^3$ clearly fails to be Zippin, since, for instance, the unit circle fails to separate this space. However, it does not take too much work to see how we might embed $K_{3,3}$ into $\mathbb{R}^3$: Figure 2.1 depicts one such embedding, if we interpret the edges interior to the hexagon as passing over and under one another appropriately. Similarly, we of course cannot embed $K_{3,3}$ into $\mathbb{R}^1$.

Heuristically, then, we can view the impossibility of embedding $K_{3,3}$ into a space as indicating that said space is, in some sense, 'two dimensional'–but this is not the only thing which can be deduced. There certainly exist 2-manifolds into which $K_{3,3}$ can be embedded without issue, with the torus and the Möbius band representing familiar examples; instances of such embeddings are sketched in Figure 2.2. Reflecting this, there exist Jordan curves which fail to separate both the torus and the Möbius band.

These observations supply us with another piece of intuition for why the 'non-embeddability' of $K_{3,3}$ into Zippin spaces is significant. We are able to embed $K_{3,3}$ into the torus and the Möbius band essentially because these spaces have a non-trivial fundamental group; thus, the inability to embed $K_{3,3}$ into a space can be interpreted as, in some vague sense, enforcing the triviality of that space's fundamental group.

In the context of compact 2-manifolds, these comments on how $K_{3,3}$ embeds into the torus and the Möbius band in fact lead to a further observation. A classical theorem, of which a proof may be found in [Lee00], classifies every compact 2-manifold as homeomorphic to either S^2, the connected sum of finitely many tori, or the connected sum of finitely many copies of the real projective plane $\mathbb{P}^2$. Knowing that $\mathbb{P}^2$ can be realized as a quotient of the Möbius band in which the horizontal edges

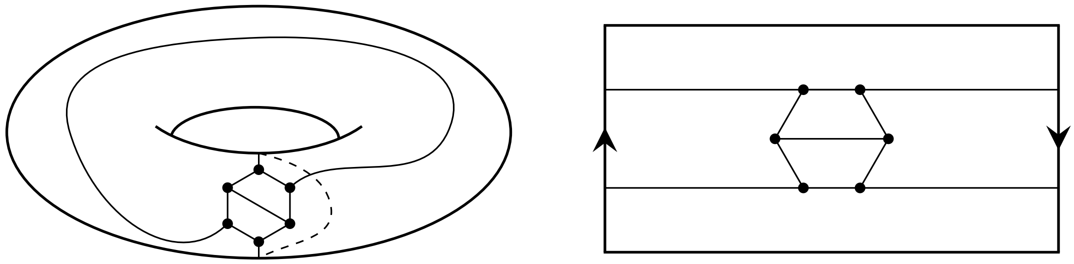

(a) (b)

Figure 2.2: Embeddings of $K_{3,3}$ into (a) the torus and (b) the Möbius band. For clarity of illustration, the Möbius band is here presented as a quotient of a rectangle in the plane, where the two vertical edges are identified with reversed orientations.

in Figure 2.2b are identified with opposing orientations, we can observe that $K_{3,3}$ can be embedded in a compact 2-manifold if and only if that manifold is not homeomorphic to S^2.

Having gained some idea why $K_{3,3}$ should be relevant to Zippin's sphere characterization, we prove what the above discussion took for granted–namely, that it is indeed impossible to embed $K_{3,3}$ into a Zippin space. We henceforth dispense with any graph-theoretic terminology, making it simpler to re-use this result in the sequel at the cost of a more complicated statement. We recall that an arc T in some space X is said to *span* a compact subspace $C \subset X$ if the intersection $T \cap C$ consists precisely of the end points of T.

Lemma 2.3.6. *Let J denote some Jordan curve in a Zippin space X, and suppose that T_1, T_2 and T are three arcs spanning J such that*

- *$]T_1[$ and $]T_2[$ lie in different components of $X - J$, and*
- *$T \cap T_1 = T \cap T_2 = \emptyset$.*

Then, the end points of T cannot be separated in J both by the end points of T_1 and by those of T_2.

Proof. We shall argue by contradiction. To this end, consider a Jordan curve $J \subset X$ together with arcs T_1, T_2 and T_3 spanning J in such a configuration that

- $]T_1[$ and $]T_2[$ lie in different components of $X - J$,
- $T \cap T_1 = T \cap T_2 = \emptyset$, and
- The end points of T are separated in J both by the end points of T_1 and by those of T_2,

as sketched in Figure 2.3. Given this arrangement of arcs, we shall construct a Jordan curve which fails to separate X.

Let x and y denote the end points of the arc T. If U_x and U_y denote the complementary domains of T_1 in J which contain x and y respectively, and if V_x and V_y denote the complementary domains of T_2 in J which contain x and y respectively, then the intersections $W_x = U_x \cap V_x$ and $W_y = U_y \cap V_y$ form a pair of disjoint open arcs in J.

We define a subspace $I = W_x \cup W_y \cup T$, sketched in Figure 2.4a, and propose that I is the complement in $J \cup T_1 \cup T_2 \cup T$ of some Jordan curve $K \subset J \cup T_1 \cup T_2 \cup T$. Indeed, we can arrive at such a curve by considering the union $T_1 \cup T_2$. If the end points of T_1 and T_2 coincide, then $T_1 \cup T_2$ is precisely the Jordan curve K; otherwise, we can form the union of $T_1 \cup T_2$ with arcs in J connecting the end points of T_1 and T_2, containing neither x nor y. Such a Jordan curve K is illustrated in Figure 2.4b.

From here, we shall show that every complementary domain of K in X intersects I. Since I is connected, it will follow that $X - K$ is also connected, which will complete the proof. With this goal in mind, let $A \subset X - K$ denote some such complementary domain.

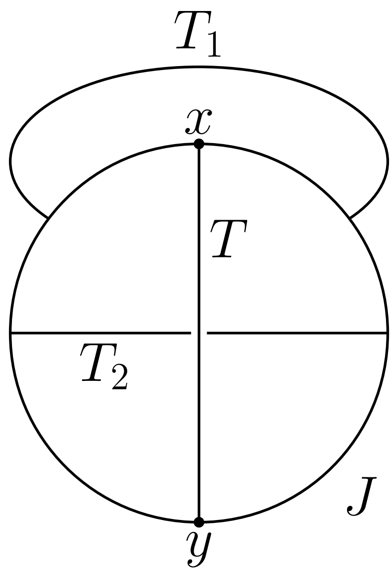

Figure 2.3

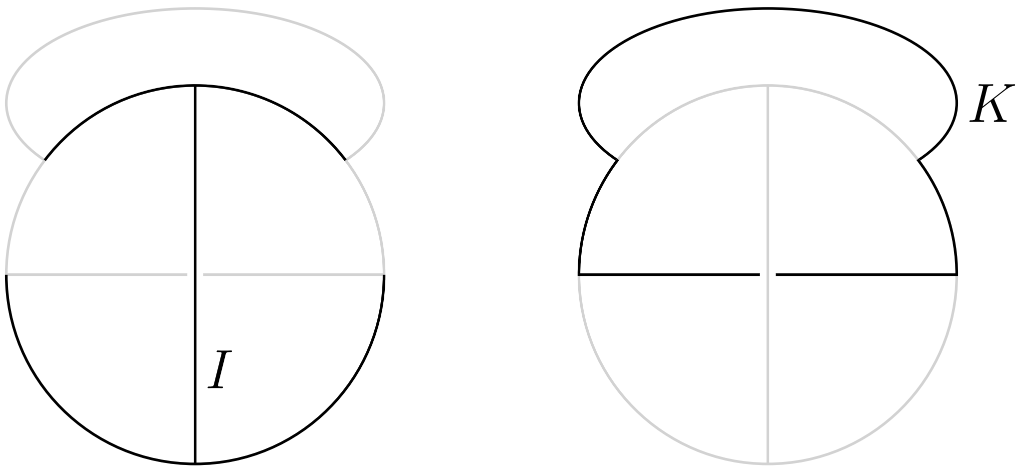

Figure 2.4

Let D_1 and D_2 denote the components of $X - J$ which contain $]T_1[$ and $]T_2[$ respectively, recalling that $D_1 \neq D_2$ by hypothesis. Given any point $z \in]T_1[$, the local connectedness of X allows us to find some neighbourhood $U \subset D_1$ of z. From Lemma 2.3.5, we know that $\mathrm{fr}\, A = K$, so that $z \in T_1 \subset \mathrm{fr}\, A$; thus, we deduce that $U \cap A \neq \emptyset$, and in particular that $D_1 \cap A \neq \emptyset$. By entirely analogous reasoning in which T_2 takes the place of T_1, we find that $D_2 \cap A \neq \emptyset$.

Now, select points $a \in D_1 \cap A$ and $b \in D_2 \cap A$. We know that A must be arc connected, because it is a connected open subspace of a generalized Peano continuum (Lemma A.5.4), allowing us to find some arc $[ab] \subset A$. Another application of Lemma 2.3.5 asserts that $\mathrm{fr}\, D_1 = J$, which combined with Lemma 2.3.3 tells us that $[ab] \cap J \neq \emptyset$. However, since $[ab] \subset X - K$ and since $(J \cup T_1 \cup T_2 \cup T) - K = I$, it in fact follows that $[ab] \cap I \neq \emptyset$. This validates our claim that $A \cap I \neq \emptyset$, and we are done: the Jordan curve K cannot separate X. □

We are now ready to prove that Zippin spaces obey a Jordan curve theorem in an appropriate sense. Given that our strategy of proof relies crucially on the notion of *arc accessibility*, discussed in Section A.5.2, we recall the relevant definition and result before tackling the proof in earnest.

Definition and Lemma A.5.8. *Let X denote some topological space with a subspace $Y \subset X$, and select some point $x \in X - Y$. We say that x is* arc accessible from Y *if there exists some arc $T \subset Y \cup \{x\}$ with x as an end point.*

Suppose that X is some Peano continuum, and $U \subsetneq X$ some open subspace. The set of points of $\operatorname{fr} U$ *which are arc accessible from U is dense in* $\operatorname{fr} U$.

Lemma 2.3.7 (Jordan curve theorem). *If X denotes some Zippin space and $J \subset X$ some Jordan curve, then the complement $X - J$ consists of precisely two connected components, each with boundary J.*

Proof. We proved the condition involving boundaries of complementary domains in Lemma 2.3.5, and we shall use Lemma 2.3.6 to prove the remaining part of the result. We know by hypothesis that J has at least two complementary domains in X, so suppose towards a contradiction that there are at least three such complementary domains, which we denote by D_1, D_2 and D_3.

Knowing that the set of points of J which are arc-accessible from D_1 is dense in J (Lemma A.5.8), we can find some arc $T_1 \subset \operatorname{cl} D_1$ spanning J (Figure 2.5a). Letting $J - T_1 = A_1 \cup A_2$ for a pair of disjoint open arcs $A_1, A_2 \subset J$, we can similarly find points $a_1 \in A_1$ and $a_2 \in A_2$ which serve as the end points of some J-spanning arc $T_2 \subset \operatorname{cl} D_2$ (Figure 2.5b). Continuing with this line of reasoning, let $A_1 - \{a_1\} = B_1 \cup B_2$ and $A_2 - \{a_2\} = B_3 \cup B_4$, for disjoint open arcs $B_1, B_2 \subset A_1$ and $B_3, B_4 \subset A_2$, and select points $b_1 \in B_1$ and $b_3 \in B_3$ which are the end points of some J-spanning arc $T_3 \subset \operatorname{cl} D_3$.

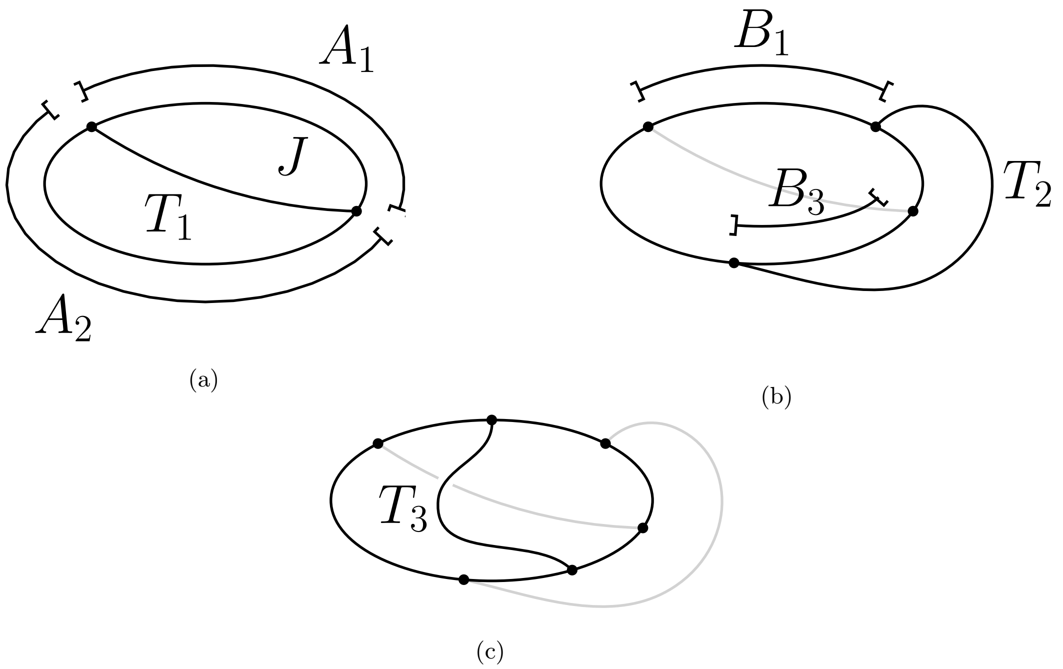

Figure 2.5

We propose that the arcs T_1, T_2 and T_3 violate Lemma 2.3.6. Indeed, $]T_1[$ and $]T_2[$ lie in distinct components of $X - J$ by construction, and likewise, our choice of T_3 ensures that T_3 is disjoint from both T_1 and T_2. Moreover, we positioned the end points of T_3 precisely to be separated in J both by the end points of T_1 and by those of T_2, in contradiction of Lemma 2.3.6. It follows that $X - J$ has exactly two components, and the Jordan curve theorem for Zippin spaces is proved. ◻

2.4 Schoenflies' Theorem

2.4.1 Lavrentieff's Theorem

We now know that, given some Jordan curve J in some Zippin space X, the complement $X - J$ consists of precisely two components, with J as their common boundary. Following the rough roadmap set out in Section 2.2, our task now is that of proving that, under the additional assumption that X is compact, the closure of each such component is a 2-cell.

As mentioned in Section 2.2, this proof shall revolve around first defining a homeomorphism on a dense subspace, and then using metric completeness to extend said homeomorphism. It is *Lavrentieff's theorem* which allows us to do so, and for the sake of comprehensiveness we prove this theorem before proceeding. (Strictly speaking, what we prove is actually less general than Lavrentieff's theorem. Lavrentieff's original result [Lav24] concerns the extension of a merely continuous–perhaps not uniformly so–homeomorphism to G_δ subspaces lying between its original domain and codomain, and the closures of the same.)

We divide our proof of Lavrentieff's theorem into two stages. In Lemma 2.4.1, we prove a preliminary extension result, and in Lemma 2.4.2, we use said result to construct the desired homeomorphism.

Lemma 2.4.1 (Kuratowski's lemma)**.** *Let (X, d_X) and (Y, d_Y) denote a pair of complete metric spaces, with subspaces $A \subset X$ and $B \subset Y$.*

Suppose that $f\colon A \to B$ denotes some uniformly continuous function. Then, f permits a uniformly continuous extension $\widetilde{f}\colon \operatorname{cl}_X A \to \operatorname{cl}_Y B$.

Proof. Given some point $x \in \operatorname{cl}_X A$, we can find some sequence $(x_n)_{n=1}^\infty$ of points of A with limit x. The uniform continuity of f ensures that the image $(f(x_n))_{n=1}^\infty$ is a Cauchy sequence of points in B, and the completeness of (Y, d_Y) ensures that this sequence has some limit in $\operatorname{cl}_Y B$. We propose first that this limit is independent of the sequence with which we start, in the sense that if $(x_n)_{n=1}^\infty$ and $(y_n)_{n=1}^\infty$ are two sequences in A with limit y, then $(f(x_n))_{n=1}^\infty$ and $(f(y_n))_{n=1}^\infty$ have the same limit in $\operatorname{cl}_X B$.

Indeed, suppose that $f(x_n) \to a$ and $f(y_n) \to b$, and consider $d_Y(a, b)$. Using the triangle inequality twice to observe that

$$d_Y(a, b) \le d_Y(a, f(x_n)) + d_Y(f(x_n), f(y_n)) + d_Y(f(y_n), b) \tag{2.5}$$

for every $n \in \mathbb{N}$, fix some $\varepsilon > 0$. We can immediately deduce that

$$d_Y(a, f(x_n)) < \frac{\varepsilon}{3} \text{ and } d_Y(f(y_n), b) < \frac{\varepsilon}{3} \tag{2.6}$$

for sufficiently-large values of n, so we wish to bound the value of the second term on the right-hand side of (2.5).

Because f is uniformly continuous, we know that there must exist some $\delta > 0$ such that if $d_X(x_n, y_n) < \delta$, then $d_Y(f(x_n), f(y_n)) < \frac{\varepsilon}{3}$. Further, since the sequences $(x_n)_{n=1}^\infty$ and $(y_n)_{n=1}^\infty$ both share the same limit, we know that we can always arrange that $d_X(x_n, y_n) < \delta$ by choosing a large enough value of n. This proves our proposition.

Thus, given a point $x \in \operatorname{cl}_X A$, we tentatively define

$$\widetilde{f}(x) = \lim_{n\to\infty} f(x_n), \text{ where } x_n \to x \text{ in } A, \tag{2.7}$$

knowing from the above that this is at least a well defined extension of f–although its uniform continuity remains to be seen. We prove this by similar reasoning to the above.

Indeed, consider some pair of points $x, y \in \operatorname{cl}_X A$ which are the limits of sequences $(x_n)_{n=1}^\infty$ and $(y_n)_{n=1}^\infty$ in A respectively. Then, we know that

$$d_Y\big(\widetilde{f}(x), \widetilde{f}(y)\big) \le d_Y\big(\widetilde{f}(x), f(x_n)\big) + d_Y\big(f(x_n), f(y_n)\big) + d_Y\big(f(y_n), \widetilde{f}(y)\big) \tag{2.8}$$

for each $n \in \mathbb{N}$. Fix some $\varepsilon > 0$. By definition, $f(x_n) \to \widetilde{f}(x)$ and $f(y_n) \to \widetilde{f}(y)$, so that

$$d_Y\big(\widetilde{f}(x), f(x_n)\big) < \frac{\varepsilon}{3} \text{ and } d_Y\big(f(y_n), \widetilde{f}(y)\big) < \frac{\varepsilon}{3} \tag{2.9}$$

for sufficiently large values of n. Moreover, we can exploit the uniform continuity of f to deduce the existence of some $\delta > 0$ such that if $d_X(x_n, y_n) < 2\delta$, then $d_Y(f(x_n), f(y_n)) < \frac{\varepsilon}{3}$, for all $n \in \mathbb{N}$.

We now suggest that if $d_X(x, y) < \delta$, then $d_X(x_n, y_n) < 2\delta$ for sufficiently large values of n; combining this with (2.8) and (2.9) will complete the proof. This is essentially immediate, for since y lies in the δ-ball in X centred at x, we know that the sequences $(x_n)_{n=1}^{\infty}$ and $(y_n)_{n=1}^{\infty}$ are both eventually in this ball. We conclude that $\widetilde{f}$ is indeed uniformly continuous. □

Lemma 2.4.2 (Lavrentieff). *Let (X, d_X) and (Y, d_Y) denote a pair of complete metric spaces, with subspaces $A \subset X$ and $B \subset Y$.*

Suppose that $f\colon A \to B$ denotes some uniform isomorphism. (That is, let f denote a uniformly continuous bijection with uniformly continuous inverse.) Then, f can be extended to a uniform isomorphism $\widetilde{f}\colon \operatorname{cl}_X A \to \operatorname{cl}_Y B$.

Proof. We can use Kuratowski's lemma twice to produce uniformly continuous extensions

$$\widetilde{f}\colon \operatorname{cl}_X A \to \operatorname{cl}_Y B \text{ and } \widetilde{g}\colon \operatorname{cl}_Y B \to \operatorname{cl}_X A \tag{2.10}$$

of f and f^{-1} respectively, so that we need only prove that these extensions are mutually inverse.

Suppose, then, that $(x_n)_{n=1}^{\infty}$ is some sequence of points in A with a limit $x \in \operatorname{cl}_X A$. We know that $f(x_n) \to \widetilde{f}(x)$, and since $f(x_n) \in B$ for every $n \in \mathbb{N}$, we also know that

$$(f^{-1} \circ f)(x_n) = x_n \to (\widetilde{g} \circ \widetilde{f})(x). \tag{2.11}$$

It follows that $\widetilde{g} \circ \widetilde{f} = \mathrm{id}_{\operatorname{cl}_X A}$, and entirely analogous reasoning reveals that $\widetilde{f} \circ \widetilde{g} = \mathrm{id}_{\operatorname{cl}_Y B}$ as well. □

2.4.2 Arc Complexes

With Lavrentieff's theorem in hand, we have a means of extending a homeomorphism from a dense subspace to an entire space–but we are yet to see how we might define an appropriate dense subspace in this setting. At the heart of our construction is, given some complementary domain Y of a Jordan curve J in a compact Zippin space X, an iterative subdivision of $\operatorname{cl}_X Y$ into ever-smaller pieces. This subdivision is made possible by the so-called *θ-curve lemma*, which essentially states that we can use arcs to bisect complementary domains of Jordan curves.

Lemma 2.4.3 (θ-curve lemma). *Let T_1, T_2 and T_3 denote three arcs in some Zippin space X, all of which have the same end points but are otherwise disjoint. The union $T_1 \cup T_2 \cup T_3$ is called a* θ-curve.

Such a θ-curve has precisely three complementary domains in X. One has boundary $T_1 \cup T_2$, another has boundary $T_2 \cup T_3$, and another has boundary $T_1 \cup T_3$.

Proof. The unions $T_1 \cup T_2$, $T_2 \cup T_3$ and $T_1 \cup T_3$ are all Jordan curves, so that we can use the Jordan curve theorem (Lemma 2.3.7) to deduce that

- $X - (T_1 \cup T_2)$ has components D_1 and E_1, where $]T_3[\subset E_1$;
- $X - (T_2 \cup T_3)$ has components D_2 and E_2, where $]T_1[\subset E_2$; and
- $X - (T_1 \cup T_3)$ has components D_3 and E_3, where $]T_2[\subset E_3$.

Notice that D_1, D_2 and D_3 are all components of $X - (T_1 \cup T_2 \cup T_3)$. To see why, observe that D_1 is certainly a connected subspace of $X - (T_1 \cup T_2 \cup T_3)$, and that $\operatorname{fr} D_1 = T_1 \cup T_2$. Were there to exist some component A of $X - (T_1 \cup T_2 \cup T_3)$ with $D_1 \subsetneq A$, then selecting points $x \in D_1$ and $y \in A - D_1$

and using Lemma A.5.4 to find an arc $[xy] \subset A$, Lemma 2.3.3 would imply that $(T_1 \cup T_2) \cap A \neq \emptyset$, which is manifestly a contradiction. The same style of reasoning works equally well for D_2 and D_3.

Thus, it will suffice for us to prove that there is no fourth complementary domain of $T_1 \cup T_2 \cup T_3$ in X. Supposing that A is some fourth component of $X - (T_1 \cup T_2 \cup T_3)$, our strategy shall be to use Lemma 2.3.6 to arrive at a contradiction. The vital ingredient in this procedure will be that there exist points interior to T_1, T_2 and T_3 which are arc-accessible from A, as we now demonstrate.

Supposing that $]T_1[\cap \operatorname{fr} A = \emptyset$ and using Corollary 2.3.2, we find that $\operatorname{fr} A \subset T_2 \cup T_3$. In fact, we can see that $\operatorname{fr} A = T_2 \cup T_3$ by a similar argument to that which proved Lemma 2.3.5. Indeed, were there to exist some point $z \in (T_2 \cup T_3) - \operatorname{fr} A$, we would be able to find some connected neighbourhood $U \subset T_2 \cup T_3$ of z disjoint from A. Then, Corollary 2.3.4 would tell us that the arc $(T_2 \cup T_3) - U$ separates X, which cannot be the case when X is Zippin.

This contradicts the Jordan curve theorem: the only components of $X - (T_2 \cup T_3)$ are D_2 and E_2, but the above shows that A is a third such component. It follows that $]T_1[\cap \operatorname{fr} A \neq \emptyset$, and two more arguments of a similar nature reveal that $]T_2[\cap \operatorname{fr} A \neq \emptyset$ and $]T_3[\cap \operatorname{fr} A \neq \emptyset$ as well; recalling Lemma A.5.8 allows us to deduce the existence of points of T_1, T_2 and T_3 which are arc-accessible from A, as claimed.

We are now ready to construct the desired violation of Lemma 2.3.6. Let S denote some arc with one end point interior to T_1 and the other interior to T_2, such that $]S[\subset A$. Further, let S_1 denote some arc such that $]S_1[\subset D_1$, the end points of which separate those of S on $T_1 \cup T_2$. If we let $S_2 = T_3$, we can see immediately that

- S_1 and S_2 lie in different components of $X - (T_1 \cup T_2)$,
- $S \cap S_1 = S \cap S_2 = \emptyset$,
- The end points of S_1 separate those of S in $T_1 \cup T_2$, because we chose S_1 precisely so that this would be so; and
- The end points of S_2 separate those of S in $T_1 \cup T_2$, because one end point of S is interior to T_1 and the other to T_2, while the end points of S_2 are those of T_1 and T_2.

This is a contradiction of Lemma 2.3.6, and the proof is complete. □

The value of the θ-curve lemma for our current purposes derives from the fact that it holds just as well in the 2-cell D^2 as it does in $\operatorname{cl}_X Y$, where we recall that Y is some complementary domain in X of a Jordan curve $J \subset X$. This will allow us to produce subdivisions of the two spaces which are, in some sense, 'combinatorially equivalent'. Let us elaborate on this idea. If we take $\operatorname{cl}_X Y$ and draw across it some J-spanning arc (Figure 2.6a), then the θ-curve lemma tells us that we have divided $\operatorname{cl}_X Y$ into two components; likewise, an S^1-spanning arc in D^2 (Figure 2.6b) splits D^2 in two.

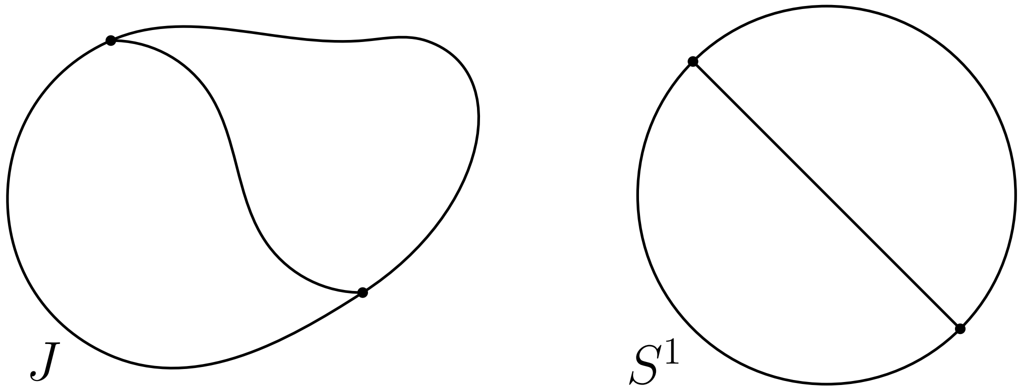

(a) (b)

Figure 2.6

Crucially, we can repeat this for the resulting complementary domains in $\mathrm{cl}_X Y$ and in D^2. We sketch the outcome of adding a single additional arc in Figure 2.7, and of adding several in Figure 2.7.

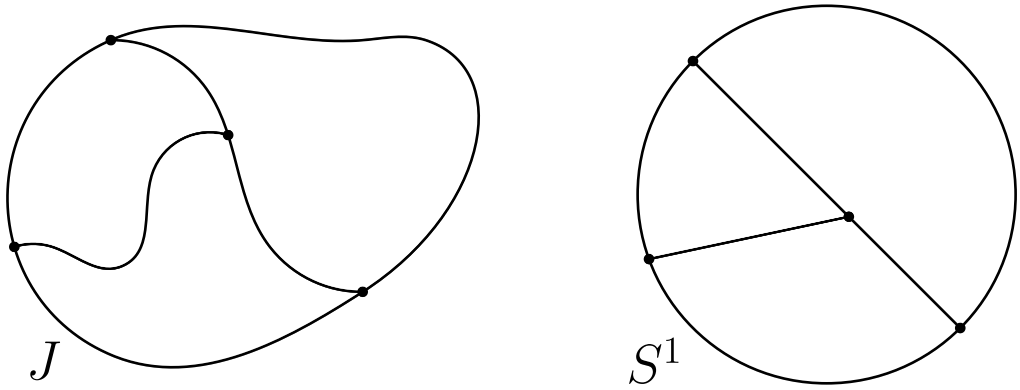

(a) (b)

Figure 2.7: An additional spanning arc can be added to each of the configurations sketched in Figure 2.6, while retaining their 'combinatorial equivalence'.

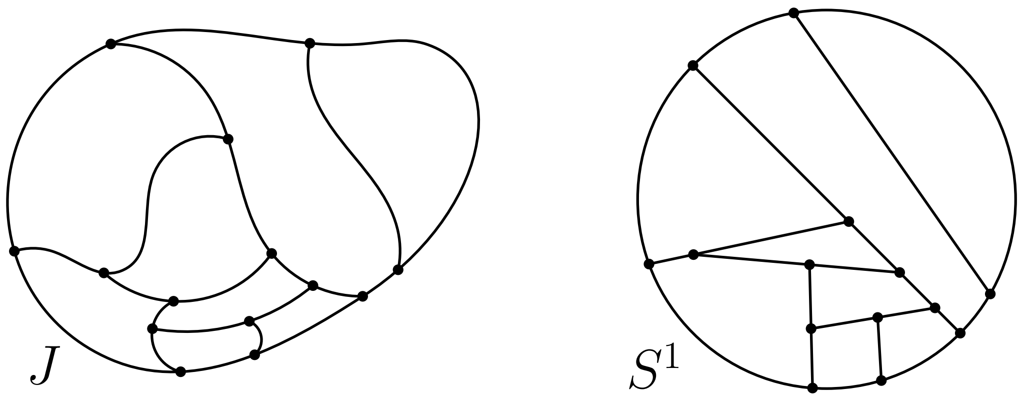

(a) (b)

Figure 2.8: Finitely many additional spanning arcs can be added to each of the configurations sketched in Figure 2.6, while retaining their 'combinatorial equivalence'.

This notion of subdivision by finitely many spanning arcs, and the fact that we can produce 'essentially identical' such subdivisions of both $\mathrm{cl}_X Y$ and of D^2, is made formal by the following definitions and lemma. Although notationally dense, these say nothing fundamentally new: the salient features have already been covered in the preceding discussion and Figures 2.6-2.8. In Definition 2.4.4, the subspace A is, for our purposes, either the Jordan curve $J \subset X$ or the boundary $S^1 \subset D^2$.

Definition 2.4.4. *Consider some topological space X with a subspace $A \subset X$. Let $\{T_1, T_2, \ldots, T_n\}$ denote some collection of finitely many arcs in X with the property that, for each $j \in \{1, 2, \ldots, n\}$, the arc T_j spans*[3] *$A \cup \bigcup_{i=1}^{j-1} T_i$. Then, the union $C = A \cup \bigcup_{i=1}^{n} T_i$ is called an* (A-spanning) arc complex *in X.*

We define the vertices*,* edges *and* domains *of such an arc complex as follows:*

- *The* vertices *are precisely the end points of the arcs $\{T_1, T_2, \ldots, T_n\}$;*
- *The* edges *are precisely the subarcs of the arcs $\{T_1, T_2, \ldots, T_n\}$ which span the arc complex C; and*

[3] We once more recall that the arc T_j is said to *span* A if the intersection $A \cap T_j$ consists precisely of the end points of T_j.

- *The* domains *are precisely the components of* $X - C$.

Suppose that D_1 *and* D_2 *denote some pair of domains of some arc complex* $C \subset X$. *We say that* D_1 *and* D_2 *are* adjacent *if* $\operatorname{fr} D_1 \cap \operatorname{fr} D_2 \neq \emptyset$.

Definition 2.4.5. *Let* X *and* Y *denote topological spaces with arc complexes* $C \subset X$ *and* $C' \subset Y$, *where the subspaces of* X *and* Y *spanned by* C *and* C' *respectively are homeomorphic to one another. If there exist incidence-preserving bijections between the vertices, edges and domains of* C *and* C', *then the arc complexes* C *and* C' *are said to be* isomorphic.

Explicitly, let C_0, C_1 *and* C_2 *denote, respectively, the sets of edges, vertices and domains of* C; *similarly, let* C_0', C_1' *and* C_2' *denote the sets of edges, vertices and domains of* C'. *Then, an* isomorphism between C and D *consists of bijections*

$$f_0 \colon C_0 \to C_0', \; f_1 \colon C_1 \to C_1' \text{ and } f_2 \colon C_2 \to C_2'$$

such that

- *If* $x \in C_0$ *and* $S \in C_1$, *then* $x \in S$ *if and only if* $f_0(x) \in f_1(S)$; *and*
- *If* $S \in C_1$ *and* $D \in C_2$, *then* $S \subset \operatorname{fr}_X D$ *if and only if* $f_1(S) \subset \operatorname{fr}_Y f_2(D)$.

Lemma 2.4.6. *Let* Y *denote some complementary domain of a Jordan curve* J *in a compact Zippin space* X. *Every* J*-spanning arc complex in* $\operatorname{cl}_X Y$ *is isomorphic to an* S^1*-spanning arc complex in* D^2; *conversely, every* S^1 *spanning arc complex in* D^2 *is isomorphic to a* J*-spanning arc complex in* $\operatorname{cl}_X Y$.

Proof. Consider some J-spanning arc complex $C = J \cup \bigcup_{i=1}^n T_i$ in $\operatorname{cl}_X Y$. We shall prove the result by induction on n.

The basis case in which $n = 0$ is immediate, so suppose that the arc complex $J \cup \bigcup_{i=1}^{n-1} T_i$ is isomorphic to some S^1-spanning arc complex in D^2. Let S_1 and S_2 denote the edges of C containing the end points of the arc T_n, and D the domain of C containing $]T_n[$.

Suppose that the aforementioned isomorphism identifies the edges S_1 and S_2 with edges S_1' and S_2' respectively, and the domain D with a domain D'. Use Lemma A.5.4 to find some arc $T' \subset \operatorname{cl}_{D^2} D'$ such that $]T'[\subset D'$, with end points on S_1' and S_2'.

Now, use the θ-curve lemma to deduce that T_n has two complementary domains in D, and T' two complementary domains in D'. Identifying these complementary domains with one another in such a way as to preserve incidence, and identifying T_n with T', extends our isomorphism to an isomorphism between C and an S^1-spanning arc complex in D^2, completing the induction. The reverse proposition follows by identical logic. □

2.4.3 Subdividing Complementary Domains

With Lavrentieff's theorem, the notion of an arc complex, and Lemma 2.4.6 in hand, we are well on our way to defining a homeomorphism between (dense subspaces of) $\operatorname{cl}_X Y$ and D^2. Our next step is to show that we can find arc complexes in $\operatorname{cl}_X Y$ with 'arbitrarily small domains', and we approach this in two distinct phases.

The first, to which this section is dedicated, is to simply show that we need only finitely many J spanning arcs in order to subdivide $\operatorname{cl}_X Y$ into components of arbitrarily small diameter. However, in doing so, we have no control over how the arcs involved intersect, and in particular they in general will fail to comprise an arc complex. Accordingly, the second stage is to show that we can always produce from such a collection of arcs an arc complex with domains also of arbitrarily small diameter; this is the subject of the next section.

As it was for Section 2.4.2, the essential tool for the subdivisions of this section is the θ-curve lemma. However, we shall also have need of some auxiliary results: namely, compact Zippin spaces are cyclic (Definition A.10.1), as are the closures of complementary domains of Jordan curves in the same. (Given how pathological arbitrary Jordan curves can be, this result is perhaps more surprising than it would first appear.)

These cyclicity results are foundational to the remainder of our proof of Zippin's sphere characterization, and the requisite machinery is developed in Sections A.8-A.10. In particular, the characterisation given in Theorem A.10.3 will be used extensively in what follows; for the sake of convenience, we restate the most important elements of said sections below before proceeding.

Definition A.10.1. *Let X denote some Peano continuum. If, for each pair of points $x, y \in X$, there exists some Jordan curve $J_{xy} \subset X$ which contains both x and y, then X is said to be* cyclically connected.

The cyclically connected subcontinua of X which are maximal with respect to set-theoretic inclusion are termed the cyclic components *of X.*

Definition A.8.1. *Define an equivalence relation $\sim$ on the set of non-cut points of some Peano continuum X by declaring that*

$$x \sim y \text{ if and only if no point of } X \text{ separates } x \text{ and } y.$$

We call this the conjugacy relation *on X, and if $x \sim y$, then the points x and y are said to be* conjugate.

The equivalence class of a point $x \in X$ with respect to $\sim$ is called the conjugacy class *of x, and we denote this by $C_x \subset X$.*

Theorem A.10.3. *We say that a Peano continuum X has the* three point property *if, for every trio $x, y, z \in X$ of distinct points, there exists some arc $[xyz] \subset X$.*

For every Peano continuum X, the following are equivalent:

1. *X has no cut points,*
2. *X has the three point property, and*
3. *X is cyclically connected.* □

Let us now prove that every Zippin space is indeed cyclic in the sense of Definition A.10.1.

Lemma 2.4.7. *Every compact Zippin space is cyclic.*

Proof. We shall argue by contraposition. Indeed, let X denote some Peano continuum which fails to be cyclic; then, it follows from Lemma A.10.3 there must exist some point $x \in X$ such that the conjugacy class C_x (Definition A.8.1) is not the entirety of X.

Consider some connected component A of the complement $X - C_x$. We know from Lemma A.8.3 that the boundary of A meets C_x in precisely one point, which we shall denote by y. Now, Lemma A.8.7 and Lemma A.10.3 together assert that y lies on some Jordan curve $J \subset C_x$, and that y is interior to some closed arc $T \subset C_x$. In particular, T separates X, and so X cannot possibly be Zippin. □

Lemma 2.4.8. *If Y denotes some complementary domain of a Jordan curve J in a compact Zippin space X, then $\text{cl}_X Y$ is cyclic.*

Proof. We shall argue by contradiction. Indeed, suppose that J is some Jordan curve in a Zippin space X, with Y some component of $X - J$ with a closure which fails to be cyclic. Then, we know from Lemma A.10.3 that $\text{cl}_X Y$ has some cut point p.

Given some separation $\text{cl}_X Y - \{p\} = Y_1 \mid Y_2$, we select points $y_1 \in Y_1$ and $y_2 \in Y_2$. According to Lemma A.5.8, we can find an arc $[ay_1] \subset \text{cl}_X Y_1$ such that $[ay_1] \cap J = \{a\}$, and an arc $[y_2b] \subset \text{cl}_X Y_2$ such that $[by_2] \cap J = \{b\}$; moreover, an arc $[y_1y_2] \subset \text{cl}_X Y$ necessarily contains the cut point p. It follows that the union $[ay_1] \cup [y_1y_2] \cup [y_2b]$ contains some J-spanning arc $[apb] \subset \text{cl}_X Y$.

Let us apply the θ-curve lemma to the union $J \cup [apb]$, yielding components

- A_1, with boundary $[apb] \cup \text{cl}_X J_1$;
- A_2, with boundary $[apb] \cup \text{cl}_X J_2$; and

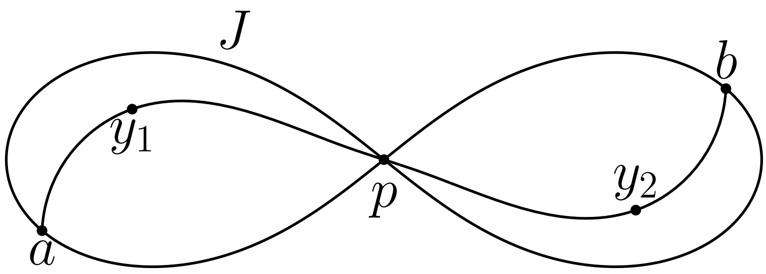

Figure 2.9

- A_3, with boundary J;

of $X - (J \cup [apb])$, where we denote by J_1 and J_2 the two complementary domains of $\{a, b\}$ in J.

Now, select any point $z_1 \in]ap[\subset [apb]$. This is a boundary point of A_1, and there must therefore exist some sequence in A_1 with limit z_1; crucially, the knowledge that Y is locally connected allows us to deduce that Y_1 is a neighbourhood of z_1, so that in particular this sequence is eventually in Y_1. Thus, we find that $A_1 \cap Y_1 \neq \emptyset$.

However, applying similar reasoning to any point $z_2 \in]pb[\subset [apb]$ reveals that $A_1 \cap Y_2 \neq \emptyset$ as well. This is a contradiction: the connected subspace $A_1 \subset Y - \{p\} = Y_1 \mid Y_2$ cannot possibly meet both Y_1 and Y_2. □

Corollary 2.4.9. *If Y denotes some complementary domain of a Jordan curve J in a compact Zippin space X, then Y is cyclic.*

Proof. The method of proof used for Lemma 2.4.8 works just as well to show that Y is cyclic. □

We now begin the process of subdivision with a mild claim: we propose that it is always possible to separate any pair of distinct points of J in $\text{cl}_X Y$ by an arc which spans J.

Lemma 2.4.10. *Let Y denote a complementary domain of a Jordan curve J in a compact Zippin space X, and consider some pair of distinct points $x, y \in J$.*

Suppose that a and b lie in different components of $J - \{x, y\}$. Then, any J-spanning arc $[ab] \subset \text{cl}_X Y$ must separate x and y in $\text{cl}_X Y$.

Proof. We know from the θ-curve lemma that, if $[ab]$ is some J-spanning arc in $\text{cl}_X Y$, then the union $J \cup [ab]$ has precisely three complementary domains in X, namely

- A_1, with boundary $[axb] \cup [ab]$ relative to X;
- A_2, with boundary $[ayb] \cup [ab]$ relative to X; and
- A_3, with boundary J relative to X;

where $[axb]$ and $[ayb]$ denote the appropriate subarcs of J, as sketched in Figure 2.10.

Relative to $\text{cl}_X Y$, the complementary domain A_1 has boundary $[ab]$. Lemma 2.3.3 therefore requires that all arcs $[xy] \subset \text{cl}_X Y$ intersect the arc $[ab]$, whence it follows that x and y belong to separate components of $\text{cl}_X Y - [ab]$. □

Refining this result, we can prove that any pair of distinct points of $\text{cl}_X Y$–whether or not they belong to J–can be separated in $\text{cl}_X Y$ by some J-spanning arc.

Lemma 2.4.11. *Let $x, y \in \text{cl}_X Y$ denote some pair of distinct points. There exists some arc in $\text{cl}_X Y$ which spans J and separates x from y in $\text{cl}_X Y$.*

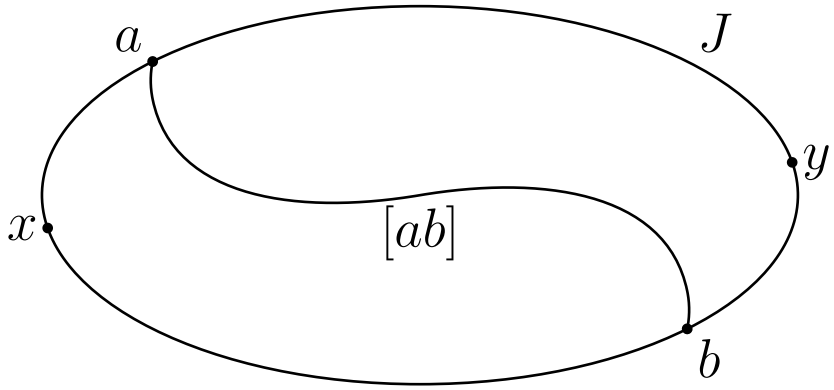

Figure 2.10

Proof. We proceed by exhausting the different options for where the points x and y may lie. There are three possibilities:

- Both x and y may lie on J, in which case we are done, thanks to Lemma 2.4.10;
- Precisely one of x and y may lie on J; or
- Both x and y may be interior to y.

Suppose, then, that $y \in J$, and notice that there exists an arc $[ayb] \subset \mathrm{cl}_X Y$ which spans J and does not contain x. (To see why, use Lemma 2.4.8 together with Lemma A.10.3 to see that x cannot be a cut point of $\mathrm{cl}_X Y$; combining this with Lemma A.3.5 allows us to find a neighbourhood $U \subset \mathrm{cl}_X Y$ of x which does not separate $\mathrm{cl}_X Y$, such that $y \notin \mathrm{cl}_Y U$. Selecting any pair of distinct points $a, b \in J$ and recalling Lemma A.10.3 yields an arc $[ayb] \subset \mathrm{cl}_X Y - \mathrm{cl}_Y U$ which, by passing to a subarc if necessary, can be assumed to span J with no loss of generality.)

Now, suppose additionally that $x \in J$. Denote by $[axb] \subset J$ the appropriate arc of J, and define the Jordan curve $K = [axb] \cup [ayb]$. Using the θ-curve lemma, we find that K is the boundary in X of some component of $X - (J \cup [ayb])$, which we denote by A. Choosing points $z_1, z_2 \in K$ which are separated by $\{x, y\}$ and are arc accessible from A, we can find some K-spanning arc $[z_1 z_2] \subset \mathrm{cl}_X A$. Again using the θ-curve lemma, we see that $[z_1 z_2]$ separates x and y in $\mathrm{cl}_X A$, and therefore in $\mathrm{cl}_X Y$. (Figure 2.11a.)

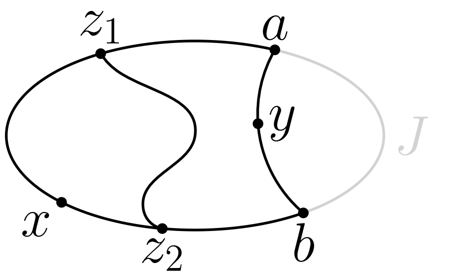

(a)

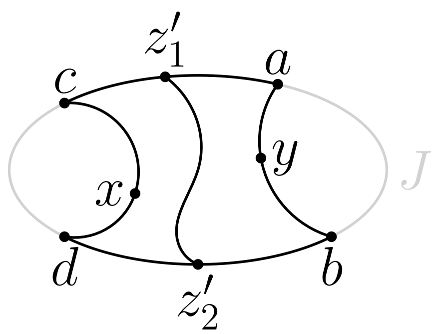

(b)

Figure 2.11

If $x \in Y$ instead, then let A be as defined above; recycling our earlier reasoning supplies us with a J-spanning arc $[cxd] \subset \mathrm{cl}_X A$. Considering the arcs $[ac], [bd] \subset J$, we find that the union $K' = [ac] \cup [ayb] \cup [bd] \cup [cxd]$ is a Jordan curve (Figure 2.11b); by the Jordan curve theorem, K' has precisely two complementary domains in X, of which one is a subspace of Y. Denote this complementary domain by A'.

Selecting points $z_1' \in \,]ac[$ and $z_2' \in \,]bd[$ for which there exists a J-spanning arc $[z_1'z_2'] \subset A'$, we can apply the θ-curve lemma to the union $K' \cup [z_1'z_2']$ in order to deduce that $[z_1'z_2']$ separates x and y in $\mathrm{cl}_X\, A'$, and therefore in $\mathrm{cl}_X\, Y$. This completes the proof. □

Now we proceed to show that it takes only finitely many J-spanning arcs to separate points of disjoint compact subspaces of $\mathrm{cl}_X\, Y$. We achieve this in two steps: first we separate single points from points of compact subspaces, and then points of compact subspaces from points of compact subspaces.

Lemma 2.4.12. *Let $C \subsetneq \mathrm{cl}_X\, Y$ denote some compact subspace, and consider a point $x \in \mathrm{cl}_X\, Y - C$. There exists a finite collection $\mathcal{T}$ of J-spanning arcs in $\mathrm{cl}_X\, Y$ which* separates x from points of C, *in the sense that for each point $y \in C$, there exists some member of $\mathcal{T}$ which separates x and y in $\mathrm{cl}_X\, Y$.*

Proof. Given any point $y \in C$, we can use the preceding lemma to find a J-spanning arc $T_y \subset \mathrm{cl}_X\, Y$ which separates x and y in $\mathrm{cl}_X\, Y$. If U_y denotes the component of $\mathrm{cl}_X\, Y - T_y$ which contains y, then the collection $\{U_y\}_{y \in C}$ comprises an open cover of C, which by compactness must be essentially finite. Letting $\{U_{y_1}, U_{y_2}, \ldots, U_{y_n}\}$ denote some finite subcover of $\{U_y\}_{y \in C}$, the collection

$$\mathcal{T} = \{T_{y_1}, T_{y_2}, \ldots, T_{y_n}\} \tag{2.12}$$

separates x from points of C. □

Lemma 2.4.13. *Let $C, D \subsetneq \mathrm{cl}_X\, Y$ denote some pair of disjoint compact subspaces. There exists a finite collection $\mathcal{T}$ of J-spanning arcs in $\mathrm{cl}_X\, Y$ which* separates points of C from points of D, *in the sense that for each $x \in C$ and for each $y \in D$, there exists some member of $\mathcal{T}$ which separates x and y in $\mathrm{cl}_X\, Y$.*

Proof. For each point $x \in C$, we can use the previous result to find some collection $\{T_1^x, T_2^x, \ldots, T_{n_x}^x\}$ of finitely many arcs which separate x from points of D. For each such arc, define U_i^x to be the component of $\mathrm{cl}_X\, Y - T_i^x$ which contains x. Then, compactness asserts that

$$\{U_i^x \mid x \in C \text{ and } i \in \{1, 2, \ldots, n_x\}\} \tag{2.13}$$

is an essentially finite open cover of C. Much as in our proof of Lemma 2.4.12, the arcs corresponding to elements of some fixed finite subcover of (2.13) form a finite collection of J-spanning arcs in $\mathrm{cl}_X\, Y$ which separate points of C from points of D, as needed. □

We can now fulfil our ambition of subdividing $\mathrm{cl}_X\, Y$ into 'arbitrarily small pieces'. All of the hard work has been done; the only thing that remains is for us to apply Lemma 2.4.13 to the right family of subspaces.

Lemma 2.4.14. *Recalling that we insisted in Definition A.2.1 that all continua be metrizable, choose any metric d which topologizes X, and fix some $\delta > 0$. There exists a finite collection $\mathcal{T}$ of J-spanning arcs in $\mathrm{cl}_X\, Y$ such that each component of $\mathrm{cl}_X\, Y - \bigcup_{T \in \mathcal{T}} T$ has diameter (with respect to d) strictly less than δ.*

Proof. If $\mathrm{diam}\, \mathrm{cl}_X\, Y < \delta$, then we have nothing to do: any single J-spanning arc in $\mathrm{cl}_X\, Y$ satisfies the conclusion of the lemma.

Suppose, then, that $\mathrm{diam}\, \mathrm{cl}_X\, Y \geq \delta$. Compactness allows us to cover $\mathrm{cl}_X\, Y$ with a collection $\mathcal{B} = \{B_1, B_2, \ldots, B_n\}$ of finitely many closed balls of radius $\frac{\delta}{4}$, where $n > 1$; moreover, there exists at least one disjoint pair of elements of $\mathcal{B}$.

To each pair of disjoint elements $B_i, B_j \in \mathcal{B}$, Lemma 2.4.13 allows us to associate some finite collection $\mathcal{T}_{ij}$ of J-spanning arcs in $\mathrm{cl}_X\, Y$ which separate points of B_i from points of B_j. Defining $\mathcal{T}$ to be the union of all such $\mathcal{T}_{ij}$, we claim that every component of $\mathrm{cl}_X\, Y - \bigcup_{T \in \mathcal{T}} T$ has diameter strictly less than δ. Indeed, select any pair of points $x, y \in \mathrm{cl}_X\, Y$ with $d(x, y) \geq \delta$; such points must belong to disjoint elements of $\mathcal{B}$, so that x and y are separated in $\mathrm{cl}_X\, Y$ by some element of $\mathcal{T}$. □

2.4.4 Subdivision by Arc Complexes

Using Lemma 2.4.14, we can fix some arbitrary $\delta > 0$ and find some finite collection $\mathcal{T}$ of J-spanning arcs which divide $\mathrm{cl}_X Y$ into pieces of diameter no greater than δ–but if we are to use Lemma 2.4.6 to make a connection with the 2-cell, then we need to produce an arc complex from $\mathcal{T}$.

Examining Definition 2.4.4, we see how $\mathcal{T}$ may fail to be an arc complex: concretely, there might exist no enumeration $\mathcal{T} = \{T_1, T_2, \ldots, T_n\}$ such that T_j spans $J \cup \bigcup_{i=1}^{j-1} T_j$ for every $j \in \{1, 2, \ldots, n\}$. (For instance, consider what happens if two members of $\mathcal{T}$ intersect along a non-degenerate subarc.)

To correct this deficiency, then, we might try to produce an arc complex from $\mathcal{T}$ by working 'one step at a time'. Deciding upon some arbitrary enumeration $\mathcal{T} = \{T_1, T_2, \ldots, T_n\}$, we first form the union $J \cup T_1$. Then, if we let $\mathcal{S}_2$ denote the collection of $(J \cup T_1)$-spanning subarcs of T_2, we could form the union $J \cup T_1 \cup \bigcup_{S \in \mathcal{S}_2} S$, continuing inductively in this manner until all elements of $\mathcal{T}$ have been accounted for.

There is, however, a subtlety here. Our definition of an arc complex, and the crucial isomorphism between arc complexes in $\mathrm{cl}_X Y$ and in D^2, relies on the assumption that arc complexes can be assembled from only finitely many arcs–and it is entirely possible that, given two arcs T_1 and T_2, there exist infinitely many T_1-spanning subarcs of T_2. As an illustration, we give an explicit example in which T_2 contains a countable infinitude of T_1-spanning subarcs.

Example 2.4.15 (An arc with infinitely many subarcs spanning another arc)**.** Consider the arcs $T_1, T_2 \subset \mathbb{R}^2$ sketched in Figure 2.12, in which T_1 is the graph of the function

$$x \mapsto \begin{cases} 0 & \text{if } x = 0 \text{ and} \\ x \sin\left(\frac{1}{x}\right) & \text{if } x \in \left]0, 1\right], \end{cases} \tag{2.14}$$

and T_2 is simply the product $[0, 1] \times \{0\}$.

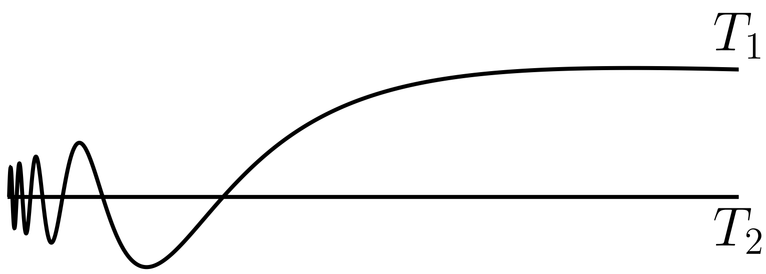

Figure 2.12

Of course, every subarc of T_2 of the form $\left[\frac{1}{(n+1)\pi}, \frac{1}{n\pi}\right] \times \{0\}$ for some $n \in \mathbb{N}$ spans T_1, and there are infinitely many such subarcs. □

We can see in this example that, although there are infinitely many spanning subarcs, they eventually become arbitrarily short, and we can prove that a suitable generalization of this observation holds in general. In particular, we are less interested in arcs with end points which are far apart in some ambient space than we are in arcs with end points which are far apart in the arc complex we seek to construct. The following definition and lemma make this more precise.

Definition 2.4.16. *Let Z denote some arc connected subspace of $\mathrm{cl}_X Y$, and $T \subset X$ some arc with end points $x, y \in Z$. If there exists some arc $[xy] \subset Z$ such that* diam $[xy] < \varepsilon$ *for some $\varepsilon > 0$, then the arc T is said to be ε-*small *in Z. Otherwise, the arc T is said to be ε-*large *in Z.*

*If, for some pair of points $x, y \in Z$, every arc $[xy] \subset Z$ is ε-large, then we shall say that the points x and y are ε-*distant *in Z.*

Lemma 2.4.17. *Fix some $\varepsilon > 0$, and denote by $A \subset \mathrm{cl}_X Y$ some compact subspace. If $T \subset \mathrm{cl}_X Y$ denotes some J-spanning arc, then there exist at most finitely many A-spanning subarcs of T which are ε-large in A.*

Proof. We first notice that there are at most countably many subarcs of T which span A, ε-large or otherwise. Indeed, the interior of each such subarc is a component of the complement $T - A$, and $T - A$ is both locally connected and second countable. Were such a space to have uncountably many components, then it could not possibly have a countable basis: the set of components of a locally connected space is a collection of disjoint open subspaces.

From here, we argue by contradiction. Suppose that there exist infinitely many ε-large subarcs of T which span A, and choose some parametrization $\gamma\colon [0,1] \to T$; according to the above observation, we can arrange the end points of the ε-large A-spanning subarcs of T into a sequence $(\gamma(t_n))_{n=1}^{\infty}$, ordered such that $t_i < t_j$ whenever $i < j$.

The corresponding sequence $(t_n)_{n=1}^{\infty}$ in $[0,1]$ is then monotone, so must have some limit $t \in [0,1]$. Appealing to the local connectedness of A, we find some arc-connected neighbourhood $U \subset A$ of $\gamma(t)$ of diameter less than $\frac{\varepsilon}{2}$, and we know that the sequence $(\gamma(t_n))_{n=1}^{\infty}$ is eventually in this neighbourhood. In particular, U must contain both end points of some ε-large A spanning subarc of T, but this is a contradiction: by definition, such end points cannot be connected by any arc in A of diameter less than ε. □

This result tells us that, so long as we are content to always work at some small-but-finite resolution, there are no obstructions to the construction of a arc complexes from $\mathcal{T}$. To facilitate further discussion, we encapsulate this observation in the definition of an *ε-approximation*.

Definition 2.4.18. *Let $\mathcal{T} = \{T_1, T_2, \ldots, T_n\}$ denote some collection of finitely many J-spanning arcs in $\operatorname{cl}_X Y$, and fix some $\varepsilon > 0$.*

Let $A_0 = J$, and for each $i \in \{1, 2, \ldots, n\}$, inductively define A_i as the union of A_{i-1} with all A_{i-1}-spanning subarcs of T_i which are ε-large in A_{i-1}. The resulting subspace A_n is the ε-approximation *of $\mathcal{T}$.*

Of course, the notion of an ε-approximation is only useful if it respects the properties of $\mathcal{T}$ which we care about. Let us verify that, for sufficiently small ε, passing from $\mathcal{T}$ to an ε-approximation thereof preserves the separation of points of disjoint compact subspaces, in the following sense.

Lemma 2.4.19. *Let $\mathcal{T}$ denote some collection of finitely many J-spanning arcs which separate points of disjoint compact subspaces $C, D \subsetneq \operatorname{cl}_X Y$.*

Let A_ε denote the ε-approximation of $\mathcal{T}$. There exists some $\varepsilon > 0$ such that A_ε separates points of C from points of D, in the sense that for each pair of points $x \in C$ and $y \in D$, there exists an arc in A which separates x and y in $\operatorname{cl}_X Y$.

Proof. We proceed by contradiction. Indeed, suppose that that no $\varepsilon > 0$ satisfies the statement of the lemma; then, for each $n \in \mathbb{N}$, we can find points $x_n \in C$ and $y_n \in D$ which are not separated by any arc of the $\frac{1}{n}$-approximation $A_{n^{-1}}$. The compactness of C and D demands that the sequences $(x_n)_{n=1}^{\infty}$ and $(y_n)_{n=1}^{\infty}$ have accumulation points $x \in C$ and $y \in D$ respectively.

By hypothesis, there must exist some arc $T = [ab] \in \mathcal{T}$ which separates x and y in $\operatorname{cl}_X Y$. Moreover, we can use Lemma A.5.8 to find some arc $X = [xz_x] \subset \operatorname{cl}_X Y - T$ such that $X \cap J = \{z_x\}$, along with some arc $Y = [yz_y] \subset \operatorname{cl}_X Y - T$ such that $Y \cap J = \{z_y\}$. (Figure 2.13a.)

Selecting some $k \in \mathbb{N}$ such that $\frac{1}{k} < d(T, X \cup Y)$, we define U as the $\frac{1}{k}$-neighbourhood of T in $\operatorname{cl}_X Y$, as roughly indicated in Figure 2.13b. Now, suppose that the points a and b belong to the same component of the intersection $U \cap A_{k^{-1}}$; then, there would exist some arc in $U \cap A_{k^{-1}}$ connecting a to b, and such an arc would separate z_x and z_y in J. The same line of reasoning used to prove Lemma 2.4.10 reveals that said arc would also separate x and y in $\operatorname{cl}_X Y$, and this would be the contradiction we seek: by construction, no arc contained in $A_{k^{-1}}$ can separate x and y in $\operatorname{cl}_X Y$.

Thus, we will be done if we can prove that a and b belong to the same component of $U \cap A_{k^{-1}}$. To this end, suppose otherwise, so that $U \cap A_{k^{-1}} = V_a \mid V_b$ is some separation in which $a \in V_a$ and $b \in V_b$. Arbitrarily orienting T from a to b, define p to be the last point at which T meets V_a, and a to be the first point at which T meets V_b (Figure 2.13c). The subarc $[pq] \subset T$ spans $A_{k^{-1}}$, but is

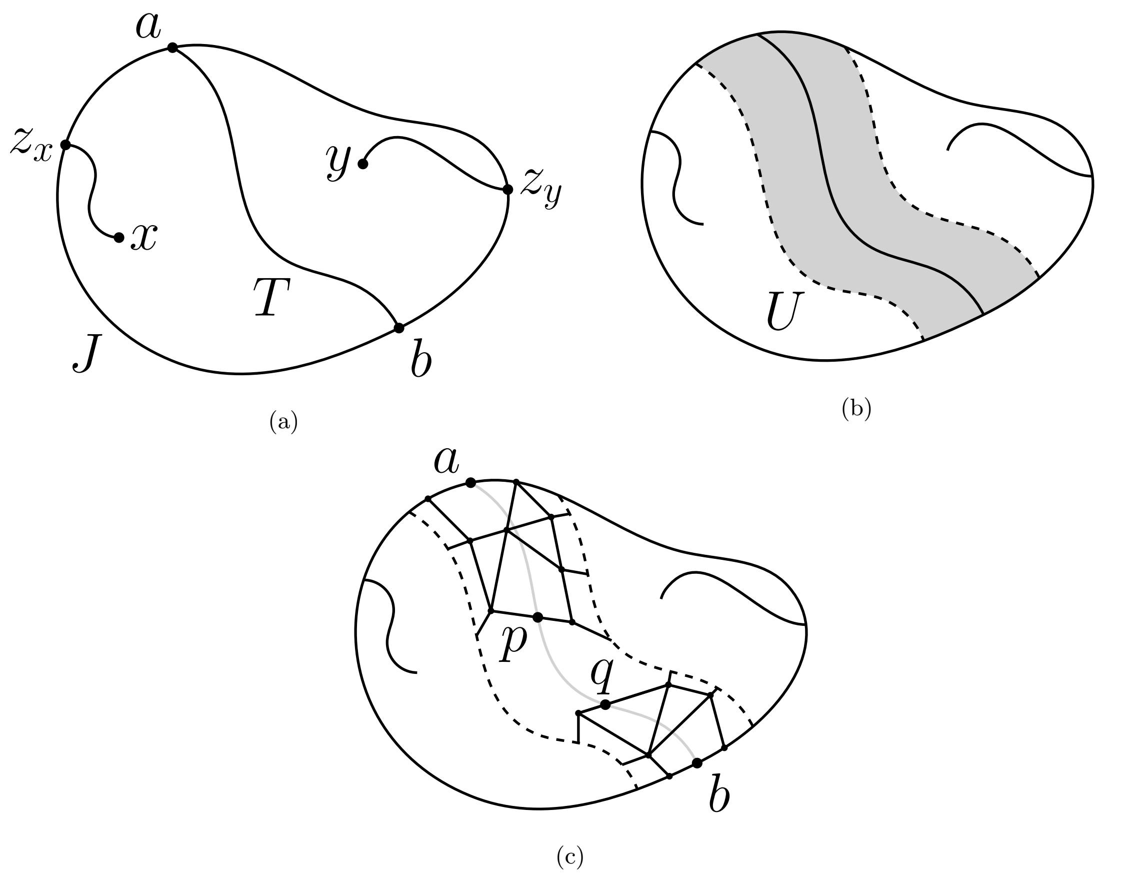

Figure 2.13

not contained within A_{k-1}; recalling the construction of the $\frac{1}{k}$-approximation of $\mathcal{T}$, it follows that the points p and q cannot be $\frac{1}{k}$-distant in A_{k-1}.

Unravelling the jargon, we find that there must exist some arc in A_{k-1} which connects p to q and has diameter less than $\frac{1}{k}$. However, such an arc must then lie in U, contradicting our assumption that a and b belong to distinct components of $U \cap A_{k-1}$, and completing the proof. $\square$

Lemma 2.4.20. *Let $\mathcal{T}$ denote some collection of finitely many J spanning arcs in $\mathrm{cl}_X\, Y$, with the property that every component of $\mathrm{cl}_X\, Y - \bigcup_{T \in \mathcal{T}} T$ has diameter strictly less than δ, for some $\delta > 0$.*

There exists some $\varepsilon > 0$ such that every domain of the ε-approximation of $\mathcal{T}$ also has diameter strictly less than δ.

Proof. We cover $\mathrm{cl}_X\, Y$ by some collection $\mathcal{B} = \{B_1, B_2, \ldots, B_n\}$ of finitely many closed balls of radius $\frac{\delta}{4}$, and define the set

$$I = \{(i,j) \,|\, B_i \cap B_j = \emptyset\}. \tag{2.15}$$

Given any $(i,j) \in I$, we know that $\mathcal{T}$ must separate points of B_i from points of B_j; otherwise, there would exist some complementary domain of $\bigcup_{T \in \mathcal{T}} T$ in $\mathrm{cl}_X\, Y$ of diameter at least δ. Lemma 2.4.19 therefore asserts that the ε_{ij}-approximation of $\mathcal{T}$ separates points of B_i from points of B_j, for some $\varepsilon_{ij} > 0$.

Now, let $\varepsilon = \min_{(i,j) \in I} \varepsilon_{ij}$; since I is a finite set, this ε is well defined. Moreover, inspection of Definition 2.4.18 shows that, for each $(i,j) \in I$, the ε_{ij}-approximation of $\mathcal{T}$ is a subspace of the ε-approximation of $\mathcal{T}$. Thus, every domain of the ε-approximation of $\mathcal{T}$ has diameter less than δ, as required. $\square$

2.4.5 Complementary Domains as 2-Cells

Lemma 2.4.20 is the final ingredient we need in order to prove that $\mathrm{cl}_X Y$ is a 2-cell. Indeed, for each $n \in \mathbb{N}$, let $\mathcal{T}_n$ denote some collection of finitely many J-spanning arcs which divide $\mathrm{cl}_X Y$ into pieces of diameter no greater than $\frac{1}{n}$, in the sense of Lemma 2.4.14. Using Lemma 2.4.20, we can produce from each such $\mathcal{T}_n$ an arc complex A_n, all domains of which also have diameter strictly less than $\frac{1}{n}$.

Via Lemma 2.4.6, we know that each arc complex A_n is isomorphic to some S^1-spanning arc complex B_n in the 2-cell D^2. We do not have any *a priori* guarantee that each domain of B_n is also bounded above by $\frac{1}{n}$, but no generality is lost in assuming this to be the case. (To see why, notice that we can reduce any domain of B_n of diameter at least $\frac{1}{n}$ to finitely many components of diameter strictly less than $\frac{1}{n}$ by, for example, repeated barycentric subdivisions. The arc complex that results is, again via Lemma 2.4.6, isomorphic to some J-spanning arc complex in $\mathrm{cl}_X Y$, with which we can replace our original arc complex A_n.)

Now, we can define subspaces

$$\begin{aligned} C &= \{x \in \mathrm{cl}_X Y \mid x \text{ is a vertex of } A_n \text{ for some } n \in \mathbb{N}\} \text{ and} \\ D &= \left\{x \in D^2 \,\middle|\, x \text{ is a vertex of } B_n \text{ for some } n \in \mathbb{N}\right\}, \end{aligned} \tag{2.16}$$

which are dense in their respective ambient spaces essentially by construction; further, the isomorphisms between the arc complexes $\{A_n\}_{n\in\mathbb{N}}$ and $\{B_n\}_{n\in\mathbb{N}}$ furnish us with a natural bijection $f\colon C \to D$. Recalling Lavrentieff's theorem (Lemma 2.4.2), we find that it will be sufficient for us to prove f uniformly bicontinuous in order to establish a homeomorphism $\mathrm{cl}_X Y \cong D^2$.

To this end, fix some $\varepsilon > 0$, and choose some $n \in \mathbb{N}$ such that $\frac{2}{n} < \varepsilon$. We notice first that the uniform continuity of f will follow from the existence of some $\delta > 0$ with the property that, if $d(x, y) < \delta$ for any pair of points $x, y \in C$, then x and y lie in the closures of some pair of (not necessarily distinct) adjacent domains of A_n. (Explicitly, suppose that we are given some such δ, and select a pair of points $x, y \in C$ for which $d(x, y) < \delta$. We then know that $f(x)$ and $f(y)$ belong to the closures of some adjacent pair of domains in B_n, and every domain of B_n has diameter less than $\frac{1}{n}$. It follows that $\|f(x) - f(y)\| \leq \frac{2}{n} < \varepsilon$, so that f is uniformly continuous.)

We shall argue that such a δ must exist by contradiction. Suppose that, for each $k \in \mathbb{N}$, there exists some pair of points $x_k, y_k \in C$ such that

- $d(x_k, y_k) < \frac{1}{k}$, but
- $f(x_k)$ and $f(y_k)$ do not lie in the closures of an adjacent pair of domains of B_n.

Appealing to compactness and passing to subsequences if necessary, we find that the sequences $(x_k)_{k=1}^{\infty}$ and $(y_k)_{k=1}^{\infty}$ have limits x and y respectively in C. Moreover, the arc complex A_n of course has only finitely many domains, so that an application of the pigeonhole principle allows us to lose no generality in assuming additionally that $(x_k)_{k=1}^{\infty}$ lies entirely in $\mathrm{cl}_X D_1$ and $(y_k)_{k=1}^{\infty}$ entirely in $\mathrm{cl}_X D_2$, for some pair of non-adjacent domains D_1 and D_2 of A_n.

Observe that

$$d(x, y) \leq d(x, x_k) + d(x_k, y_k) + d(y_k, y) \tag{2.17}$$

for each $k \in \mathbb{N}$. Since $x_k \to x$ and $y_k \to y$, we can of course arrange that $d(x, x_k) < \frac{\lambda}{3}$ and $d(y, y_k) < \frac{\lambda}{3}$ for any $\lambda > 0$, by choosing some sufficiently large value of k. If $\frac{1}{k} < \frac{\lambda}{3}$, then our definition of the sequences $(x_k)_{k=1}^{\infty}$ and $(y_k)_{k=1}^{\infty}$ also ensures that $d(x_k, y_k) < \frac{\lambda}{3}$, whence it follows that $x = y$. However, this is a contradiction: we chose the domains D_1 and D_2 to be non-adjacent, so that $\mathrm{cl}_X D_1 \cap \mathrm{cl}_X D_2 = \emptyset$. Thus, they cannot possibly have a common boundary point. The uniformly continuity of f is therefore established, while that of f^{-1} follows by exchanging the roles of $\mathrm{cl}_X Y$ and D^2 in the above reasoning.

All of this is summarised by the following result.

Lemma 2.4.21 (Schoenflies' theorem)**.** *Let Y denote some complementary domain of a Jordan curve J in a compact Zippin space X. Then, there exists some homeomorphism $f\colon \mathrm{cl}_X Y \to D^2$ such that $f(J) = S^1$.*

Proof. We did not show in the above that $f(J) = S^1$, but this follows immediately from our construction of the requisite arc complex isomorphisms in Lemma 2.4.6. □

2.5 Completing the Proof

So far, we have shown that if J is some Jordan curve in a compact Zippin space X, then $X - J$ consists of exactly two components A_1 and A_2, both with boundary J (Lemma 2.3.7), and we have deduced the existence of homeomorphisms $f_1\colon \operatorname{cl}_X A_1 \to D^2$ and $f_2\colon \operatorname{cl}_X A_2 \to D^2$ such that $f_1(J) = f_2(J) = S^1$ (Lemma 2.4.21). We are almost ready to conclude, using Lemma 2.2.2 to view $\operatorname{cl}_X A_1$ and $\operatorname{cl}_X A_2$ as two hemispheres, that X is homeomorphic to S^2, but there is one remaining subtlety which must be tackled.

It was crucial in Lemma 2.2.2 that we attached the pair of 2-cells along their boundary circles 'without twisting', but as it stands, we only know that f_1 and f_2 both map the Jordan curve J to the circle S^1; given a single point $x \in J$, we have no assurance that $f_1(x) = f_2(x)$. In other words, it is possible that f_1 and f_2 differ on S^1 by a non-trivial homeomorphism of the circle. Fortunately, however, a classical trick due to Alexander[4] [Ale23] gives us a way to circumvent the issue.

Lemma 2.5.1 (Alexander's lemma)**.** *Let $\varphi\colon S^1 \to S^1$ denote some homeomorphism. Then, φ can be extended to a homeomorphism $\widetilde{\varphi}\colon D^2 \to D^2$.*

Proof. With the exception of the origin, each point of D^2 of course permits a unique representation in the form rx, where $r \in [0, 1]$ and $x \in S^1$. Defining

$$\widetilde{\varphi}(rx) = rf(x) \tag{2.18}$$

yields the desired extension. □

Let us apply Alexander's lemma to the composition $f_1 \circ f_2^{-1}\colon S^1 \to S^1$, producing an extension $G\colon D^2 \to D^2$. The composition $G \circ f_2\colon \operatorname{cl}_X A_2 \to D^2$ is then a homeomorphism with the property that, given any point $x \in J$,

$$(G \circ f_2)(x) = \big(f_1 \circ f_2^{-1} \circ f_2\big)(x) = f_1(x), \tag{2.19}$$

and this is precisely what we need. Identifying $\operatorname{cl}_X A_1$ with one copy of D^2 via the homeomorphism f_1, and $\operatorname{cl}_X A_2$ with a second copy of D^2 via the homeomorphism $G \circ f_2$, we invoke Lemma 2.2.2 to conclude that every compact Zippin space is homeomorphic to S^2, completing our proof of Zippin's sphere characterization.

[4]We remark in passing that an analogous result holds in all finite dimensions, and can be proved by exactly the same methods. This fact sometimes goes by the moniker of *Alexander's trick*, but we avoid that terminology here; 'Alexander's trick' may also refer to a much more general result, stating that if two homeomorphisms of the n-cell D^n are isotopic on the boundary S^{n-1}, then they are in fact isotopic throughout D^n. We shall have no need of this more powerful machinery here.

Appendix A

Results from Continuum Theory

A.1 Overview

The theory of continua has a lengthy and storied history, with its genesis in the late 19^{th} century and Cantor's explorations [Can83] of perfect subspaces[1] of $\mathbb{R}^n$ satisfying an additional property[2] which reduces to connectedness in the compact setting. (Interestingly, it was also in [Can83] that Cantor first introduced the now-classical space bearing his name, which we will encounter in Definition A.6.4.) However, this would not evolve into the modern definition of continua (Definition A.2.1) until the notion of compactness was given its modern formulation by Alexandroff and Urysohn [AU24, Wil70], and that of connectedness by Lennes and Riesz [Len05, Rie06, Wil78].

From its infancy, continuum theory has facilitated research in adjacent fields of mathematics, especially in topological dynamics [Wes18], and spawned numerous problems–many of which remain open, with a representative cross-section given in [vMR90]. One example, counted by Bing as among the most interesting open problems in contemporary topology [Bin69], is the *plane fixed point problem.* Its statement seems relatively innocuous: if $C \subset \mathbb{R}^2$ denotes some continuum such that $\mathbb{R}^2 - C$ is connected, then does every continuous function $C \to C$ necessarily have a fixed point? Nevertheless, the question remains unanswered, although Bellamy has provided an example [Bel79] of a potentially non-planar continuum without this property, while Fokkink et al. answer the question in the affirmative for a restricted class of continua and functions in [FMOT08].

Here, we shall restrict ourselves to a narrow transect of continuum theory, developing only what is necessary for our treatment of Zippin's sphere characterization in Chapter 2. With arbitrary continua, we shall prove what we need in order to characterize arcs and the circle (Lemmas A.4.5 and A.4.6). With this complete, we shall add local connectedness to our hypotheses and begin to study the so-called *Peano continua.*

After proving in Section A.5 that Peano continua enjoy far stronger connectivity properties than arbitrary continua, we have two main objectives in our exploration of Peano continua. The first is a complete characterization of Peano continua without explicit mention of local connectedness (Theorem A.6.8), while the second (Theorem A.10.3) is a description of *cyclic* Peano continua (Definition A.10.1) which is indispensable in our proof of Zippin's sphere characterization. Along the way, we shall also encounter a decomposition of Peano continua explicitly originating in Whyburn's work [Why27a], but also implicit in the work of Ayres [Ayr27].

A.2 Basic Definitions

A *continuum*, from the perspective of contemporary general topology, is always some form of compact, connected space, although there is some variation in the literature as to whether any additional

[1] That is, closed subspaces with no isolated points.
[2] In modern language, this additional property is that of *well chainedness.*

hypotheses should be included in this definition.

Here, we follow Nadler's convention [Nad92] of always requiring our continua to be metrizable, and we also consider Whyburn's notion [Why45] of *generalized continua*. We make these definitions explicit below.

Definition A.2.1. *A* continuum *is a non-empty topological space which is compact, connected, and metrizable; a* Peano continuum *is a continuum which is also locally connected.*

A generalized continuum *is any non-empty topological space which is locally compact, connected, and metrizable; a* generalized Peano continuum *is a generalized continuum which is also locally connected.*

Any continuum or generalized continuum which contains more than one point is referred to as non-degenerate*.*

It is worthwhile to note that, by this definition, continua comprise a strict subclass of the compact, connected spaces: as a straightforward example, endowing a two-point set with the trivial topology yields a space which is of course compact and connected, but fails to be Hausdorff. However, there are also many examples of compact, connected Hausdorff spaces which Definition A.2.1 excludes from being continua, of which we give just one example.

Example A.2.2 (The cofinite topology on $\mathbb{R}$)**.** The *cofinite topology on* $\mathbb{R}$ is defined by declaring that a subset $U \subset \mathbb{R}$ is open if and only if either

- U is empty, or
- The complement $\mathbb{R} - U$ is finite.

First, we shall prove that the space $\mathbb{R}_{\text{cof}}$ produced by equipping the set $\mathbb{R}$ with this topology is both compact and connected; then, we argue that the space $\mathbb{R}_{\text{cof}}$ is not first countable, and *a fortiori* not metrizable.

For compactness, let us denote by $\mathcal{U}$ any open cover of $\mathbb{R}_{\text{cof}}$. If $\mathbb{R} \in \mathcal{U}$, then of course $\mathcal{U}$ contains a finite subcover, so let us suppose this not to be the case, and distinguish some arbitrary element $U_0 \in \mathcal{U}$. Our definition of the cofinite topology tells us that the complement $\mathcal{R} - U_0$ consists of finitely many points, say $\{x_1, x_2, \ldots, x_n\}$. Now, for each $i \in \{1, 2, \ldots, n\}$, there must of course exist some $U_i \in \mathcal{U}$ such that $x_i \in U_i$, and it follows immediately that the collection $\{U_0, U_1, U_2, \ldots, U_n\}$ also covers $\mathbb{R}_{\text{cof}}$.

As regards connectedness, it will suffice for us to prove that the only subspaces of $\mathbb{R}_{\text{cof}}$ which are both closed and open are the empty set and $\mathbb{R}_{\text{cof}}$ itself. Indeed, let $A \subset \mathbb{R}_{\text{cof}}$ denote some subspace which is both open and closed, and suppose that $A \notin \{\emptyset, \mathbb{R}_{\text{cof}}\}$; then, our definition of the cofinite topology forces us to conclude that both $\mathbb{R} - A$ and A contain only finitely many points–but since $\mathbb{R}$ contains infinitely many points, this is impossible. It follows, therefore, that $\mathbb{R}_{\text{cof}}$ is connected.

To establish that the space $\mathbb{R}_{\text{cof}}$ is not first countable (and therefore not metrizable), despite being both compact and connected, we shall argue by contradiction. To this end, suppose that there were to exist some countable neighbourhood basis for the point $0 \in \mathbb{R}$ in this topology, say

$$\mathcal{B} = \{\mathbb{R} - A_1, \mathbb{R} - A_2, \ldots\}. \tag{A.1}$$

Then, the union $C = \{0\} \cup \bigcup_{n \in \mathbb{N}} A_n$ is countable, so that there must exist some point $x \in \mathbb{R} - C$. The complement $\mathbb{R} - \{x\}$ is, by definition, a neighbourhood of 0, but by construction this neighbourhood can contain no element of $\mathcal{B}$. □

Before proceeding, we also give an example of a compact, connected, non-metrizable (and non-Hausdorff) space which is of significant interest elsewhere in mathematics, especially in algebraic geometry–namely, the *spectrum of* $\mathbb{Z}$ (or, more generally, of any principal ideal domain).

Example A.2.3 (The spectrum of $\mathbb{Z}$)**.** Let $\operatorname{Spec} \mathbb{Z}$ denote the set of prime ideals of $\mathbb{Z}$, and declare a subspace $C \subset \operatorname{Spec} \mathbb{Z}$ to be closed if and only if

$$C = \{P \in \operatorname{Spec} \mathbb{Z} \mid I \subset P\} \tag{A.2}$$

for some ideal $I \subset \mathbb{Z}$. This is the so-called *Zariski topology* on $\operatorname{Spec}\mathbb{Z}$, and using this to topologize $\operatorname{Spec}\mathbb{Z}$ produces a space known as the *spectrum of* $\mathbb{Z}$, again denoted by $\operatorname{Spec}\mathbb{Z}$.

To see that $\operatorname{Spec}\mathbb{Z}$ is compact, recall that every ideal of $\mathbb{Z}$ is principal. Thus, given some ideal $(x) \subset \mathbb{Z}$ and some prime ideal $(p) \in \operatorname{Spec}\mathbb{Z}$, we know that if $(x) \subset (p)$, then p is a factor of x. Since an integer has only finitely many prime factors, we find that every closed subspace of $\operatorname{Spec}\mathbb{Z}$ is finite. Thus, given some arbitrary open cover $\mathcal{U}$ of $\operatorname{Spec}\mathbb{Z}$, we can distinguish some element $U \in \mathcal{U}$, safe in the knowledge that the complement $\operatorname{Spec}\mathbb{Z} - U$ is at most finite. This complement can of course be covered by at most finitely many elements of $\mathcal{U}$, establishing the compactness of $\operatorname{Spec}\mathbb{Z}$.

Moreover, the definition (A.2) ensures that no closed subspace of $\operatorname{Spec}\mathbb{Z}$ can contain the trivial ideal (0); thus, $\operatorname{Spec}\mathbb{Z}$ permits no disjoint open subspaces, and must therefore be connected, but cannot be Hausdorff. □

Given that our primary interest in the main body of this work lies with the 2-sphere, however, we do not lose much in insisting that our continua be metrizable–and in doing so, we are able to simplify certain proofs, and make possible others which would not hold without this assumption.

A useful means of constructing new continua from some prescribed collection of nested continua is to take their intersection, which we detail in the following lemma. In order to apply this result to our proof of Lemma A.3.3, we are forced to couch this lemma in the language of nets, which we do not formally introduce, in order to avoid a lengthy detour into the Moore-Smith theory[3] [MS22] of 'generalized convergence' which is lucidly described in [Cla16].

The fact that we here work exclusively in first countable spaces, however, allows the reader unfamiliar with nets to replace arbitrary directed sets by the natural numbers, and nets by sequences, without changing the core idea of the following proof. However, when we prove Lemma A.3.3, we must consider the intersection of a collection of subcontinua of a given continuum, and this collection is not, in general, indexed by the natural numbers; nevertheless, it is certainly directed by reverse set-theoretic inclusion.

Lemma A.2.4. *Consider some collection $\{X_i\}_{i\in I}$ of continua, directed by reverse set-theoretic inclusion. (That is, impose some preorder $\preceq$ on $\{X_i\}_{i\in I}$ by declaring that $X_i \preceq X_j$ if and only if $X_j \subset X_i$.)*

Then, the intersection $\bigcap_{i\in I} X_i$ is itself a continuum.

Proof. Let us arbitrarily select some index $j \in I$. Observing that $\bigcap_{i\in I} X_i = \bigcap_{i\in I}(X_i \cap X_j)$, we see that no generality is lost in assuming that $X_i \subset X_j$ for each $i \in I$. Moreover, we notice that $\bigcap_{i\in I} X_i$ is a closed subspace of X_j, and therefore must be compact.

As for connectedness, we shall argue by contraposition, showing that if $\bigcap_{i\in I} X_i$ is disconnected, then so too must one of the spaces $\{X_i\}_{i\in I}$ be. Suppose, then, that $\bigcap_{i\in I} X_i = A \cup B$, where the disjoint subspaces A and B are non-empty and closed. Then, since the intersection $\bigcap_{i\in I} X_i$ is compact Hausdorff, we know that A and B are both compact. Knowing that A and B are disjoint compact subspaces of the compact Hausdorff (hence normal) space X_j, we can find some pair U and V of disjoint neighbourhoods of A and B respectively in X_j.

Now, if all of the spaces $\{X_i\}_{i\in I}$ are to be connected, then for each $i \in I$, there must exist some point $x_i \in X_i - (U \cap V)$: otherwise, $X_i = (U \cap X_i) \mid (V \cap X_i)$ would be a separation of a connected space. We assemble the points $\{x_i\}_{i\in I}$ into a net $\Phi\colon I \to X_j$ defined by $\Phi(i) = x_i$. (Here, we have used our above assumption that $X_i \subset X_j$ for each $i \in I$.)

The compactness of X_j ensures that this net has some accumulation point $x \in X_j$. If $W \subset X_j$ denotes any neighbourhood of this accumulation point, then for each $i \in I$, there must by definition exist some $i' \in I$ such that $X_{i'} \subset X_i$ and $x_{i'} \in W$. In other words, we have shown that every neighbourhood of the accumulation point x intersects all of the continua $\{X_i\}_{i\in I}$, so that

$$x \in \bigcap_{i\in I} \operatorname{cl}_{X_j} X_i = \bigcap_{i\in I} X_i. \tag{A.3}$$

[3]Somewhat confusingly, this 'Moore' is E.H. Moore, not the R.L Moore of the Moore's theorem discussed in Chapter 1. Muddying the waters further, R.L Moore was at one point supervised by E.H. Moore. [Wil76]

This in turn implies that the union $U \cup V$ is a neighbourhood of the accumulation point x–but this is a contradiction. By construction, the net Φ never takes values in $U \cup V$, so *a fortiori* it is impossible for Φ to frequently be in $U \cup V$. □

A.3 Cut Points and Separators

A great deal of insight into the structure of a continuum can be gleaned by studying subspaces which separate the continuum on their removal: indeed, the entirety of Chapter 2 was devoted to a characterization of the 2-sphere in terms of the (dis)connectedness of certain complementary domains. We now investigate these ideas in greater generality, first introducing some terminology.

Definition A.3.1. *Let X denote some topological space. If $A \subset X$ denotes some subspace such that $X - A$ is disconnected, then we say that A is a* separator *of X, or that A* separates *X.*

If a singleton $\{x\} \subset X$ is a separator of X, then we call the point x a cut point *of X. Dually, any point of X which is not a cut point is termed a* non-cut point *of X.*

The following reasonable-sounding result will be of significant utility in our study of continua. To indicate why we need to include in Lemma A.3.2 the hypothesis that the separator C be connected, consider the unions $A \cup C$ for the planar continua sketched in Figure A.1; one is manifestly itself connected, while the other is just as clearly disconnected.

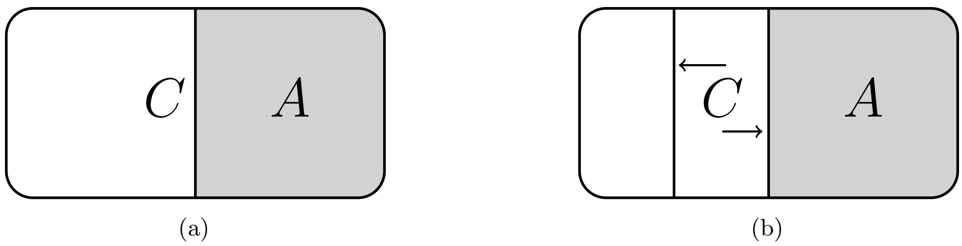

Figure A.1

Lemma A.3.2. *Denote by C some connected separator of a continuum X. Then, if $X - C = A \mid B$ is some separation, the unions $A \cup C$ and $B \cup C$ are both continua.*

Proof. To prove that $A \cup C$ is connected, we argue by contradiction. If $A \cup C$ is disconnected, then we can produce some separation $A \cup C = U \mid V$, and the connectedness of C allows us to assume without loss of generality that $C \subset V$. It follows that $U \subset A$, for were there to exist a point $u \in U - A$, this point would lie in C, and therefore in V. This, however, implies that $X = U \mid (B \cup V)$ is a separation of X, which is a contradiction. Thus, $A \cup C$ is connected, and analogous reasoning shows us that $B \cup C$ must be as well.

To complete the proof, we can see that $A \cup C$ and $B \cup C$ are closed, and therefore compact, by observing that the complements $X - (A \cup C) = B$ and $X - (B \cup C) = A$ are open. □

Another powerful result which may at first seem innocuous concerns the existence of non-cut points: namely, that every continuum contains at least two.[4] We can gain some intuition for why this should be the case by considering the closed unit interval $[0, 1]$. If we delete, say, the origin from $[0, 1]$, then we arrive at a space which has only a single non-cut point, but we have sacrificed compactness in doing so; on the other hand, if we attempt to force the non-cut points to coalesce by passing to the quotient space S^1, then we can retain compactness, but produce a space with no cut points at all.

Lemma A.3.3. *Let p denote some cut point of a continuum X, and $X - \{p\} = A \mid B$ some separation. Then, each of A and B contain a non-cut point of X.*

[4] In passing, we remark that this fact was first proved in [Moo20] by R.L. Moore, originator of Theorem 1.1.2.

Proof. Towards a contradiction, suppose that every point of A is a cut point of X, so that for each $a \in A$ we can find some separation $X - \{a\} = U(a) \mid V(a)$. Notice that it is impossible for both $U(a)$ and $V(a)$ to intersect $B \cup \{p\}$, for otherwise we would have a separation

$$B \cup \{p\} = \big(U(a) \cap (B \cup \{p\})\big) \mid \big(V(a) \cap (B \cup \{p\})\big) \tag{A.4}$$

of a space which is, according to Lemma A.3.2, connected. Thus, one of $U(a)$ and $V(a)$ must be contained within A; without loss of generality, we assume that $U(a) \subset A$.

Now, the collection $\{U(a) \cup \{a\}\}_{a \in A}$ is a collection of continua, all contained within A, which we can direct by set-theoretic inclusion. Using Lemma A.2.4, we see that the intersection

$$C = \bigcap_{a \in A} (U(a) \cup \{a\}) \subset A \tag{A.5}$$

is itself a continuum.

Selecting any point $c \in C$, consider the associated separation $X - \{c\} = U(c) \mid V(c)$, where $U(c) \subset A$. Choosing some point $d \in U(c)$, observe that c cannot be a point of $U(d)$, for otherwise both $U(d)$ and $V(d)$ would intersect $V(c) \cup \{c\}$, inducing a separation of a continuum as in (A.4). We have arrive at a contradiction: on the one hand, we require that $c \in \bigcap_{a \in A}(U(a) \mid \{a\})$, but on the other, we have just demonstrated that $c \notin U(d) \cup \{d\}$ for the point $d \in U(c) \subset A$.

Applying the same reasoning to points of B completes the proof. □

Corollary A.3.4. *Every non-degenerate continuum has at least two non-cut points.* □

One heuristic for visualizing cut points of a continuum is to consider them as where that continuum is 'pinched' to a point. The following lemma makes this slightly more precise: if we take a non-cut point of a continuum X, then we can always find some neighbourhood of that point which fails to separate X.

Lemma A.3.5. *If x denotes some non-cut point of a continuum X, then there exists some neighbourhood $U \subsetneq X$ of x such that the complement $X - U$ is connected.*

Proof. Arguing by contraposition, let us assume that every neighbourhood of some point $x \in X$ separates X. Then, arbitrarily choosing some metric which induces the topology of X, we know in particular that the open balls $\{B(x; n^{-1})\}_{n \in \mathbb{N}}$ all separate X. Let

$$X - B\left(x; \frac{1}{n}\right) = Y_n \mid Z_n \tag{A.6}$$

denote some separation for each $n \in \mathbb{N}$, and notice that, since

$$X - B\left(x; \frac{1}{n}\right) \subset X - B\left(x; \frac{1}{n+1}\right) \text{ for each } n \in \mathbb{N}, \tag{A.7}$$

we lose no generality if we also assume that $Y_n \subset Y_{n+1}$ and $Z_n \subset Z_{n+1}$ for every $n \in \mathbb{N}$.

We know that

$$\begin{aligned} X - \{x\} &= X - \left(\bigcap_{n \in \mathbb{N}} B\left(x; \frac{1}{n}\right)\right) \\ &= \bigcup_{n \in \mathbb{N}} \left(X - \left(x; \frac{1}{n}\right)\right) \\ &= \left(\bigcup_{n \in \mathbb{N}} Y_n\right) \cup \left(\bigcup_{n \in \mathbb{N}} Z_n\right). \end{aligned} \tag{A.8}$$

However, the unions $\bigcup_{n \in \mathbb{N}} Y_n$ and $\bigcup_{n \in \mathbb{N}} Z_n$ are both open and disjoint, since for each $n \in \mathbb{N}$, both Y_n and Z_n are open and disjoint, with $Y_1 \subset Y_2 \subset Y_3 \subset \cdots$ and $Z_1 \subset Z_2 \subset Z_3 \subset \cdots$. It follows that x is a cut point of X, which completes the proof. □

A.4 Characterizing the Closed Unit Interval

In this section, we shall identify the closed unit interval $[0, 1]$ as being, up to homeomorphism, the only continuum with precisely two non-cut points. This result is not only elegant and intuitively plausible, but also of great use, giving us a tool for identifying arcs in other continua without explicitly constructing a homeomorphism.

This result was originally proved by the same Moore who proved Theorem 1.1.2, and–much as we found in Chapter 1–Moore's original exposition is at best challenging for many modern readers to follow. Our proof here follows [Wil70], and is enabled by Whyburn's notion [Why28] of a *separation order*. Roughly, we can summarise the separation order as stating that if a point $x \in X$ separates a point $y \in X$ from a, then we can view x as being, in some sense, 'closer' to a than y is.

A result of Cantor's [Can95], also of interest in set theory and in model theory [Roi90, Mar00], will allow us to identify a subspace of X with a subspace of $[0, 1]$ order-theoretically. After constructing a natural extension of this order isomorphism to the entirety of X and $[0, 1]$, we shall see that the separation order on X in fact induces the topology of X; thus, we can promote this order isomorphism to a topological homeomorphism, and in doing so we arrive at the desired characterization.

Having sketched how we are to proceed, we now make the arguments involved explicit. We first use the classical *back-and-forth argument* to prove Cantor's order-theoretic result, which is of a slightly different flavour to the topological results that follow it.

Lemma A.4.1 (Cantor)**.** *Let $(A, \leq_A)$ and $(B, \leq_B)$ denote some pair of countable total orders with neither greatest nor least elements. Suppose additionally that these orders are* dense*, in the sense that given any pair of points $a_1, a_2 \in A$ and $b_1, b_2 \in B$, we can find $a_3 \in A$ and $b_3 \in B$ such that $a_1 <_A a_3 <_A a_2$ and $b_1 <_B b_3 <_B b_2$.*

Then, $(A, \leq_A)$ and $(B, \leq_B)$ are order isomorphic to one another.

Proof. Arbitrarily choose enumerations $A = \{a_1, a_2, \ldots\}$ and $B = \{b_1, b_2, \ldots\}$, assuming without loss of generality that whenever $i \neq j$, we have both that $a_i \neq a_j$ and that $b_i \neq b_j$. Further, define the finite subsets $A_n = \{a_1, a_2, \ldots, a_n\}$ and $B_n = \{b_1, b_2, \ldots, b_n\}$ for each $n \in \mathbb{N}$.

To define an order isomorphism $f\colon A \to B$, we define the restrictions $f|_{A_n}$ and $f^{-1}|_{B_n}$ inductively. For the basis case of this induction, we simply set $f(a_1) = b_1$ and $f^{-1}(b_1) = a_1$, so suppose that we have successfully defined the restrictions $f|_{A_{n-1}}\colon A_{n-1} \to B$ and $f^{-1}|_{B_{n-1}}$ for some $n > 1$. We wish to define $f(a_n)$ and $f^{-1}(b_n)$ in terms of these restrictions.

Our requirement that f be an order isomorphism means that, given the restrictions $f|_{A_{n-1}}$ and $f^{-1}|_{B_{n-1}}$, we have defined f on a domain $C = A_{n-1} \cup f^{-1}(B_{n-1})$. Now, if $a_n \in C$, then there is nothing for us to do, so suppose that $a_n \notin C$. There are then three possibilities for where the point a_n may lie with respect to C in the order $\leq_A$, which we shall exhaust. Namely,

- a_n may be strictly greater than every element of C,
- a_n may be strictly less than every element of C, or
- a_n may have both an immediate successor and an immediate predecessor in C. (Notice that this successor and predecessor must exist, since C is finite.)

In the first case, our assumption that $(B, \leq_B)$ has no greatest element allows us to find some point $b \in B$ such that $f(\max C) <_B b$. Knowing that f is an order isomorphism, we deduce that $b \notin f(C)$, and so we define $f(a_n) = b$.

We can deal with the second case in essentially the same fashion. There exists a point $b' \in B$ such that $b' <_B f(\min C)$, because $(B, \leq_B)$ has no least element, and since f is an order isomorphism, we know that $b' \notin f(C)$. Thus, we define $f(a_n) = b'$.

As for the third case, let s and p denote respectively the immediate successor and predecessor in C of the point a_n. Using the density of $(B, \leq_B)$ together with the fact that f preserves order allows us to find a point $b'' \in B$ such that $f(p) < b'' < f(s)$, and by the same logic as above, we know that $b'' \notin f(C)$, allowing us to define $f(a_n) = b''$.

In order to define $f^{-1}(b_n)$, we need only exchange the roles of f and f^{-1} and those of $(A, \leq_A)$ and $(B, \leq_B)$ in the above. □

Corollary A.4.2. *Every countable total order with neither a greatest nor a least element which is dense in the sense of Lemma A.4.1 is order isomorphic to $\mathbb{Q} \cap \,]0,1[$ with its standard order.* □

Now, let us formally define the *separation order*, and show that–at least in the setting we are interested in here–this order induces the right topology.

Definition and Lemma A.4.3. *Let X denote some continuum with precisely two non-cut points $a, b \in X$. Define a relation $\preceq$ on X by declaring that*

$$x \preceq y \text{ if and only if either } x = y, \text{ or } x \text{ separates } a \text{ from } y \text{ in } X.$$

This relation is a total order on X, which we call the separation order*.*

Proof. For any point $x \in X - \{a, b\}$, we shall use the notation

$$X - \{x\} = A(x) \mid B(x) \tag{A.9}$$

to refer to a separation for which $a \in A(x)$ and $b \in B(x)$.

We start by proving that $\preceq$ is antisymmetric. Towards a contradiction, suppose that there exist distinct points $x, y \in X$ such that both $x \preceq y$ and $y \preceq x$. Then, we would have separations $X - \{x\} = A(x) \mid B(x)$ and $X - \{y\} = A(y) \mid B(y)$ such that $y \in B(x)$ and $x \in B(y)$. However, we know (Lemma A.3.2) that both $B(x) \cup \{x\}$ and $B(y) \cup \{y\}$ are connected subspaces of X which contain the point b, so that we have inclusions

$$B(x) \cup \{x\} \subset B(y) \text{ and } B(y) \cup \{y\} \subset B(x). \tag{A.10}$$

Taken together, these inclusions imply that

$$y \in B(x) \cup \{x\} \subset B(y), \tag{A.11}$$

and this is a contradiction: by definition, $y \notin B(y)$. The antisymmetry of $\preceq$ is therefore established.

We can prove that $\preceq$ is transitive by similar means. Given points $x, y, z \in X$ such that $x \preceq y$ and $y \preceq z$, we know that $B(z) \cup \{z\} \subset B(y)$, and likewise that $B(y) \cup \{y\} \subset B(x)$. It follows that $z \in B(x)$, so that $x \preceq z$, as needed.

Finally, we turn towards the question of whether or not $\preceq$ is connex–that is, whether or not, for each pair of points $x, y \in X$, $x \preceq y$ or $y \preceq x$. Selecting some pair of distinct points $x, y \in X$, consider the separation $X - \{x\} = A(x) \mid B(x)$. If $y \in B(x)$, then by definition we must have that $x \preceq y$, so suppose instead that $y \in A(x)$. Then, we know that $B(x) \cup \{x\} \subset B(y)$, and in particular that $x \in B(y)$, and we conclude that $\preceq$ is connex. Putting everything together, we deduce that $\preceq$ is a total order on X. □

Lemma A.4.4. *Let X denote some continuum with precisely two non-cut points $a, b \in X$. The order topology induced on X by the separation order $\preceq$ is precisely the topology of X.*

Proof. Recall that the order topology induced on X by $\preceq$ is generated from a subbasis of rays of the form

$$L_x = \{y \in X \mid y \prec x\} \text{ and } U_x = \{y \in X \mid x \prec y\}, \tag{A.12}$$

for points $x \in X$.

Consider some point $x \in X - \{a, b\}$, and retain the notation of (A.9). If $y \in X - \{x\}$, then by definition we know that $y \in A(x)$ if and only if either $y = a$, or if y separates x from a–but this is precisely the definition of the ray L_x. Similarly, we find that $B(x) = U_x$. Knowing that $A(x)$ and $B(x)$ are by open in X by definition, and observing also that $L_a = \emptyset$, $U_a = X - \{a\}$, $L_b = X - \{b\}$ and $U_b = \emptyset$, we find that the topology of X refines the separation order topology: every subbasis element for the latter is open in the former.

Conversely, consider some arbitrary open subspace $U \subset X$, end select a point $x \in U - \{a, b\}$. We shall argue by contradiction that there exist points $p, q \in X$ such that

$$x \in \,]p, q[\, = \{y \in X \mid p \prec y \prec q\} \subset U. \tag{A.13}$$

To this end, suppose that there exists no such interval, so that $]p, q[\cap (X - U) \neq \emptyset$ for all distinct points $p, q \in X$. *A fortiori*, we find that the closed interval

$$[p, q] = \{y \in X \mid p \preceq y \preceq q\} \tag{A.14}$$

also intersects $X - U$, yielding a collection

$$\mathcal{C} = \{[p, q] \cap (X - U) \mid p, q \in X \text{ and } p \prec x \prec q\} \tag{A.15}$$

of non-empty subspaces of X. Knowing as we now do that the topology of X is a refinement of the separation order topology, we also find that every member of $\mathcal{C}$ is closed in X.

Now, notice that, given a pair of closed intervals $[p_1, q_1]$ and $[p_2, q_2]$ such that $p_1 \prec q_1$ and $p_2 \prec q_2$,

$$[p_1, q_1] \cap [p_2, q_2] = [\max\{p_1, p_2\}, \min\{q_1, q_2\}]. \tag{A.16}$$

from which it follows that $\mathcal{C}$ is in fact closed under finite intersections. Knowing that X is compact, we find that there exists at least one point of X common to all elements of $\mathcal{C}$, so that

$$\bigcap_{C \in \mathcal{C}} C = (X - U) \cap \bigcap \{[p, q] \mid p, q \in X \text{ and } p \prec x \prec q\} \neq \emptyset. \tag{A.17}$$

However, we also know that

$$\bigcap \{[p, q] \mid p, q \in X \text{ and } p \prec x \prec q\} = \{x\} \subset U, \tag{A.18}$$

and the combination of (A.17) with (A.18) produces a contradiction.

We have successfully shown that if $U \subset X$ is open and $x \in U - \{a, b\}$, then U contains an interval, open with respect to the separation order on X, which contains the point x. If we can show similar results hold when $a \in U$ or $b \in U$, then we will be able to conclude that the separation order topology also refines the topology of X–but the reasoning for this is essentially identical to the above. (Explicitly, we repeat the same reasoning, replacing the open intervals $]p, q[$ with rays L_q or with rays U_p.) □

From here, all that needs to be done in order to complete our characterization of $[0, 1]$ is to apply Lemma A.4.1 to an appropriate subspace, and then extend the resulting order isomorphism to a homeomorphism.

Lemma A.4.5. *For any continuum X, the following are equivalent:*

- *X has precisely two non-cut points, and*
- *X is homeomorphic to the closed unit interval $[0, 1]$.*

Proof. One implication is trivial, since of course the closed unit interval has exactly two non-cut points: to wit, 0 and 1.

Thus, consider some continuum X with precisely two non-cut points $a, b \in X$. Throughout, we shall denote by $X - \{x\} = A(x) \mid B(x)$ a separation such that $a \in A(x)$ and $b \in B(x)$, for any point $x \in X - \{a, b\}$.

The space X is compact and metrizable, and therefore separable, allowing us to find some countable dense subspace $P \subset X$, where we lose no generality in assuming that $a, b \notin P$. Letting $\preceq$ denote the separation order on X, we propose that $(P, \preceq)$ satisfies the hypotheses of Lemma A.4.1.

That $(P, \preceq)$ has neither a greatest nor a least element is immediate: the greatest and least elements of $(X, \preceq)$ are precisely the points a and b, which we have explicitly excluded from P. Moreover, given points $p_1, p_2 \in P$ such that $p_1 \prec p_2$, the density of P in X ensures that there exists some point p_3 common to P and the separation-order-open interval $]p_1, p_2[$; in other words, $p_1 \prec p_3 \prec p_2$, so that the second condition of Lemma A.4.1 is also satisfied.

Invoking Lemma A.4.1, we arrive at some order isomorphism $f\colon (P, \preceq) \to \mathbb{Q} \cap]0,1[$. We wish to extend this to an order isomorphism $\widetilde{f}\colon (X, \preceq) \to [0,1]$, which, in light of Lemma A.4.4, will be the homeomorphism we seek.

Let us define $\widetilde{f}(a) = 0$ and $\widetilde{f}(b) = 1$. Now, given some point $x \in X - \{a, b\}$, we know that

$$X - \{x\} = [a, x[\cup]x, b], \tag{A.19}$$

where both intervals are defined in terms of the separation order $\preceq$. Using the fact that f is an order isomorphism, we deduce that the pair $(f([a,x[\cap P), f(]x,b]))$ comprises a Dedekind cut of $\mathbb{Q} \cap]0,1[$, uniquely determining a point of $[0,1]$ which we choose to be $\widetilde{f}(p)$.

To demonstrate that this extension $\widetilde{f}$ of f is an order isomorphism, notice first that of course $\widetilde{f}(a)$ and $\widetilde{f}(b)$ are, respectively, the greatest and least elements of $[0,1]$. Given some pair of points $x, y \in X - \{a,b\}$ such that $x < y$, we also know that $[a, x[\subset [a, y[$. It therefore follows that $f([a,x[\cap P) \subset f([a,y[\cap P)$, so that $\widetilde{f}(x) < \widetilde{f}(y)$ by the standard ordering of Dedekind cuts. (See, for example, [Pug10].) □

A.4.1 Characterizing the Circle

It is not hard to promote Lemma A.4.5's characterization of $[0,1]$ to a characterization of S^1. Indeed, our strategy is essentially that used in our proof of Zippin's sphere characterization writ small: we show that any continuum satisfying the right hypotheses is homeomorphic to a pair of unit intervals attached at their end points, which is of course a homeomorph of S^1.

Lemma A.4.6. *For any continuum X, the following are equivalent:*

- *X is separated by every non-degenerate pair of its points, and*
- *X is homeomorphic to the circle S^1.*

Proof. One implication is trivial, so let us suppose that every non-degenerate pair of points separates the continuum X.

Given any pair of distinct points $x, y \in X$, by hypothesis we can produce some separation $X - \{x,y\} = U \mid V$, and Lemma A.3.2 asserts that both $U' = U \cup \{x,y\}$ and $V' = V \cup \{x,y\}$ are continua. If we can prove there to exist homeomorphisms $U' \cong [0,1]$ and $V' \cong [0,1]$, both mapping $x \mapsto 0$ and $y \mapsto 1$, then we will be able to deduce that $X \cong S^1$ by invoking Lemma 2.2.2, and we shall achieve this by availing ourselves of Lemma A.4.5.

For the sake of a contradiction, let us suppose that $U' \ncong [0,1]$; then, Lemma A.4.5 allows us to find some point $u \in U$ such that $U' - \{u\}$ is connected. At this point, we distinguish two possibilities–V' may be an arc with end points a and b, or it may not be–and derive a contradiction in each case separately.

If V' is indeed an arc with end points a and b, then we can arbitrarily select a point $v \in V$ corresponding to some separation $V' - \{v\} = V_a \mid V_b$, where

- $a \in V_a$ and $b \in V_b$, and
- Both V_a and V_b are connected, each being a homeomorph of a half-open interval.

Then,

$$X - \{u, v\} = V_a \cup (U' - \{u\}) \cup V_b \tag{A.20}$$

must be connected, because $a \in V_a \cap (U' - \{u\})$ and $b \in V_b \cap (U' - \{u\})$, which is the desired contradiction.

If, instead, V' is not an arc with end points a and b, then Lemma A.4.5 enables us to choose some point $v \in V'$ such that $V' - \{v\}$ is connected. Then,

$$X - \{u, v\} = (U' - \{u\}) \cup (V' - \{v\}) \tag{A.21}$$

is connected, because $\{a, b\} \subset (U' - \{u\}) \cap (V' - \{v\})$, which again supplies us with our contradiction. □

A.5 Local Arc Connectedness of Peano Continua

There are numerous examples of continua which fail to be path connected; one such space, the topologists' sine curve, is briefly mentioned in Section 1.4.3. However, introducing the hypothesis of local connectedness and passing to the study of Peano continua eliminates this pathology in a striking fashion. Peano continua are not merely path connected, but rather arc connected–and even locally so.

It is worth pointing out that, in general, path connectedness is a strictly weaker notion than arc connectedness, as we can see by considering the spectrum of $\mathbb{Z}$.

Example A.5.1 (Spec $\mathbb{Z}$ is path connected, but not arc connected)**.** Recall our definition of the topological space Spec $\mathbb{Z}$ from Example A.2.3.

In that example, we showed that the trivial ideal $(0) \in \operatorname{Spec}\mathbb{Z}$ is common to every open subspace of Spec $\mathbb{Z}$, so that this point is *generic* in the sense that $\operatorname{cl}_{\operatorname{Spec}\mathbb{Z}}\{(0)\} = \operatorname{Spec}\mathbb{Z}$. Now, given any pair of points $P, Q \in \operatorname{Spec}\mathbb{Z}$, let us define a function $\gamma\colon [0,1] \to \operatorname{Spec}\mathbb{Z}$ by

$$\gamma(t) = \begin{cases} P & \text{if } t = 0, \\ (0) & \text{if } t \in\,]0,1[, \text{ and} \\ Q & \text{if } t = 1. \end{cases} \tag{A.22}$$

The function γ is continuous, for given any open subspace $U \subset \operatorname{Spec}\mathbb{Z}$, inspection of (A.22) together with our above comment that the point (0) is generic reveals that

$$\gamma^{-1}(U) = \begin{cases} [0,1] & \text{if } P, Q \in U; \\ [0,1[& \text{if } P \in U \text{ and } Q \notin U, \\ ]0,1] & \text{if } P \notin U \text{ and } Q \in U, \text{ and} \\ ]0,1[& \text{otherwise.} \end{cases} \tag{A.23}$$

In other words, the preimage of any open subspace of Spec $\mathbb{Z}$ under γ is open in $[0, 1]$, so that Spec $\mathbb{Z}$ is path connected. (In fact, because (0) is common to all open subspaces of Spec $\mathbb{Z}$, the space is actually locally path connected.)

Nevertheless, Spec $\mathbb{Z}$ fails to be arc connected, and again, it is the existence of the generic point (0) which is the crucial feature. Were there to exist some arc $\gamma\colon [0,1] \to \operatorname{Spec}\mathbb{Z}$ from (0) to some point $P \in \operatorname{Spec}\mathbb{Z}$, then the image $\gamma(]0,1])$ would be an open neighbourhood of P not containing (0). □

In essence, our proof that Peano continua are (locally) arc connected is reasonably intuitive. Given some pair of points x and y of a Peano continuum X, we will use local connectedness to construct a nested sequence of 'ever-smaller' subcontinua of X, each containing x and y, then use Lemma A.2.4 to pass to the intersection of this sequence. Showing that this intersection is an arc then amounts to a straightforward application of our characterization of $[0, 1]$ in Lemma A.4.5.

Despite the conceptual simplicity of the proof, however, there is some terminological overhead involved. Particularly, our construction of the aforementioned nested continua will rely on the idea of a *simple chain*, constructed from a finite sequence of overlapping connected subspaces. (Figure A.2).

Definition A.5.2. *Let x and y denote points of a topological space X. A* simple chain from x to y in X *is a collection $\{U_1, U_2, \ldots, U_n\}$ of open subspaces of X such that*

- *$x \in U_i$ if and only if $i = 1$;*
- *$y \in U_i$ if and only if $i = n$; and*
- *$U_i \cap U_j \neq \emptyset$ if and only if $|i - j| \leq 1$.*

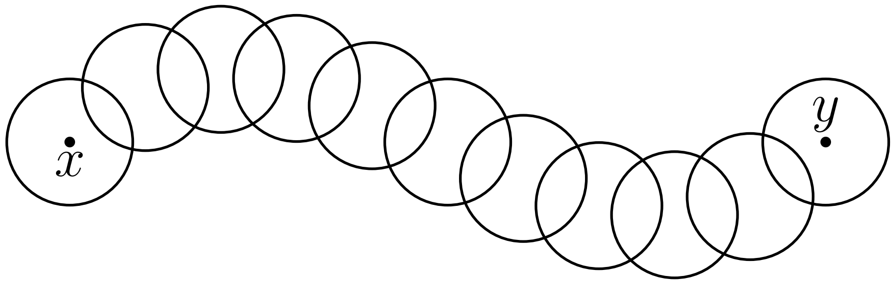

Figure A.2: A simple chain from x to y.

Before we can prove Peano continua to be locally arc connected, we need a preliminary result. Morally, this result tells us that, in a connected space, it is easy to find chains which enjoy prescribed properties: we just need to find an open cover of our space by subspaces enjoying those same properties.

Lemma A.5.3. *Let $\mathcal{U}$ denote some open cover of a connected space X. Then, given any pair of distinct points $x, y \in X$, there exists some simple chain from x to y in X consisting of elements of $\mathcal{U}$.*

Proof. Distinguishing some arbitrary point $x \in X$, we define a subspace

$$Y = \{y \in X \mid x \text{ and } y \text{ are connected by a simple chain of elements of } \mathcal{U}\}. \tag{A.24}$$

Given that X is connected, it will suffice for us to prove that Y is both open and closed in X. Openness is immediate, since if $\{U_i\}_{i=1}^n$ is some simple chain of elements of $\mathcal{U}$ connecting x to a point $y \in X$, then by definition we must have that $U_n \subset Y$. Thus, suppose that $y \in \mathrm{cl}_X Y$.

Let V denote some element of $\mathcal{U}$ which contains y; then, there must exist some point $z \in Y \cap V$. Selecting some chain $\{U_i\}_{i=1}^n$ of elements of $\mathcal{U}$ which connects x to z, we are done if $V = U_i$ for some $i \in \{1, 2, \dots, n\}$, so suppose that this is not the case.

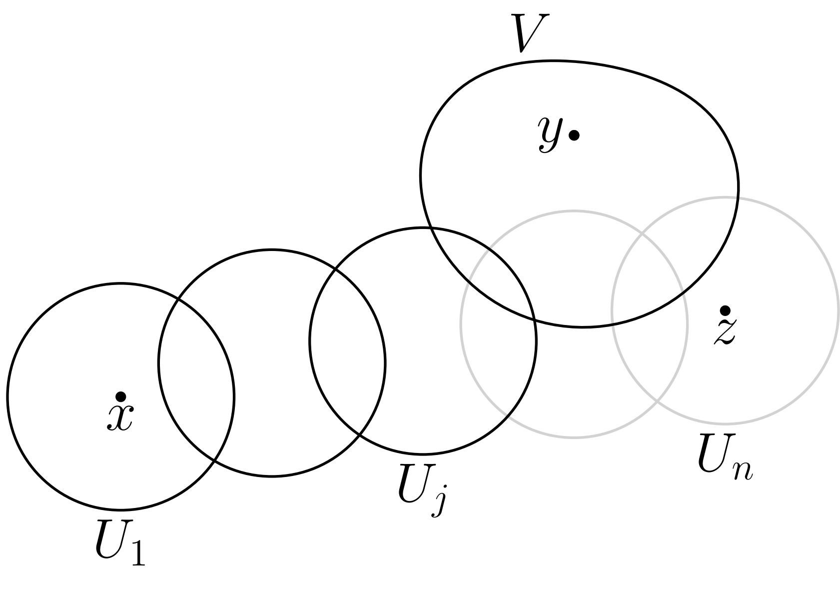

Figure A.3

By hypothesis, the set

$$I = \{i \in \{1, 2, \dots, n\} \mid U_i \cap V \neq \emptyset\} \tag{A.25}$$

is non-empty, so let us define $j = \min I$. Then, the collection $\{U_1, U_2, \dots, U_j, V\}$ comprises a simple chain of elements of $\mathcal{U}$ connecting x to y, so that $y \in Y$; thus, we conclude that $Y = X$. □

Lemma A.5.4. *Every connected open subspace of any Peano continuum is arc connected.*

Proof. Let Y denote some connected open subspace of a Peano continuum X, and distinguish some pair of distinct points $x, y \in Y$; we wish to find an arc in Y which connects x to y.

Knowing that the open subspace Y inherits the local connectedness of X, we deduce the existence of some open cover of Y by connected subspaces, all of diameter strictly less than 1 with respect to some metric d inducing the topology of X. According to Lemma A.5.3, then, there exists some simple chain $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ comprised of such subspaces, which connects x to y in Y.

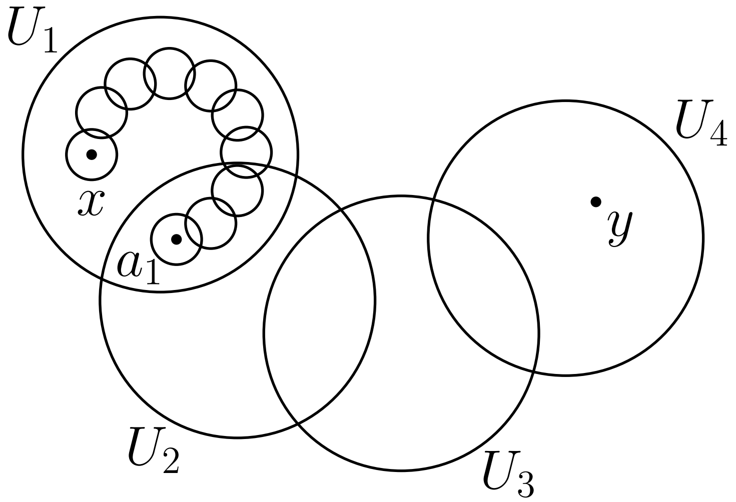

Figure A.4

Now, select some point $a_1 \in U_1 \cap U_2$, recalling that our definition of a simple chain ensures that this intersection is non-empty. Again relying on local connectedness, we can cover U_1 by connected open subspaces of diameter strictly less than $\frac{1}{2}$, as measured by the metric d, and so Lemma A.5.3 delivers a simple chain of such subspaces, connecting x to a_1 in U_1, which we shall denote by $\mathcal{V}_1$ (Figure A.4). Moreover, the regularity of X means that we are free to assume without loss of generality that $\mathrm{cl}_X V \subset U_1$ for each element $V \in \mathcal{V}_1$.

Similarly, we can select a point $a_2 \in U_2 \cap U_3$, and by identical reasoning produce a simple chain $\mathcal{V}_2$ of connected open subspaces, all of diameter strictly less than $\frac{1}{2}$ and all with closures contained within U_2, which connects a_1 to a_2 in U_2 (Figure A.5).

We should like to concatenate the chains $\mathcal{V}_1$ and $\mathcal{V}_2$, in some suitable sense, to produce a simple chain from x to a_2, but there is a technical obstruction to simply taking the union $\mathcal{V}_1 \cup \mathcal{V}_2$. Figure A.5 indicates what can go wrong with this naïve approach: if the chains $\mathcal{V}_1$ and $\mathcal{V}_2$ intersect anywhere except at their first and final elements respectively, then their union cannot possibly be a simple chain.

Fortunately, we are able to sidestep this issue. If $\mathcal{V}_1 = \{V_1, V_2, \dots, V_k\}$ and $\mathcal{V}_2 = \{V_1', V_2', \dots, V_{k'}'\}$, then let us define

$$\begin{aligned} l &= \max\{i \in \{1, 2, \dots, k\} \mid V_i \text{ intersects any element of } \mathcal{V}_2\} \text{ and} \\ l' &= \min\{i \in \{1, 2, \dots, k'\} \mid V_i' \cap V_l \neq \emptyset\}. \end{aligned} \tag{A.26}$$

With these definitions, we see readily that $\{V_1, V_2, \dots, V_l, V_{l'}', \dots, V_{k'}'\}$ constitutes a simple chain connecting x to a_2. Continuing in this fashion for each $i \in \{1, 2, \dots, n-1\}$ yields a chain $\mathcal{V}_{n-1}$ in Y connecting x to y such that, for each $V \in \mathcal{V}_{n-1}$,

- V is connected and open;

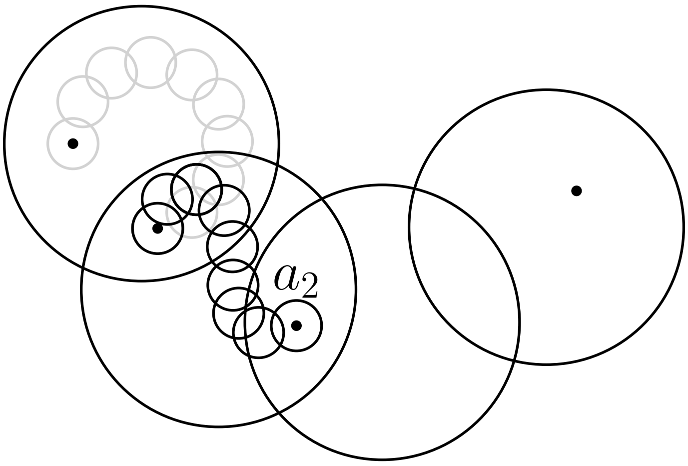

Figure A.5

- $\operatorname{diam}_d V < \frac{1}{2}$; and
- $\operatorname{cl}_X V \subset \bigcup_{i=1}^{n} U_i$.

Such a chain is depicted schematically in Figure A.6.

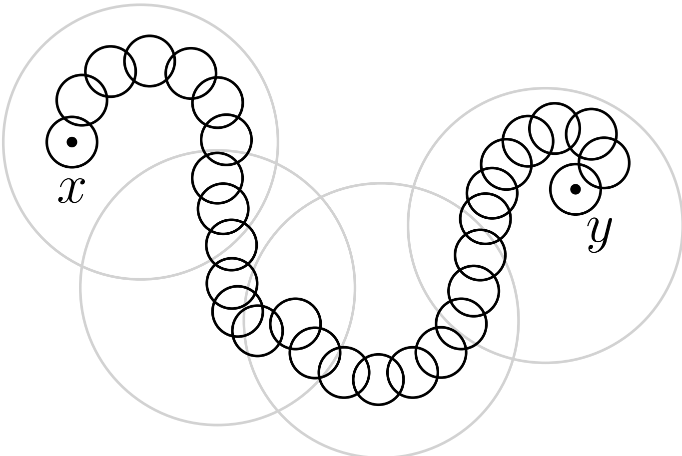

Figure A.6

Crucially, we can iterate this construction. This yields, for each $n \in \mathbb{N}$, some simple chain $\mathcal{U}_n$ connecting x to y in X such that, for each $U \in \mathcal{U}_n$,

- U is connected,
- $\operatorname{diam}_d U < \frac{1}{n}$, and
- $\operatorname{cl}_X U \subset \bigcup_{V \in \mathcal{U}_{n-1}} V$, when $n > 1$.

Introducing the notation

$$C_n = \operatorname{cl}_X \left(\bigcup_{U \in \mathcal{U}_n} U \right) \tag{A.27}$$

for each $n \in \mathbb{N}$, we arrive at a descending sequence

$$C_1 \supset C_2 \supset C_3 \supset \cdots \tag{A.28}$$

of continua in Y, each of which contains the points x and y. Lemma A.2.4 asserts that the intersection

$$C = \bigcap_{n \in \mathbb{N}} C_n \tag{A.29}$$

is itself a continuum containing x and y, so that we need only demonstrate that every point of C other than x and y is a cut point in order to complete the proof by way of an appeal to Lemma A.4.5.

Selecting any point $z \in C - \{x, y\}$, we know that for each $n \in \mathbb{N}$, the point z belongs to either one or two elements of the chain $\mathcal{U}_n$; thus, for each $n \in \mathbb{N}$, we define

$$\begin{aligned} W_n &= \bigcup\{U \in \mathcal{U}_n \mid U \text{ is strictly prior to all elements of } \mathcal{U}_n \text{ which contain } z\} \text{ and} \\ W'_n &= \bigcup\{U \in \mathcal{U}_n \mid U \text{ is strictly later than all elements of } \mathcal{U}_n \text{ which contain } z\}, \end{aligned} \tag{A.30}$$

as sketched in Figure A.7. We propose that the open subspaces $C \cap \bigcup_{n\in\mathbb{N}} W_n$ and $C \cap \bigcup_{n\in\mathbb{N}} W'_n$ cover C, which will prove $C - \{z\}$ to be disconnected, since these subspaces are disjoint by construction.

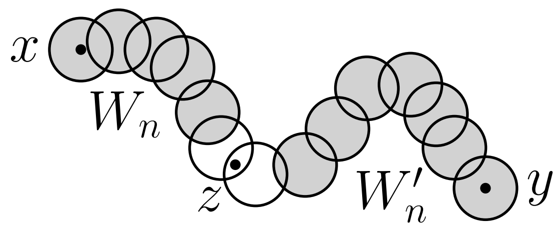

Figure A.7

Consider any point $c \in C - \{z\}$, and choose some $N \in \mathbb{N}$ such that $\frac{1}{2N} < d(c, z)$. Then, having arranged that $\operatorname{diam}_d U < \frac{1}{N}$ for every $U \in \mathcal{U}_N$, we find that c and z must belong to different elements of the chain $\mathcal{U}_N$–but this implies that either $c \in W_N$ or $c \in W'_N$. It follows that the subspaces $C \cap \bigcup_{n\in\mathbb{N}} W_n$ and $C \cap \bigcup_{n\in\mathbb{N}} W'_N$ do indeed cover C, so that $C \subset Y$ is an arc from x to y. □

Corollary A.5.5. *Every Peano continuum is locally arc connected.*

Proof. We have just demonstrated that every connected open subspace of a Peano continuum is locally arc connected, so that this is an immediate consequence of the fact that Peano continua are, by definition, locally connected. □

As an aside, our newfound knowledge that Peano continua are arc connected allows us to quickly prove that pathologies of the kind discussed in Example A.5.1 are only possible in a non-Hausdorff setting. That is, arc connectedness is precisely equivalent to path connectedness for Hausdorff spaces.

Lemma A.5.6. *For any Hausdorff space X, the following are equivalent:*

- *X is path connected, and*
- *X is arc connected.*

Proof. One implication is trivial, so suppose that X is path connected. Given any pair of distinct points $x, y \in X$ and any path $\gamma\colon [0, 1] \to X$ from x to y, the image $\gamma([0, 1]) \subset X$ is of course compact and connected. Moreover, the same technique as that used in our proof of Lemma 1.4.4 allows us to see that $\gamma([0, 1])$ is also locally connected.

Thus, Corollary A.5.5 tells us that there exists some arc from x to y in $\gamma([0, 1])$, and it follows that X is arc connected. □

A.5.1 Uniform Local Arc Connectedness of Peano Continua

In fact, we need not content ourselves merely with local arc connectedness: Peano continua turn out to be, in a suitable sense, *uniformly locally arc connected.* We shall need this fact only once in what follows, but its use cannot be avoided: it is the essential technical ingredient in our proof of the Hahn-Mazurkiewicz-Sierpiński theorem in Section A.6.3.

Definition and Lemma A.5.7. *Peano continua are* uniformly locally arc connected. *That is, if X denotes some Peano continuum topologized by a metric d, then there exists for each $\varepsilon > 0$ some $\delta > 0$ such that, if $d(x, y) < \delta$ for some pair of distinct points $x, y \in X$, then there exists some arc $[xy] \subset X$ for which* $\operatorname{diam}[xy] < \varepsilon$.

Proof. We start by proving that any Peano continuum X is uniformly locally connected, in an appropriate sense. Fixing some $\varepsilon > 0$, we can find some finite open cover $\mathcal{U}$ of X by connected subspaces, each of diameter strictly less than $\frac{\varepsilon}{2}$; because X is compact Hausdorff, this cover must have some Lebesgue number δ. It follows that, if $x, y \in X$ are points such that $d(x, y) < \delta$, then there exists a connected subspace $U \subset X$ of diameter strictly less than $\frac{\varepsilon}{2}$ which contains both x and y: namely, any element $U \in \mathcal{U}$ which contains x.

To promote this to uniform local arc connectedness, we rely on Lemma A.5.4. Each point $z \in U$ permits some connected neighbourhood $U_z \subset X$ of diameter strictly less than $\frac{\varepsilon}{4}$, so that the union $V = \bigcup_{z \in U} U_z$ is a connected open subspace of X of diameter strictly less than ε; according to Lemma A.5.4, such a subspace is in fact arc connected.

This construction ensures that $U \subset V$, and we know that if $d(x, y) < \delta$, then $x, y \in U$; thus, the uniform local arc connectedness of X is established. □

A.5.2 Arc Accessible Boundary Points

One relatively direct–but extremely useful–consequence of the local arc connectedness of Peano continua is that, in such spaces, we can 'almost always' find an arc from a point of a given open subspace to a point on that subspace's boundary. More precisely, we have the following fact, which is used extensively throughout Chapter 2.

Definition and Lemma A.5.8. *Let X denote some topological space with a subspace $Y \subset X$, and select some point $x \in X - Y$. We say that x is* arc accessible from Y *if there exists some arc $T \subset Y \cup \{x\}$ with x as an end point.*

Suppose that X is some Peano continuum, and $U \subsetneq X$ some open subspace. The set of points of $\operatorname{fr} U$ *which are arc accessible from U is dense in* $\operatorname{fr} U$.

Proof. Select some point $x \in \operatorname{fr} U$, and denote by $V \subset X$ some arc connected neighbourhood of x, which Corollary A.5.5 ensures must exist. Choosing some point $y \in U \cap V$, we know that there must exist some arc $[yx] \subset V$, and we define p to be the point at which this arc first meets $\operatorname{fr} U$; essentially by definition, p is then arc accessible from U. It follows that the subspace of $\operatorname{fr} U$ which is arc accessible from U is indeed dense in $\operatorname{fr} U$. □

A.6 Characterizing Peano Continua

Although local connectedness is a formally simple property, it can in practice be far from straightforward to verify whether or not a given continuum is Peano. For later use, we prove in this section the *Hahn-Mazurkiewicz-Sierpiński theorem* [Maz13a, Maz13b, Hah14, Sie20], a classical and striking characterization of Peano continua as, equivalently, those continua which are the continuous image of the closed unit interval, or as those continua which can be assembled from finitely many arbitrarily small Peano subcontinua.

Before we tackle the Hahn-Mazurkiewicz-Sierpiński theorem head-on, however, we will need a handful of preliminary results, revolving around Sierpiński's notion of *Property S*, and around the *Alexandroff-Hausdorff theorem* [Ale27, Hau27]. Property S affords us a useful alternative characterization of local connectedness in the compact Hausdorff setting, whereas the Alexandroff-Hausdorff theorem supplies a continuous surjection from the Cantor space to any compact metrizable space.

A.6.1 Property S

Introduced by Sierpiński [Sie20] precisely for the purpose of identifying Peano continua amongst the class of arbitrary continua, Property S can be described fairly straightforwardly in modern language, and shown to be equivalent to local connectedness for compact Hausdorff spaces. In this work, we are interested in Property S primarily for its utility in proving that a connected union of finitely many Peano continua is itself a Peano continuum.

Definition A.6.1. *Let X denote some topological space, with the property that every open cover of X permits a refinement by finitely many connected subspaces. Then, the space X is said to have* Property S*.*

Lemma A.6.2. *For any compact Hausdorff space X, the following are equivalent:*

- *X is locally connected, and*
- *X has Property S.*

Proof. Supposing that the compact Hausdorff space X is locally connected, let $\mathcal{U}$ denote some arbitrary open cover of X. Each point $x \in X$ belongs to at least one element $U_x \in \mathcal{U}$, and the assumed local connectedness of X asserts the existence of some corresponding connected neighbourhood $V_x \subset U_x$ of x. Compactness demands that the collection $\{V_x\}_{x \in X}$ of all such neighbourhoods be essentially finite, which tells us that X has Property S.

Conversely, suppose that the compact Hausdorff space X enjoys Property S, and select some arbitrary point $x \in X$, along with some neighbourhood $U \subset X$ of x. Every compact Hausdorff space is regular, and so we can find some neighbourhood $V \subset X$ of the point x such that $\operatorname{cl} V \subset U$; then, the pair $\{U, X - \operatorname{cl} V\}$ comprises an open cover of X. Our definition of Property S requires that this open cover have some refinement by finitely many connected subspaces, and choosing any element of this refinement containing x supplies us with a connected neighbourhood of x which is contained within our original neighbourhood U. □

Lemma A.6.3. *Let $\{X_1, X_2, \ldots, X_n\}$ denote some collection of Peano continua for which the union $X = X_1 \cup X_2 \cup \cdots \cup X_n$ is connected. Then, the union X is a Peano continuum.*

Proof. It is clear immediately that X is a continuum, and so we need only concern ourselves with local connectedness. We are free to restrict ourselves to the setting in which $n = 2$, with the general case following inductively, and in light of Lemma A.6.2, we need only prove that the union $X = X_1 \cup X_2$ has Property S.

Indeed, denote by $\mathcal{U}$ some open cover of X; then,

$$\mathcal{U}_1 = \{U \cap X_1 \,|\, U \in \mathcal{U}\} \text{ and } \mathcal{U}_2 = \{U \cap X_2 \,|\, U \in \mathcal{U}\} \tag{A.31}$$

are open covers of X_1 and X_2 respectively. The fact that both X_1 and X_2 have Property S provides us with finite refinements $\mathcal{V}_1$ and $\mathcal{V}_2$ of $\mathcal{U}_1$ and $\mathcal{U}_2$, respectively, by connected subspaces; then, the union $\mathcal{V}_1 \cup \mathcal{V}_2$ is a refinement of the original cover $\mathcal{U}$ by finitely many connected subspaces. In other words, the union X has Property S, and must therefore be a Peano continuum. □

A.6.2 The Alexandroff-Hausdorff Theorem

Before proving the Alexandroff-Hausdorff theorem we first recall two characterizations of the Cantor space, well-known [AB06] to be equivalent to one another, along with the fact that the Cantor space retracts onto each of its closed subspaces.

Definition A.6.4. *The* Cantor space *is the product $C = \{0,1\}^{\mathbb{N}}$, where each instance of $\{0,1\}$ is equipped with the discrete topology. Where it is necessary to consider this product as a metric space, we shall use the metric*

$$d(\mathbf{c}, \mathbf{d}) = \sum_{n=1}^{\infty} \frac{|c_n - d_n|}{3^n},$$

where $\mathbf{c} = (c_1, c_2, c_3, \ldots)$ and $\mathbf{d} = (d_1, d_2, d_3, \ldots)$.

Equivalently, the Cantor space is the subspace of the closed unit interval $[0,1]$ *produced by the standard* middle-thirds *construction, detailed in, for instance, [Wil70]. The map*

$$(c_1, c_2, c_3, \ldots) \mapsto \sum_{n=1}^{\infty} \frac{2c_n}{3^n}$$

is a homeomorphism between the two realizations of the Cantor space.

Lemma A.6.5. *Every non-empty closed subspace of the Cantor space* C *is a retract of* C*. That is, if* $X \subset C$ *is closed and non-empty, then there exists some continuous function* $r\colon C \to X$ *such that* $r|_X = \mathrm{id}_X$*.*

Proof. Endowing C with the metric d of Definition A.6.4, let $X \subset C$ denote some non-empty compact subspace. We propose that, for each point $\mathbf{c} \in C$, there exists precisely one point $r(\mathbf{c}) \in X$ such that

$$d(\mathbf{c}, r(\mathbf{c})) = \inf_{\mathbf{x} \in X} d(\mathbf{c}, \mathbf{x}) = d(\mathbf{c}, X). \tag{A.32}$$

The existence of at least one such point is straightforward, since for any fixed point $\mathbf{c} \in C$, the function $\mathbf{x} \mapsto d(\mathbf{c}, \mathbf{x})$ is a continuous function on the compact space X, which must attain its greatest lower bound.

As for uniqueness, fix some point $\mathbf{c} \in C$, and suppose that $\mathbf{x}, \mathbf{y} \in X$ are some pair of points such that

$$d(\mathbf{c}, \mathbf{x}) = d(\mathbf{c}, \mathbf{y}) = d(\mathbf{c}, X). \tag{A.33}$$

Then, we know that

$$d(\mathbf{c}, \mathbf{x}) = \sum_{n=1}^{\infty} \frac{|c_n - x_n|}{3^n} = \sum_{n=1}^{\infty} \frac{|c_n - y_n|}{3^n} = d(\mathbf{c}, \mathbf{y}), \tag{A.34}$$

which can be so only if $c_n - x_n = c_n - y_n$ for each $n \in \mathbb{N}$, justifying our claim. Thus, (A.32) specifies a well defined function $r\colon C \to X$, which by construction restricts to the identity map on X. If we can additionally show this function to be continuous, then the proof will be complete.

To this end, consider some sequence $(\mathbf{c}_n)_{n \in \mathbb{N}}$ of points of C with a limit $\mathbf{c} \in C$, and suppose for the sake of a contradiction that the sequence $(r(\mathbf{c}_n))_{n \in \mathbb{N}}$ fails to converge to $r(\mathbf{c})$. By using compactness to pass to a subsequence if necessary, we lose no generality in assuming that $r(\mathbf{c}_n) \to \mathbf{x}$ for some point $\mathbf{x} \in X$.

Now, we know by definition that $d(\mathbf{c}, r(\mathbf{c})) = d(\mathbf{c}, X)$; further, knowing that the function $\mathbf{c} \mapsto d(\mathbf{c}, X)$ is continuous on C, it follows that

$$d(\mathbf{c}, r(\mathbf{c})) = \lim_{n \to \infty} d(\mathbf{c}_n, X) = \lim_{n \to \infty} d(\mathbf{c}_n, r(\mathbf{c}_n)) = d(\mathbf{c}, \mathbf{x}). \tag{A.35}$$

This is the contradiction we seek for we assumed the points $r(\mathbf{c})$ and $\mathbf{x}$ to be distinct, but this together with (A.35) violates the uniqueness result proved above. □

Although the statement of the Alexandroff-Hausdorff theorem is perhaps at first surprising, it is not especially complicated to prove. Relatively direct constructions of the requisite surjection in terms of inverse limit spaces are possible [Wil70], but here we opt for a more 'hands-off' but less technically intricate construction, fundamentally facilitated by Lemma A.6.5.

Theorem A.6.6 (Alexandroff-Hausdorff)**.** *Every compact metrizable space is a continuous image of the Cantor space.*

Proof. Our strategy shall be to relate an arbitrary compact metrizable space to the Cantor space by mapping both into the Hilbert cube.

Letting C denote the Cantor space, recall that the so-called *Devil's staircase* may defined as the function $\varphi\colon C \to [0,1]$ acting by

$$\varphi(c_1, c_2, c_3, \ldots) = \sum_{i=1}^{\infty} \frac{c_i}{2^i}. \tag{A.36}$$

Using the metric d of Definition A.6.4 allows us to see immediately that φ is continuous. Moreover, we notice that φ must be surjective, since every point of $[0,1]$ permits at least one dyadic expansion.

The function φ of course induces a continuous surjection $\widetilde{\varphi}\colon C^{\mathbb{N}} \to [0,1]^{\mathbb{N}}$, explicitly defined by

$$\widetilde{\varphi}(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \ldots) = (\varphi(\mathbf{c_1}), \varphi(\mathbf{c_2}), \varphi(\mathbf{c_3}), \ldots). \tag{A.37}$$

However, the observation that

$$C^{\mathbb{N}} = \left(\{0,1\}^{\mathbb{N}}\right)^{\mathbb{N}} \cong \{0,1\}^{\mathbb{N}\times\mathbb{N}} \cong \{0,1\}^{\mathbb{N}} \cong C \tag{A.38}$$

allow us to formally identify the Cantor space C with the product $C^{\mathbb{N}}$, and so we arrive at a continuous surjection $C \to [0,1]^{\mathbb{N}}$. Abusing notation slightly, we use $\widetilde{\varphi}$ to denote this function also.

Now, let X denote some compact metrizable space. We know that there must exist[5] some embedding $\iota\colon X \to [0,1]^{\mathbb{N}}$, so that the image $\iota(X)$ is closed in $[0,1]$; thus, the preimage $(\widetilde{\varphi}^{-1} \circ \iota)(X)$ is a closed subspace of C. Using Lemma A.6.5, we deduce the existence of some retract $r\colon C \to (\widetilde{\varphi}^{-1} \circ \iota)(X)$. The composition $\widetilde{\varphi} \circ r$ is then the desired continuous surjection from C to (a homeomorph of) X. □

We remark in passing that the hypothesis of metrizability cannot be omitted from Theorem A.6.6; the Stone-Čech compactification $\beta\mathbb{N}$ of $\mathbb{N}$ serves as an explicit counterexample.

Example A.6.7 (The Stone-Čech compactification of $\mathbb{N}$)**.** We do not concern ourselves with the details of how one might define the Stone-Čech compactification of a given locally compact Hausdorff space, referring to [Wil70]. For the purposes of this example, we concern ourselves only with $\beta\mathbb{N}$, the Stone-Čech compactification of the natural numbers $\mathbb{N}$.

The Cantor space C is well-known to have the cardinality $\mathfrak{c}$ of the continuum, while one can show [Wal74] that $\beta\mathbb{N}$ has cardinality $2^{\mathfrak{c}}$; in particular, there can exist no surjection $C \to \beta\mathbb{N}$. However, the space $\beta\mathbb{N}$ cannot be metrizable, because it is not even first countable: any separable, first countable Hausdorff space has a cardinality of at most $2^{\aleph_0}$. (To see why, notice that each point of such a space is the limit of some sequence taking values in a countable subspace.) □

A.6.3 The Hahn-Mazurkiewicz-Sierpiński Theorem

We now have all of the building blocks necessary for a proof of the Hahn-Mazurkiewicz-Sierpiński theorem; the only thing that remains is for us to put them all together. We shall see that most of the requisite implications are straightforward consequences of our work so far, with only one (namely, that every Peano continuum is a quotient of the closed unit interval $[0,1]$) being especially involved.

Theorem A.6.8 (Hahn-Mazurkiewicz-Sierpiński)**.** *For any continuum X, the following are equivalent:*

- *X is Peano;*
- *X is a continuous image of the closed unit interval $[0,1]$; and*
- *X can be expressed as the union of finitely many Peano continua of diameter strictly less than ε, given any $\varepsilon > 0$.*

Proof. We shall organise each implication of this proof in roughly ascending order of complexity. Thus, we proceed as follows:

1. We start by showing that every continuous image of $[0,1]$ is Peano;
2. Using this, we prove that every continuous image of $[0,1]$ satisfies the third condition in the theorem's statement;

[5] One means of constructing such an embedding is to invoke Urysohn's lemma, producing from some countable basis for the topology of X a collection $\{f_1, f_2, f_3, \ldots\}$ of countably many continuous functions $X \to [0,1]$. Then, it can be shown that the evaluation map $x \mapsto (f_1(x), f_2(x), f_3(x), \ldots)$ is the desired embedding of X into $[0,1]$. This procedure is discussed in explicit detail in, for instance, [Wil70].

3. Then, the machinery of Property S will allow us to demonstrate that every continuum satisfying the third condition above is Peano; and finally
4. We construct a continuous surjection from $[0, 1]$ to an arbitrary Peano continuum X.

The conceptual work for the first of these steps is already done: by re-using the method employed for our proof of Lemma 1.4.4 and recalling that every continuous function from a compact space to a Hausdorff space is closed, we find that every continuous image of $[0, 1]$ is a Peano continuum.

For the second step, suppose that $\pi\colon [0, 1] \to X$ is some continuous surjection and select any $\varepsilon > 0$, along with any metric d which topologizes the continuum X. Knowing that every continuous function on a compact Hausdorff space is uniformly continuous, we can find some $\delta > 0$ such that if $|x - y| < \delta$ for points $x, y \in [0, 1]$, then $d(\pi(x), \pi(y)) < \varepsilon$. Thus, we decompose $[0, 1]$ as a union

$$[0, 1] = I_1 \cup I_2 \cup \cdots \cup I_n \tag{A.39}$$

of closed intervals of diameter strictly less than δ, so that

$$X = \pi([0, 1]) = \pi(I_1) \cup \pi(I_2) \cup \cdots \cup \pi(I_n). \tag{A.40}$$

Each of the closed intervals $\{I_1, I_2, \ldots, I_n\}$ is of course homeomorphic to $[0, 1]$, and we have already demonstrated that continuous images of $[0, 1]$ are Peano; the second step is therefore dealt with.

As for the third, we need only invoke Lemma A.6.3, and so all that remains is to find, given a Peano continuum X topologized by a metric d, a continuous surjection $[0, 1] \to X$. Using Theorem A.6.6, we can produce a continuous surjection $f\colon C \to X$, where we denote by $C \subset [0, 1]$ the standard middle-thirds Cantor set; from here, we shall use Lemma A.5.7's guarantees of uniform local arc connectedness to extend f to a continuous surjection $\widetilde{f}\colon [0, 1] \to X$.

We know that the complement $[0, 1] - C$ has countably many components, all of which are open intervals, which we denote by $\{I_1, I_2, I_3, \ldots\}$, where $I_n =]a_n, b_n[$ for each $n \in \mathbb{N}$. Moreover, we enumerate these intervals such that, if $n < m$, then either $\operatorname{diam} I_m < \operatorname{diam} I_n$, or else $\operatorname{diam} I_m = \operatorname{diam} I_n$ and $b_n < a_m$.

Now, Lemma A.5.7 allows us to find, for each $n \in \mathbb{N}$, some $\delta_n > 0$ such that if $x, y \in X$ are distinct points such that $d(x, y) < \delta_n$, then there exists an arc in X from x to y of diameter strictly less than 2^{-n}. Moreover, the surjection $f\colon C \to X$ is uniformly continuous, so that for each $n \in \mathbb{N}$, we can also select some $\zeta_n > 0$ such that, if $|c - d| < \zeta_n$ for points $c, d \in C$, then $d(f(c), f(d)) < \delta_n$.

Of course, the interval $[0, 1]$ is bounded, so that only finitely many of the intervals $\{I_n\}_{n\in\mathbb{N}}$ have diameter at least ζ_1; let us denote these intervals by $\{I_1, I_2, \ldots, I_{k_1}\}$. For each $i \in \{1, 2, \ldots, k_1\}$, we define the restriction $\widetilde{f}|_{\operatorname{cl} I_i}$ by

- $\widetilde{f}(\operatorname{cl} I_i) = f(a_i) = f(b_i)$, if $f(a_i) = f(b_i)$; or otherwise
- $\widetilde{f}|_{\operatorname{cl} I_i} = \gamma_i$, where γ_i is any homeomorphism from the closed interval $[a_i, b_i]$ to any arc from $f(a_i)$ to $f(b_i)$ in X.

Similarly, only a finite subcollection $\{I_{k_1+1}, I_{k_1+2}, \ldots, I_{k_2}\}$ of the arcs $\{I_n\}_{n\in\mathbb{N}}$ have diameters in $[\zeta_2, \zeta_1[$. For each $i \in \{k_1 + 1, k_1 + 2, \ldots, k_2\}$, we define $\widetilde{f}|_{\operatorname{cl} I_i}$ similarly to the above case, but with an additional condition; since we now know that $\operatorname{diam} I_i < \zeta_1$, we also know that $d(f(a_i), f(b_i)) < \delta_1$. In particular, we are able to join $f(a_i)$ and $f(b_i)$ (when the two points do not coincide) by an arc of diameter strictly less than $\frac{1}{2}$. Explicitly, we define

- $\widetilde{f}(\operatorname{cl} I_i) = f(a_i) = f(b_i)$, if $f(a_i) = f(b_i)$; or otherwise
- $\widetilde{f}|_{\operatorname{cl} I_i} = \gamma_i$, for any homeomorphism $\gamma_i\colon [a_i, b_i] \to T_i$, where $T_i \subset X$ is some arc from $f(a_i)$ to $f(b_i)$ such that $\operatorname{diam} T_i < \frac{1}{2}$. Notice that our definition of ζ_1 ensures the existence of such an arc.

We can continue in this fashion for each $n \in \mathbb{N}$. That is, for each component I_i of $[0, 1] - C$ such that $\operatorname{diam} I_i \in [\zeta_{n+1}, \zeta_n[$, we define $\widetilde{f}|_{\operatorname{cl} I_i}$ by

- $\widetilde{f}(\operatorname{cl} I_i) = f(a_i) = f(b_i)$, if $f(a_i) = f(b_i)$; or otherwise

- $\widetilde{f}|_{\mathrm{cl}\, I_i} = \gamma_i$, for any homeomorphism $\gamma_i\colon [a_i, b_i] \to T_i$, where $T_i \subset X$ is some arc from $f(a_i)$ to $f(b_i)$ such that $\operatorname{diam} T_i < 2^{-n}$.

Thus, we produce an extension $\widetilde{f}\colon [0,1] \to X$ of f, which of course must be surjective. If we are able to show that $\widetilde{f}$ is also continuous, then we will be done, and to achieve this, we prove continuity at points of $[0,1] - C$ and points of C separately.

Given a point $x \in [0,1] - C$, we know that $x \in I_i$ for some $i \in \mathbb{N}$. Denoting by $(x_n)_{n\in\mathbb{N}}$ some sequence in $[0,1]$ converging to x, and by $U \subset X$ some arbitrary neighbourhood of $\widetilde{f}(x)$, we know that $(x_n)_{n\in\mathbb{N}}$ eventually lies in $I_i \cap \widetilde{f}^{-1}(U)$. Having defined $\widetilde{f}|_{\mathrm{cl}\, I_i}$ to be a homeomorphism, we find immediately that $\widetilde{f}(x_n) \in U$ for all sufficiently large values of n; it follows that $\widetilde{f}(x_n) \to \widetilde{f}(x)$, so that $\widetilde{f}$ is continuous on $[0,1] - C$.

If, instead, $x \in C$, we can distinguish three possibilities:

- $x = a_i$ for some $i \in \mathbb{N}$;
- $x = b_i$ for some $i \in \mathbb{N}$; or
- $x \notin \{a_1, b_1, a_2, b_2, \ldots\}$.

When $x = a_i$ for some $i \in \mathbb{N}$, we know from the above that $\widetilde{f}$ is right-continuous at x, so that it will suffice for us to prove that $\widetilde{f}$ is also left-continuous at x. With this aim, prescribe some $\varepsilon > 0$, select some $N \in \mathbb{N}$ such that $2^{1-N} < \varepsilon$, and choose some $\zeta \in \,]0, \zeta_N[$ such that, for any point $y \in \,]x - \zeta, x[$, either

- $y \in C$, or
- $y \in I_n$ for some $n \in \mathbb{N}$, where $\operatorname{diam} I_n < \zeta_N$.

From here, we can see that $\widetilde{f}$ is indeed left-continuous at x. Explicitly, when $y \in C$, our definition of ζ_N guarantees that

$$d\big(\widetilde{f}(x), \widetilde{f}(y)\big) = d(f(x), f(y)) < \frac{1}{2^N} < \frac{1}{2^{N-1}} < \varepsilon, \tag{A.41}$$

whereas if $y \in I_n$ for some $n \in \mathbb{N}$, we can deduce that

$$\begin{aligned} d\big(\widetilde{f}(x), \widetilde{f}(y)\big) &\le d\big(\widetilde{f}(x), f(b_n)\big) + d\big(f(b_n), \widetilde{f}(y)\big) \\ &\le d(f(x), f(b_n)) + d\big(f(b_n), \widetilde{f}(y)\big) \\ &< \frac{1}{2^N} + \frac{1}{2^N} \\ &\le \frac{1}{2^{N-1}} \\ &< \varepsilon. \end{aligned} \tag{A.42}$$

When $x = b_i$ for some $i \in \mathbb{N}$, the continuity of $\widetilde{f}$ follows by a formally similar argument, in which the roles of left- and right-continuity are exchanged. Should it be the case that $x \notin \{a_1, b_1, a_2, b_2, \ldots\}$, then we deduce the continuity of $\widetilde{f}$ at x by combining the argument for the left-continuity of $\widetilde{f}$ on $\{a_1, a_2, a_3, \ldots\}$ with that for the right-continuity of $\widetilde{f}$ on $\{b_1, b_2, b_3, \ldots\}$, and this completes the proof. □

A.7 The 2-Bogensatz

Menger's theorem [Men27] is a classical result in graph theory which, roughly, states that if u and v are some pair of vertices in a finite graph G, then the number of disjoint paths in G from u to v counts the number of vertices which must be removed from G in order to separate u and v. A number of contrasting modern proofs of this fact are presented in [Die05].

The 2-*Bogensatz*[6] which we prove in this section is readily seen to be a continuum-theoretic analogue of a special case of Menger's theorem, and will play an important role in our work in Section A.10.

[6] German for the 2 *arc theorem*.

It is interesting to note that the 2-Bogensatz can be used as the basis case for an inductive proof [Nöb32, Zip33, Why48] of a complete analogue of Menger's theorem, appropriately known as the *n-Bogensatz*, although this is highly non-trivial and not necessary for our purposes. We also remark that the relationship between continuum theory and graph theory extends deeper than a superficial parallel to Menger's theorem; in [Nad93], Nadler makes productive use of this connection, and in [Nad92], the same author gives a pedagogical overview of how the theories of continua and graphs are intertwined.

Our method of proof for Lemma A.7.1 is essentially an expansion of that presented by Whyburn in [Why45].

Lemma A.7.1 (2-Bogensatz)**.** *Let X denote some Peano continuum with no cut points, and consider some pair $A, B \subset X$ of disjoint, non-degenerate closed subspaces. There then exists some pair of disjoint arcs from A to B.*

Proof. We define a subspace $S \subset X$ which consists of all points $x \in X$ for which there exist arcs $[ab], [px] \subset X$ such that

- $a, p \in A$ and $b \in B$; and
- $[ab] \cap [px] = \emptyset$.

If we can show that $S = X$, the result will follow, and because X is connected, it will suffice to prove that S is non-empty, open and closed in X.

For non-emptiness, we use the assumed non-degeneracy of A to select some distinct pair of points $a_1, a_2 \in A$, along with some point $b \in B$. The arc connectedness of X yields some arc $[a_1b] \subset X$, and local arc connectedness supplies us with some arc connected neighbourhood $U \subset X$ of a_2 which is disjoint from the arc $[a_1b]$. It follows that $U \subset S$, so that S is at least non-empty. We sketch this in Figure A.8a.

Similar reasoning[7] indicates that S is also open. Indeed, given some point $x \in S$, there must exist some arc $[a_1b]$ from A to B and some arc $[a_2x]$ such that $[a_1b] \cap [a_2x] = \emptyset$. Choosing some arc connected neighbourhood $V \subset X$ of x which is disjoint from $[a_1b]$ and arbitrarily selecting a point $y \in V$, there of course exists some arc $[xy] \subset V$; the union $[a_2x] \cup [xy]$ then contains an arc $[a_2y]$ which is disjoint from $[a_1b]$, as indicated in Figure A.8b.

It remains for us to prove that S is closed. Choosing any point $x \in \operatorname{cl} S$ and some arc connected neighbourhood $W \subset X$ of x, there must exist some point $y \in W \cap S$; using our definition of S, we deduce the existence of arcs $[a_1b]$ and $[a_2y]$ of the above form.[8] Supposing that W and $[a_1b]$ are disjoint, we can find some arc $[yx] \subset W$, necessarily not meeting $[a_1b]$. Then, the union $[a_2y] \cup [yx]$ contains an arc $[a_2x]$ disjoint from $[a_1b]$; thus, it follows that $x \in S$ in this case.

Suppose instead that $W \cap [a_1b] \neq \emptyset$. Our hypotheses ensure that x is not a cut point of X, so that there must exist some arc $[cd] \subset X - \{x\}$ such that $[cd] \cap A = \{c\}$ and $[cd] \cap B = \{b\}$; moreover, the local arc connectedness of X ensures that we lose no generality in assuming that $W \cap A = W \cap [cd] = \emptyset$. Notice that we are done if $[cd] \cap [a_2y] = \emptyset$: in this case, we can simply find an arc $[yx] \subset W$, so that the union $[a_2y] \cup [yx]$ contains an arc $[a_2x]$ disjoint from $[cd]$. Inspecting our definition of S reveals that this means that $x \in S$.

Thus, let us assume that $[cd]$ and $[a_2y]$ fail to be disjoint. We define the following points, indicated in Figure A.8c.

- r_1: the point at which $[a_1b]$ first intersects $\operatorname{cl} W$;
- r_2: the point at which $[a_2y]$ first intersects $\operatorname{cl} W$; and

[7] In an effort to economise on notation, we re-use the symbols a_1, a_2 and b here, although there is no need for any of these points to coincide with those with the same names in the previous paragraph. This overloading of notation is depicted in Figure A.8.

[8] Again, we recycle the symbols a_1, a_2 and b for points which need not be the same as those in the preceding portions of the proof.

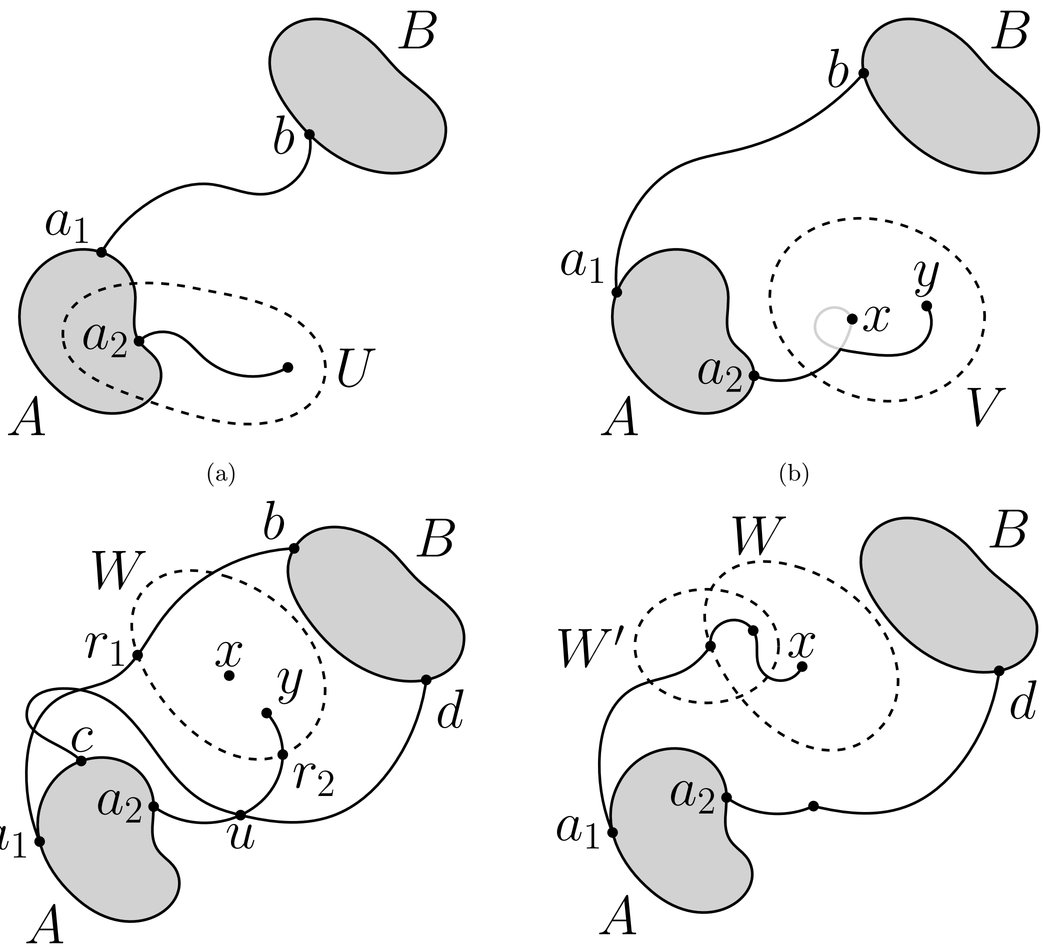

(a) (b) (c) (d)

Figure A.8

- u: the point at which $[cd]$ last intersects the union $[a_1r_1] \cup [a_2r_2]$. (Our assumption that $[cd] \cap [a_2y] \neq \emptyset$ ensures the existence of such a point.)

There are two mutually exclusive possibilities for where the point u may lie: either $u \in [a_2r_2]$, as sketched in Figure A.8c, or $u \in [a_1r_1]$.

If $u \in [a_2r_2]$, then the union $[a_2u] \cup [ud]$ is an arc from A to B which is disjoint from $[a_1r_1] \cup \operatorname{cl} W$. Finding some arc connected neighbourhood $W' \subset X$ of r_1, disjoint from $A \cup [a_2u] \cup [ud]$, allows us to find an arc $[r_1z] \subset W'$ and an arc $[zx] \subset W$, where $z \in W \cap W'$. Then, the union $[a_1r_1] \cup [r_1z] \cup [zx]$ contains an arc $[a_1x]$ disjoint from $[a_2d]$, whence it follows that $x \in S$, as needed. This construction is illustrated in Figure A.8d.

If, on the other hand, $u \in [a_1r_1]$, then the same reasoning with the roles of $[a_1r_1]$ and $[a_2r_2]$ exchanged produces disjoint arcs $[a_1d]$ and $[a_2x]$, so that $x \in S$ in this case also. This completes the proof. □

A.8 Conjugacy

Whyburn's theory of *cyclic connectedness*, which we shall encounter in earnest in Section A.10, has had a significant impact on research into the structure of Peano continua, some of which is summarised in [McA66]. This theory, first touched upon in [Why27a] and [Ayr27], is for our current purposes most naturally couched in the language of *conjugacy*, introduced by Kuratowski and Whyburn [KW30] and, according to the cited work, inspired by Moore's discussion of upper semicontinuity in [Moo29].

In essence, conjugacy is an equivalence relation on the set of non-cut points of a Peano continuum, defined as follows.

Definition A.8.1. *Define an equivalence relation $\sim$ on the set of non-cut points of some Peano continuum X by declaring that*

$$x \sim y \textit{ if and only if no point of } X \textit{ separates } x \textit{ and } y.$$

We call this the conjugacy relation *on X, and if $x \sim y$, then the points x and y are said to be* conjugate.

The equivalence class of a point $x \in X$ with respect to $\sim$ is called the conjugacy class *of x, and we denote this by $C_x \subset X$.*

In general, the relation $\sim$ extended to include cut points cannot be an equivalence relation, since it is not typically transitive.[9] Radó and Reichelderfer [RR47] embraced this failure of transitivity and developed a generalized notion of conjugacy based on what they called *cyclic transitivity*, but to our current aims this is only relevant as a curiosity.

What follows is a series of elementary properties of conjugacy classes, adapted from [WD79], which we shall use throughout the remainder of this chapter.

Lemma A.8.2. *If x denotes a non-cut point of some Peano continuum X, then the conjugacy class $C_x \subset X$ is closed.*

Proof. Consider some arbitrary point $y \notin C_x$. By definition, there must exist some point $p \in X$ with a corresponding separation $X - \{p\} = X_1 \mid X_2$, where $x \in X_1$ and $y \in X_2$. Notice that $C_x \subset X_1$, for any point of X_2 is necessarily separated from x by the point p.

Now, we know that X_1 is closed, so that every convergent sequence of points of C_x must have its limit in X_1. In particular, we see that no sequence of points of C_x can possibly converge to the point y, or, phrased differently, that $y \notin \mathrm{cl}_X C_x$. Repeating this reasoning for each point $y \in X - C_x$, we deduce that $(\mathrm{cl}\, C_x) \cap (X - C_x) = \emptyset$–or, more transparently, that C_x is closed. □

Lemma A.8.3. *If x denotes a non-cut point of some Peano continuum X, then each component of $X - C_x$ has precisely one boundary point in C_x.*

Proof. This is trivial if C_x is a singleton, so let us assume that C_x has at least two points. Denoting by Y some component of $X - C_x$, we suppose towards a contradiction that there exist two distinct points $a, b \in C_x \cap \mathrm{fr}\, Y$.

Let U and V denote disjoint, arc connected neighbourhoods in X of the points a and b respectively; then, the union $Y \cup U \cup V$ is a connected open subspace of X, and therefore arc connected (Lemma A.5.4). Thus, there exists some arc $[ab] \subset Y \cup U \cup V$, and such an arc must have some subarc $[a'b'] \subset [ab]$ such that $a', b' \in C_x$ and $]a'b'[\subset Y$.

Now, choose any point $y \in\,]a'b'[$. Having arranged that $]a'b'[\subset Y \subset X - C_x$, we know from Definition A.8.1 that there must exist some point $p \in X$ such that $X - \{p\} = X_1 \mid X_2$ is a separation with $x \in X_1$ and $y \in X_2$. By similar reasoning to that used to prove Lemma A.8.2, we deduce that $C_x - \{p\} \subset X_1$, which implies that at least one[10] of the points a' and b' lies in X_1. r However, if $a' \in X_1$, then $[ay[\subset X_1$; if $b' \in X_1$, then $]yb] \subset X_1$ instead. Either way, this is a contradiction: we know that $y \in X_2$, and a point of X_2 cannot possibly be a boundary point of X_1. □

Lemma A.8.4. *Let K denote some connected subspace of a Peano continuum X. For each non-cut point $x \in X$, the intersection $C_x \cap K$ is connected.*

[9] For instance, if p is a cut point of a Peano continuum X which separates points $x, y \in X$, it is easy to think of configurations in which $x \sim p$ and $p \sim y$, even though $x \not\sim y$.

[10] It is in principle possible that $p = a'$ or $p = b'$.

Proof. We shall argue by contraposition, supposing that $C_x \cap K$ is disconnected for some $K \not\subset C_x$, and showing that K must be disconnected as well. (Notice that the result is immediate when $K \subset C_x$.)

Let $C_x \cap K = A \mid B$ denote some separation. From Lemma A.8.3, we know that each component of $K - C_x$ has precisely one boundary point in C_x, which must lie in either A or B. Thus, we define subspaces

$$\begin{aligned} K_A &= A \cup \bigcup \{\text{components of } K - C_x \text{ with a boundary point in } A\} \text{ and} \\ K_B &= B \cup \bigcup \{\text{components of } K - C_x \text{ with a boundary point in } B\}. \end{aligned} \tag{A.43}$$

By construction, $K_A \cup K_B = K$ and $K_A \cap K_B = \emptyset$. We contend that

$$K_A \cap (\mathrm{cl}_K\, K_B) = (\mathrm{cl}_K\, K_A) \cap K_B = \emptyset, \tag{A.44}$$

for this will imply that K is disconnected, completing the proof.

Towards a contradiction, suppose that there exists some point $y \in K_A \cap (\mathrm{cl}_K\, K_B)$. Such a point is of course a limit of some sequence $(y_n)_{n\in\mathbb{N}}$ of points of K_B, and this limit must lie in A. (To see why, notice that our definition in (A.43) implies that every point of $K_A - A$ belongs to an open subspace disjoint from K_B: namely, a component of $K - C_x$ which has a boundary point in A.)

Consider some neighbourhood $U \subset X$ of y which is disjoint from B. Since the sequence $(y_n)_{n\in\mathbb{N}}$ is eventually in this neighbourhood, we see that U must intersect at least one component of $K - C_x$ with a boundary point in B, and we denote this component by L. However, this implies that L contains a sequence with limit $y \in A$; in other words, L must have boundary points in both A and B alike, and this contradicts Lemma A.8.3. □

Corollary A.8.5. *Every conjugacy class in a Peano continuum is connected.* □

Lemma A.8.6. *If x denotes a non-cut point of some Peano continuum X, then the conjugacy class $C_x \subset X$ is itself a Peano continuum.*

Proof. The metrizability of C_x is trivial, while its compactness and connectedness are demonstrated by Lemma A.8.2 and Corollary A.8.5 respectively. To demonstrate that C_x is also locally connected, we recall the machinery of Property S developed in Section A.6.1.

If $\mathcal{U}$ denotes some open cover of C_x by open subspaces of C_x, then we know that $\mathcal{U}$ is of the form

$$\mathcal{U} = \{V_i \cap C_x \mid i \in I\} \tag{A.45}$$

for some collection $\{V_i\}_{i\in I}$ of open subspaces of X. This implies that the collection

$$\mathcal{V} = \{V_i\}_{i\in I} \cup (X - C_x) \tag{A.46}$$

is an open cover of X, and Lemma A.6.2 requires that such a cover permit some refinement

$$\mathcal{W} = \{W_1, W_2, \ldots, W_n\} \tag{A.47}$$

by finitely many connected subspaces. Using Lemma A.8.4, we see that

$$\mathcal{W}' = \{W_1 \cap C_x, W_2 \cap C_x, \ldots, W_n \cap C_x\} \tag{A.48}$$

is a refinement of our original cover $\mathcal{U}$ by finitely many connected subspaces–but another application of Lemma A.6.2 tells us that C_x must therefore be locally connected. □

Lemma A.8.7. *If x denotes a non-cut point of some Peano continuum X, then the conjugacy class C_x has no cut points.*

Proof. With an eye towards deriving a contradiction, suppose that $C_x - \{p\} = A \mid B$ is a separation for some point $p \in C_x$. Notice first that p and x cannot be the same point, for by hypothesis, $X - \{x\}$ is connected, so that Lemma A.8.4 forces $C_x - \{x\} = C_x \cap (X - \{x\})$ to be connected.

Select some pair of points $a \in A$ and $b \in B$. We propose that p cannot separate a and b in X, and indeed, suppose that there were to exist some separation $X - \{p\} = X_1 \mid X_2$ with $a \in X_1$ and $b \in X_2$. We lose no generality in assuming that $x \in X_1$, and noticing that p therefore separates x and b reveals that $b \notin C_x$. This validates our proposal.

In light of this, denote by Y the component of $X - \{p\}$ which contains both a and b. According to Lemma A.8.4, the intersection

$$Y \cap C_x = (Y - \{p\}) \cap C_x = Y \cap (C_x - \{p\}) \tag{A.49}$$

is connected and contains both a and b–but this is a contradiction, for we chose a and b specifically to lie in different components of $C_x - \{p\}$. □

A.9 End Points and the Whyburn Decomposition

The non-cut points of a given Peano continuum X can naturally be divided into two classes: namely those with non-degenerate conjugacy classes, and those whose conjugacy class consists of that point alone. In this section, we prove that the latter class of points is precisely that of the *end points* of X, and in doing so arrive at Whyburn's decomposition [Why27a] of Peano continua into collections of non-degenerate conjugacy classes and end points, all connected by cut points.

Definition and Lemma A.9.1. *Let X denote some Peano continuum. A non-cut point $x \in X$ is said to be an* end point *of X if there exists a neighbourhood basis $\{U_i\}_{i \in I}$ for x such that $\mathrm{fr}_X\, U_i$ contains precisely one point for each $i \in I$.*

For any non-cut point x of a Peano continuum X, the following are equivalent:

- *$C_x = \{x\}$, and*
- *x is an end point of X.*

Proof. Suppose first that $C_x = \{x\}$, and consider some arbitrary neighbourhood $U \subset X$ of x. By hypothesis, x is a non-cut point of X, so that there must exist some neighbourhood $V \subset X$ of x such that $V \subset U$ and $X - V$ is connected (Lemma A.3.5).

Using Lemma A.5.8, we find some point $y \in X - V$ and some arc $[xy] \subset X$ such that $[xy[\subset V$. Our assumption that C_x is a singleton asserts the existence of some point $z \in X$ separating x and y, and manifestly $z \in [xy[$, for otherwise some component of $X - \{z\}$ would contain the arc $[xy]$.

Thus, let $X - \{z\} = X_1 \mid X_2$ denote some separation with $x \in X_1$ and $y \in X_2$. Having chosen V such that $X - V$ is connected, we find that $X - V \subset X_2$, so that $X_1 \subset V$. Noticing that $\mathrm{fr}_X\, X_1 = \{z\}$, we find that x is an end point of X.

Conversely, suppose that x is an end point of X, and select any point $y \in X - \{x\}$. Finding some open neighbourhood $U \subset X$ of x such that $y \notin \mathrm{cl}\, U$ and $\mathrm{fr}\, U = \{z\}$, we can use Lemma 2.3.3 to deduce that every arc $[xy] \subset X$ meets z. Thus, there exists no arc joining x to y in $X - \{z\}$, so that $y \notin C_x$. □

From the above result, we can see that there exists a natural decomposition of any Peano continuum as follows.

Definition and Lemma A.9.2. *Let X denote some Peano continuum. Each point $x \in X$ is either*

- *A cut point of X;*
- *An end point of X; or*
- *An member of a non-degenerate conjugacy class in X.*

We call this partition of the Whyburn decomposition *of X.* □

A.10 Cyclically Connected Peano Continua

The idea of cyclic connectedness was mentioned briefly in the introduction to Section A.8, but it is only now that we have the necessary machinery in place to extract what we need from it. We start with a definition originating in [Why27a, Why27b, Why27c].

Definition A.10.1. *Let X denote some Peano continuum. If, for each pair of points $x, y \in X$, there exists some Jordan curve $J_{xy} \subset X$ which contains both x and y, then X is said to be* cyclically connected.

The cyclically connected subcontinua of X which are maximal with respect to set-theoretic inclusion are termed the cyclic components *of X.*

Inspecting Definition A.10.1, we deduce immediately that cyclic connectedness precludes the existence of cut points.[11] Less obvious, however, is that the converse is in fact true: every Peano continuum with no cut points is cyclically connected, and it is to a proof of this non-trivial fact that we dedicate this section.

With what we have learned so far, it is not clear how we might directly prove that, if X is some Peano continuum with no cut points, then X is cyclically connected. The issue is that, *a priori*, we have no means of constructing Jordan curves in X, and with the following lemma we move towards addressing this shortcoming.

Lemma A.10.2. *If X denotes some Peano continuum without cut points, then each point of x is interior to some arc. That is, given any point $x \in X$, we can find points $a, b \in X - \{x\}$ for which there exists an arc $[axb] \subset X$.*

Proof. We distinguish two possibilities, and prove the result in each case separately: either the point x may be a cut point of one of its connected neighbourhoods in X, or x may not separate any of its neighbourhoods.

The first of these is more straightforward to deal with. If $U \subset X$ denotes some connected open neighbourhood of x such that $U - \{x\} = A \mid B$ is some separation, let us select points $a \in A$ and $b \in B$. Lemma A.5.4 asserts that U is arc connected, yielding arcs $[ax], [xb] \subset U$; moreover, since a and b belong to different components of $U - \{x\}$, we can deduce that $[ax] \cap [xb] = \{x\}$. It follows that the union $[ax] \cup [xb] = [axb]$ is the arc we seek.

Suppose now that no connected neighbourhood of x has x as a cut point, fix some metric d topologizing X, and choose some arbitrary pair of points $a, b \in X - \{x\}$. We lose no generality in assuming that

$$\min\{d(x, a), d(x, b)\} = 1, \tag{A.50}$$

for if this is not the case, we can simply rescale d by an appropriate constant factor.

Using local connectedness, we are able to find some connected open neighbourhood $U_1 \subset X$ such that $\operatorname{diam} U_1 < \frac{1}{2}$. Additionally, the Hahn-Mazurkiewicz-Sierpiński theorem (Theorem A.6.8) allows us to find some cover $\mathcal{A}_1$ of X by finitely many Peano subcontinua, all of diameter strictly less than $\frac{1}{4}$. We define

$$Y_1 = \bigcup\{A \in \mathcal{A}_1 \mid A \cap U_1 \neq \emptyset\}, \tag{A.51}$$

which, recalling Lemma A.6.3, we see to be a Peano continuum containing x; further, this construction ensures that $\operatorname{diam} Y_1 < 1$, so that in particular $a, b \notin Y_1$.

By construction, Y_1 is a neighbourhood of x in X, so that x is a non-cut point of Y_1. Further, were x to be an end point of Y_1, then x would also be an end point of X–but since we have assumed that X lacks cut points, this is impossible. Appealing to Lemma A.9.2, we deduce that the conjugacy class of x in Y_1 is non-degenerate. We also know from Lemma A.8.6 that this conjugacy class is itself a Peano continuum, which we shall denote by C_1. Lemma A.8.7 allows us to apply the 2-Bogensatz (Lemma A.7.1) to the disjoint closed subspaces $\{a, b\}$ and C_1, producing a pair of disjoint arcs $[aa_1], [bb_1] \subset X$, with $[aa_1] \cap C_1 = \{a_1\}$ and $[bb_1] \cap C_1 = \{b_1\}$.

[11]Indeed, were X some Peano continuum with points $a, b \in X$ separated by some third point $p \in X$, no pair of arcs from a to b in X can possibly be disjoint: all such arcs must pass through the point p.

This construction can now be iterated with a_1, b_1 and C_1 taking the roles of a, b and X respectively. This produces a Peano subcontinuum $C_2 \subset C_1$ such that $\operatorname{diam} C_2 < \frac{1}{2}$, which contains x and has no cut points; likewise, we now have disjoint arcs $[a_1 a_2], [b_1 b_2] \subset C_1$ such that $[a_1 a_2] \cap C_2 = \{a_2\}$ and $[b_1 b_2] \cap C_2 = \{b_2\}$. More generally, we can construct, for each $n \in \mathbb{N}$,

- Peano continua $X = C_0 \supset C_1 \supset \cdots \supset C_n \supset \cdots$ such that $\operatorname{diam} C_n < \frac{1}{n}$ and $x \in \bigcap_{n\in\mathbb{N}} C_n$, all of which lack cut points; and
- Disjoint pairs of arcs $[a_{n-1}a_n], [b_{n-1}b_n] \subset C_n$, each with the property that $[a_{n-1}a_n] \cap C_k = a_k$ and $[b_{n-1}b_n] \cap C_k = b_k$, for $k \in \{n-1, n\}$ where we have identified $a = a_0$ and $b = b_0$.

Defining the unions

$$\begin{aligned} T_a &= \bigcup_{n\in\mathbb{N}} [a_{n-1}a_n], \\ T_b &= \bigcup_{n\in\mathbb{N}} [b_{n-1}b_n] \text{ and} \\ T &= T_a \cup \{x\} \cup T_b, \end{aligned} \tag{A.52}$$

we suggest that T is the arc we seek. Thanks to Lemma A.4.5, it will suffice to prove that T is a continuum with precisely two non-cut points.

First, consider $\operatorname{cl}_X T_a$. Noticing that the sequence $(a_n)_{n\in\mathbb{N}}$ converges to x, we find that $x \in \operatorname{cl}_X T_a$; however, no sequence of points of T_a can have a limit in T_b. To see this, let $(y_n)_{n\in\mathbb{N}}$ denote some sequence in T_a with a limit $y \in T_b$. For some $n \in \mathbb{N}$, it must be the case that $y \in]b_{n-1}, b_n]$, so that $]b_{n-1}, b_{n+1}[\subset T_b$ is a neighbourhood of y in T. The sequence $(y_n)_{n\in\mathbb{N}}$ can never take values in such a neighbourhood, and we conclude that $y_n \not\to y$.

Since it is clear that T_a can have no boundary points in $X - T$, we find that

$$\operatorname{cl}_X T_a = \operatorname{cl}_T T_a = T_a \cup \{x\}, \tag{A.53}$$

with an analogous result for $\operatorname{cl}_X T_b$. Knowing that T_a and T_b are both connected (for otherwise, one of the intervals $\{[a_{n-1}, a_n], [b_{n-1}, b_n]\}_{n\in\mathbb{N}}$ would be disconnected), we deduce that so too are their closures $\operatorname{cl}_T T_a$ and $\operatorname{cl}_T T_b$; thus, we see that T is compact and connected, with x as a cut point.

The same style of reasoning allows us to conclude that every other point of $T - \{a, b\}$ also separates T, and so the proof is complete. □

Notice that, in the presence of cut points, Lemma A.10.2 cannot possibly hold. (Consider, for instance, the point 0 in the closed unit interval $[0, 1]$.)

With Lemma A.10.2, we are able to prove the equivalence between cyclic connectedness and the absence of cut points, which is vital throughout the latter portions of Chapter 2. This proof is clarified conceptually by the introduction of the *three point property*, which is in essence a strengthening of Lemma A.10.2: in a space with this property, not only is every point interior to some arc, but we are free to choose the end points of that arc as well.

Theorem A.10.3. *We say that a Peano continuum X has the* three point property *if, for every trio $x, y, z \in X$ of distinct points, there exists some arc $[xyz] \subset X$.*

For every Peano continuum X, the following are equivalent:

1. *X has no cut points,*
2. *X has the three point property, and*
3. *X is cyclically connected.*

Proof. We shall prove first that an absence of cut points is equivalent to the presence of the three point property, then that any Peano continuum enjoying the three point property must be cyclically connected, and finally that no cyclically connected Peano continuum can have a cut point.

Suppose, then, that X is a Peano continuum without cut points, and choose any three distinct points $x, y, z \in X$. We know from Lemma A.10.2 that there exists some arc $[ayb] \subset X$ for some pair of points $a, b \in X$, and the arc connectedness of X supplies us with arcs $[xa], [bz] \subset X$. The union $[xa] \cup [ayb] \cup [bz]$ then contains an arc $[xyz]$, so that X has the three point property.

Conversely, suppose for the sake of a contradiction that X has the three point property, and that $X - \{p\} = A \mid B$ is some separation for some point $p \in X$. Selecting any pair of points $a \in A$ and $b \in B$, we know that any arc from a to b in X must contain the point p, but this implies that there can exist no arc $[pab] \subset X$. Thus, X is forbidden to have any cut points.

Supposing now that X has the three point property, we suggest that every point of X lies on some Jordan curve in X, and shall use this as a stepping stone towards a proof of the cyclic connectedness of X. Indeed, given any point $x \in X$, we can find some arc $[axb] \subset X$, and knowing from the preceding that X lacks cut points, we can also find some arc $[ab] \subset X - \{x\}$. Then, the union $[axb] \cup [ab]$ must contain some Jordan curve on which the point x lies.

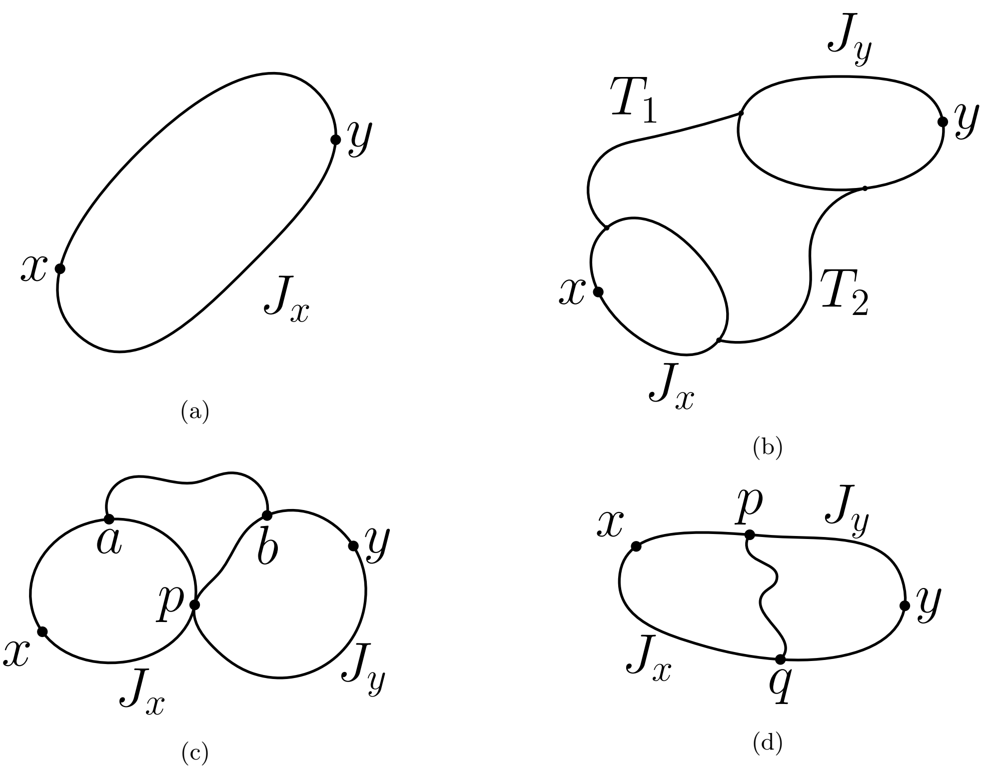

Figure A.9

Thus, given any pair of distinct points $x, y \in X$, we can find a pair of Jordan curves $J_x, J_y \subset X$ such that $x \in J_x$ and $y \in J_y$. From here, we can proceed by exhausting the possible forms of the intersection $J_x \cap J_y$:

- If $J_x = J_y$ (Figure A.9a), then there is nothing further to be done.
- If $J_x \cap J_y = \emptyset$, then an application of the 2-Bogensatz (Lemma A.7.1) allows us to find disjoint arcs T_1, T_2 from J_x to J_y, each meeting J_x and J_y only at their endpoints. The union $J_x \cup J_y \cup T_1 \cup T_2$ then contains a Jordan curve upon which both x and y lie. (Explicitly, denote by A_x the complementary domain of $T_1 \cup T_2$ in J_x which contains x, and A_y that of the same in J_y which contains y. The union $A_x \cup T_1 \cup A_y \cup T_2$ is then the desired Jordan curve; see Figure A.9b.)
- If J_x and J_y meet at precisely one point p, then our knowledge that X has no cut points allows us to find points $a \in J_x - \{p\}$ and $b \in J_y - \{p\}$ for which there exists an arc $[ab] \subset X - \{p\}$, such that $[ab] \cap J_x = \{a\}$ and $[ab] \cap J_y = \{b\}$. Then, there exists a Jordan curve containing both x and y in the union $J_x \cup J_y \cup [ab]$. (Explicitly, if A_x denotes the complementary domain of $\{a, p\}$ in J_x which contains x, and A_y that of $\{b, p\}$ in J_y which contains y, then $A_x \cup [ab] \cup A_y \cup \{p\}$

is a suitable Jordan curve. See Figure A.9c.)

- Finally, if $J_x \cap J_y$ contains multiple points, but nevertheless $J_x \neq J_y$, then we can find points $p, q \in J_x \cap J_y$ such that there exists some arc $[pyq] \subset J_y$, and such that $[pyq] \cap J_x = \{p, q\}$. Denoting by A the complementary domain of $\{p, q\}$ in J_x containing x, we see that the union $A_x \cup [pyq]$ is a Jordan curve containing both x and y. (Figure A.9d.)

We conclude that if X enjoys the three point property, then X must be cyclically connected.

All that remains for us to prove is that a cyclically connected Peano continuum has no cut points, and we shall proceed by contradiction. Suppose that X is a cyclically connected Peano continuum with some cut point $p \in X$. If $a, b \in X - \{p\}$ lie in different components, then it is impossible for there to exist a Jordan curve in X which contains both a and b. Indeed, were J some such curve, then the complement $J - \{a, b\}$ would consist of two disjoint open arcs with ends in different components of $X - \{p\}$–but according to Lemma 2.3.3, any two such arcs must contain the point p, and so cannot possibly be disjoint. □

Appendix B

Alexander Duality

B.1 Overview

In full generality, Alexander duality[1] can be viewed as a relationship between the Čech cohomology groups of a given compact subspace of the n-sphere and the singular homology groups of the complement in the n-sphere of said subspace. This is treated in many algebraic topology textbooks, such as [Spa66], and the necessity of using the Čech cohomology theory in general is discussed in [Mas78].

Here, however, we have no need of the far-reaching generality of Alexander duality as described by the preceding works. In Chapter 1, we use only an extremely restricted form of the result, which we restate below.

Theorem 1.4.1 (Alexander duality)**.** *Let $C \subset S^2$ denote some compact subspace with $n \in \mathbb{N}$ connected components. Then, there exists an isomorphism $H_1\big(S^2 - C\big) \cong \mathbb{Z}^{n-1}$.* □

In this appendix, we give a reasonably self-contained proof of Theorem 1.4.1 by elementary means. Essentially, we first prove an analogous result in the plane, and then we promote this to the sphere by a Mayer-Vietoris argument.

Our proof of the planar form of Theorem 1.4.1 is inspired by Newman's treatment [New85] of a 'homology-like' theory in the plane, and hinges upon the fact that, from the perspective of singular homology, we lose no generality in assuming that paths are the so-called *grid paths* of Definition B.2.2. The first section of this chapter is given over to proving this: in Section B.2.1, we demonstrate that all paths in the plane are path-homotopic to grid paths, and in Section B.2.2, we recall how path homotopy interacts with singular homology.

In Sections B.3 and B.4, we develop the machinery of winding number which enables our proof of the planar form of Theorem 1.4.1; it turns out that a particularly convenient formulation of winding number for these purposes is in terms of covering spaces, whose vital properties we review without proof in Lemma B.3.1. The most important–and most technical–aspect here is the relationship between winding number and homology presented in Lemma B.4.2.

Finally, in Section B.5 we make good on our promises, proving Theorem 1.4.1 with reference to some well-known facts from algebraic topology.

Before proceeding, we establish notation for the concatenation of paths. Particularly, we comment that γ_1 and γ_2 in the below definition are reversed compared to the convention of [Lee00]. (This change was made that the notation for path concatenation might better parallel that for function composition.)

Definition B.1.1. *Let $\gamma_1, \gamma_2\colon [0,1] \to X$ denote some pair of paths in a topological space X such*

[1]So named for Alexander's 1915 work [Ale15], before the advent of algebraic topology in its modern sense.

that $\gamma_1(1) = \gamma_2(0)$. The path in X defined by

$$(\gamma_2 * \gamma_1)(t) = \begin{cases} \gamma_1(2t) & \text{if } t \in \left[0, \frac{1}{2}\right] \text{ and} \\ \gamma_2(2t-1) & \text{if } t \in \left[\frac{1}{2}, 1\right]. \end{cases}$$

is then said to be the concatenation of γ_1 and γ_2.

B.2 Paths in the Plane

Ultimately, our goal in this section is to show that, when working with singular homology, we are free to assume that paths in the plane are of a particularly simple form. We divide our strategy into two pieces: first, we prove in Section B.2.1 that each path in the plane is path-homotopic to a *grid path*, and in Section B.2.2, we prove that if a given pair of paths are path-homotopic, then they must also be homologous.

It will be in Section B.4 that we see the payoff of this work. Roughly speaking, we shall use the results of this section to rephrase a claim about arbitrary singular 2-chains as a claim about grid paths, about which we can then reason combinatorially.

B.2.1 Grid Paths

A particular simple class of path in $\mathbb{R}^2$ is what we shall call the *grid path*, which consists of a concatenation of finitely many vertical and horizontal paths. Especially relevant here is that–at least from a path-homotopical perspective–we do not lose any generality in assuming that an arbitrary path in $\mathbb{R}^2$ is of this form, as we now prove.

Definition B.2.1. *If $\gamma\colon [0,1] \to X$ denotes some path in a topological space X, then the* support of γ *is defined to be the image* $\operatorname{supp}\gamma = \gamma([0,1])$.

Definition B.2.2. *A path in $\mathbb{R}^2$ is said to be a* grid path *if it is the concatenation of finitely many paths, each of which is either vertical or horizontal.*

To prove that an arbitrary path in (an open subspace of) $\mathbb{R}^2$ is path-homotopic[2] to a grid path, we first prove that it is path-homotopic to a piecewise linear one. With this result in hand, it will suffice to prove that every straight-line path in (an open subspace of) $\mathbb{R}^2$ is path-homotopic to a grid path.

Lemma B.2.3. *Path homotopy respects concatenation in the following sense.*

Let $\gamma_1, \gamma_2\colon [0,1] \to X$ denote some pair of paths in a topological space X, for which $\gamma_1(1) = \gamma_2(0)$. If γ_1 is path-homotopic in X to a path γ_1', and γ_2 to a path γ_2', then the concatenations $\gamma_2 * \gamma_1$ and $\gamma_2' * \gamma_1'$ are also path-homotopic in X.

Proof. Let $\varphi\colon [0,1]^2 \to X$ denote some path homotopy from γ_1 to γ_2, and $\psi\colon [0,1]^2 \to X$ some path homotopy from $\gamma_1' \to \gamma_2'$. Then, we see immediately that the function $[0,1]^2 \to X$ acting by

$$(s,t) \mapsto \begin{cases} \varphi(2s, t) & \text{if } s \leq \frac{1}{2} \text{ and} \\ \psi(2s-1, t) & \text{if } s \geq \frac{1}{2} \end{cases} \tag{B.1}$$

is a path homotopy from $\gamma_2 * \gamma_1$ to $\gamma_2' * \gamma_1'$. □

Lemma B.2.4. *Let $\gamma\colon [0,1] \to U$ denote some path, where $U \subset \mathbb{R}^2$ is some open subspace. Then, γ is path-homotopic in U to a piecewise linear path.*

[2]We recall that paths $\gamma_1, \gamma_2\colon [0,1] \to X$ are said to be *path-homotopic in X* if there exists some homotopy $\varphi\colon [0,1]^2 \to X$ from γ_1 to γ_2 such that $\varphi(0,t) = \gamma_1(0) = \gamma_2(0)$ and $\varphi(1,t) = \gamma_1(1) = \gamma_2(1)$ for every $t \in [0,1]$. Perhaps more transparently, a path homotopy is a homotopy between paths which leaves their endpoints fixed throughout.

Proof. The openness of U allows us to find, for each $t \in [0,1]$, some $\varepsilon_t > 0$ such that $B(\gamma(t);\varepsilon_t) \subset U$. Thus, we arrive at an open cover

$$\mathcal{U} = \{B(\gamma(t);\varepsilon_t) \mid t \in [0,1]\} \tag{B.2}$$

of U, from which we construct an open cover

$$\mathcal{V} = \{\gamma^{-1}(U) \,|\, U \in \mathcal{U}\} \tag{B.3}$$

of $[0,1]$.

Now, the cover $\mathcal{V}$ permits some Lebesgue number $\delta > 0$, because $[0,1]$ is compact and metrizable. Let us select some $N \in \mathbb{N}$ such that $\frac{1}{N} < \delta$, and subdivide $[0,1]$ as a union

$$[0,1] = \bigcup_{k=1}^{N} I_k, \text{ where } I_k = \left[\frac{k-1}{N}, \frac{k}{N}\right] \text{ for each } k \in \{1,2,\ldots,N\}. \tag{B.4}$$

By construction, for each $k \in \{1,2,\ldots,N\}$, there must exist some $U_k \in \mathcal{U}$ which contains the path segment $\gamma(I_k)$. From here, our strategy shall be to 'straighten' each segment $\gamma(I_k) \subset U_k$, and then to concatenate the results.

Explicitly, for each $k \in \{1,2,\ldots,N\}$, we define a path $\gamma_k\colon [0,1] \to U_k$ by

$$\gamma_k(t) = \gamma(Nt - k), \tag{B.5}$$

which is essentially nothing more than the restriction of γ to I_k. We know that $\operatorname{supp}\gamma_k \subset U_k$, and we define $L_k\colon [0,1] \to U_k$ as the straight-line path from $\gamma_k(0)$ to $\gamma_k(1)$. The convexity of U_k ensures that γ_k and L_k are path-homotopic in U so that the concatenations $\gamma_N * \gamma_{N-1} * \cdots * \gamma_1$ and $L_N * L_{N-1} * \cdots * L_1$ are also path-homotopic in U, thanks to Lemma B.2.3. Observing that $\gamma_N * \gamma_{N-1} * \cdots * \gamma_1$ is nothing more than a reparametrization of the original path γ completes the proof. ☐

Lemma B.2.5. *Let $L\colon [0,1] \to U$ denote some straight-line path, where $U \subset \mathbb{R}^2$ is some open subspace. Then, L is path-homotopic in U to a grid path.*

Proof. Conceptually, this proof is extremely similar to that of the preceding result. Again, we find some cover $\mathcal{U}$ of $\operatorname{supp} L$ by open balls contained within U, and pull this back through L to an open cover of $[0,1]$ with some Lebesgue number $\delta > 0$; we decompose $[0,1]$ as a union

$$[0,1] = \bigcup_{k=1}^{N} I_k, \text{ where } I_k = \left[\frac{k-1}{N}, \frac{k}{N}\right] \text{ for each } k \in \{1,2,\ldots,N\}. \tag{B.6}$$

We now select for each $k \in \{1,2,\ldots,N\}$ some $U_k \in \mathcal{U}$ such that $L(I_k) \subset U$, and define the path $L_k\colon [0,1] \to U_k$ as the straight-line path from $L\left(\frac{k-1}{N}\right)$ to $L\left(\frac{k}{N}\right)$.

Fixing some $k \in \{1,2,\ldots,N\}$, let (x_k, y_k) denote the coordinates of the centre of the ball U_k, and introduce the notation

$$L_k(0) = (a_k, b_k) \text{ and } L_k(1) = (a'_k, b'_k). \tag{B.7}$$

We can define the following straight-line paths, all of which must by convexity lie within U_k:

- v_1, starting at (a_k, b_k) and ending at (a_k, x_k);
- h_1, starting at (a_k, x_k) and ending at (x_k, y_k);
- v_2, starting at (x_k, y_k) and ending at (x_k, b'_k); and
- h_2, starting at (x_k, b'_k) and ending at (a'_k, b'_k).

Once more using the convexity of U_k, we find that the concatenation $h_2 * v_2 * h_1 * v_1$ is path-homotopic in U_k to L_k, which is a grid path. Since L itself is path-homotopic to the concatenation $L_N * L_{N-1} * \cdots * L_1$, we appeal to Lemma B.2.3 in order to conclude that L is path-homotopic to a grid path. ☐

Corollary B.2.6. *If $\gamma\colon [0,1] \to U$ denotes some path, where $U \subset \mathbb{R}^2$ is some open subspace, then γ is path-homotopic in U to a grid path.* □

B.2.2 Path Homotopy and Homology

Having shown that, as far as homotopy is concerned, we do not lose any generality in restricting our consideration of paths in $\mathbb{R}^2$ to grid paths, we now go on to prove that the same holds true from the point of view of singular homology. This is a consequence of the following result, stating that if two paths are path-homotopic, then they are necessarily homologous.

Lemma B.2.7. *Let $\gamma_1, \gamma_2\colon [0,1] \to X$ denote some pair of path-homotopic paths in a topological space X. Then, γ_1 and γ_2 are homologous.*

Proof. We introduce the notation

$$\gamma_1(0) = \gamma_1(0) = x \text{ and } \gamma_1(1) = \gamma_2(1) = y, \tag{B.8}$$

and let $\varphi\colon [0,1]^2 \to X$ denote some path homotopy from γ_1 to γ_2. Our strategy shall be to exhibit a quotient of $[0,1]^2$ which is homeomorphic to a 2-simplex Δ_2, to which φ descends; this will yield a singular 2-simplex $\Delta_2 \to X$, and we shall prove that $\gamma_1 - \gamma_2$ is homologous to the boundary of this 2-simplex.

To this end, realize a 2-simplex Δ_2 as the convex hull of points $\{(0,0),(1,0),(1,1)\} \subset \mathbb{R}^2$, and define the continuous surjection $\pi\colon [0,1]^2 \to \Delta_2$ by

$$\pi(x,y) = (x - xy, xy) \tag{B.9}$$

As a continuous function from a compact space to a Hausdorff space, the surjection π is closed, and therefore must be a quotient map, which we can interpret as collapsing the leftmost edge of the unit square to a point. (Figure B.1.)

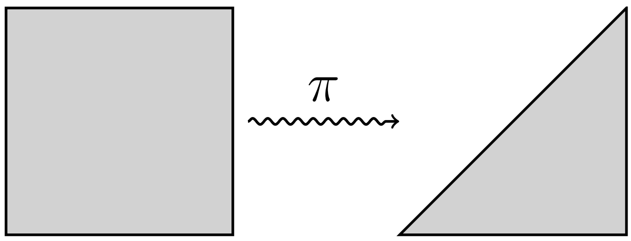

Figure B.1

Since φ is a path homotopy, we know that $\varphi(0,t) = x$; in other words, φ descends through the quotient map π, inducing a singular 2-simplex $\widetilde{\varphi}\colon \Delta_2 \to X$.

Notice that

$$\partial\widetilde{\varphi} = \kappa_y - \gamma_1 + \gamma_2, \tag{B.10}$$

where we use κ_y to denote the constant path at $y \in X$. Since $\partial\widetilde{\varphi}$ and κ_y are of course both boundaries, we deduce that γ_1 and γ_2 are indeed homologous. □

We can also show that the concatenation of paths interacts in a natural way with the notion of homology.

Lemma B.2.8. *Let $\gamma_1, \gamma_2\colon [0,1] \to X$ denote paths in a topological space X for which $\gamma_1(1) = \gamma_2(0)$. Then, the concatenation $\gamma_2 * \gamma_1$ is homologous to the singular 2-chain $\gamma_1 + \gamma_2$.*

Proof. Realize a 2-simplex $\Delta_2 \subset \mathbb{R}^2$ as the convex hull of the vertices $(0,0)$, $(1,0)$ and $(1,1)$, and define a function $\Sigma\colon \Delta_2 \to X$ by

$$\Sigma(x,y) = \begin{cases} \gamma_1(x+y) & \text{if } x+y \leq 1, \text{ and} \\ \gamma_2(x+y-1) & \text{if } x+y \geq 1. \end{cases} \tag{B.11}$$

Notice that this is well defined, since when $x + y = 1$,

$$\gamma_1(x + y) = \gamma_1(0) = \gamma_2(1) = \gamma_2(x + y - 1). \tag{B.12}$$

The continuity of Σ is manifest, and we therefore have a singular 2-simplex $\Sigma\colon \Delta_2 \to X$, the boundary of which we now compute.

We define the face maps $F_1, F_2, F_3\colon [0, 1] \to \Delta_2$ as follows:

$$\begin{aligned} F_1(t) &= (0, 1 - t), \\ F_2(t) &= (t, t) \text{ and} \\ F_3(t) &= (1 - t, 1). \end{aligned} \tag{B.13}$$

Comparing (B.13) to (B.11) reveals that

$$\begin{aligned} \partial\Delta_2 &= (\Delta_2 \circ F_1) - (\Delta_2 \circ F_2) + (\Delta_2 \circ F_3) \\ &= \gamma_1 - (\gamma_2 * \gamma_1) + \gamma_2, \end{aligned} \tag{B.14}$$

whence we conclude that $\gamma_1 + \gamma_2$ and $\gamma_2 * \gamma_1$ are indeed homologous. □

B.3 Covering Spaces and Winding Number

It is convenient for us to couch our definition of winding number in the language of covering spaces, freely making use of the equivalence between $\mathbb{R}^2$ and $\mathbb{C}$. We recall the relevant definitions and properties without proof, referring to [Lee00] for a more detailed treatment.

Definition and Lemma B.3.1. *Let $\pi\colon C \to X$ denote some continuous surjection between topological spaces. If there exists some subspace $U \subset X$ such that*

- *The fibre $\pi^{-1}(U)$ is a disjoint union of some collection $\{\widetilde{U}_i\}_{i \in I}$ of open subspaces of C, and*
- *For each $i \in I$, the restriction $\pi|_{\widetilde{U}_i}\colon \widetilde{U}_i \to U$ is a homeomorphism,*

then we say that U is evenly covered *(by π).*

If every point of x has some neighbourhood in X which is evenly covered by π, then we call π a covering map *(with base space X), and say that C is a* covering space *of X. Suppose that $\pi\colon C \to X$ is a covering map in what follows.*

Given some continuous function $f\colon Y \to X$ between topological spaces, a lift of f *(through π) is a continuous function $\widetilde{f}\colon Y \to C$ such that $\pi \circ \widetilde{f} = f$.*

If Y is connected, and if $\widetilde{f}_1$ and $\widetilde{f}_2$ are two lifts of f such that $\widetilde{f}_1(y) = \widetilde{f}_2(y)$ for some $y \in Y$, then $\widetilde{f}_1 = \widetilde{f}_2$.

If $\gamma\colon [0, 1] \to X$ denotes some path, let $x = \gamma(0)$, and select any point $c \in \pi^{-1}(x)$. There then exists precisely one lift $\widetilde{\gamma}$ of γ through π such that $\widetilde{\gamma}(0) = x$.

Let $\gamma_1, \gamma_2\colon [0, 1] \to X$ denote some pair of path-homotopic paths. If $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$ are lifts of γ_1 and γ_2 respectively through π, both of which start at the same point, then $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$ are path-homotopic in C. □

Lemma B.3.2. *The exponential map* $\exp\colon \mathbb{C} \to \mathbb{C} - \{0\}$ *is a covering map.* □

In terms of Definition B.3.1 and Lemma B.3.2, we define winding number as follows. Part of the reason we choose to frame this definition in terms of covering spaces is the ease with which Corollary B.3.4 follows from the fact that path homotopies lift to covering spaces.

Definition and Lemma B.3.3. *Let $\gamma\colon [0, 1] \to \mathbb{C} - \{0\}$ denote some path, and $\widetilde{\gamma}\colon [0, 1] \to \mathbb{C}$ some lift of γ through the exponential map. The* winding number of γ about 0 *is defined as*

$$n_\gamma(0) = \frac{1}{2\pi i}(\widetilde{\gamma}(1) - \widetilde{\gamma}(0)).$$

This is well defined, in the sense that the value of $n_\gamma(0)$ does not depend on our choice of lift.

Suppose more generally that $\gamma\colon [0,1] \to \mathbb{C}$ is some path, and select some point $z \in \mathbb{C} - \operatorname{supp}\gamma$. The winding number of γ about z *is defined to be*

$$n_\gamma(z) = n_{\gamma - z}(0), \tag{B.15}$$

where in an abuse of notation we denote by $\gamma - z$ the path defined by $t \mapsto \gamma(t) - z$.

Proof. Suppose that $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$ are two lifts of γ through the exponential map. Then, for each $t \in [0,1]$, we know that

$$\exp(\widetilde{\gamma}_1(t) - \widetilde{\gamma}_2(t)) = \gamma(t) - \gamma(t) = 0, \tag{B.16}$$

whence we deduce the existence of some function $f\colon [0,1] \to \mathbb{Z}$ such that

$$\widetilde{\gamma}_1(t) - \widetilde{\gamma}_2(t) = 2i\pi f(t). \tag{B.17}$$

Inspecting (B.17), we see that the function f must be continuous; since $\mathbb{Z}$ is totally disconnected, it follows that f must in fact be a constant function with some value $k \in \mathbb{Z}$. We can now observe that

$$\begin{aligned}(\widetilde{\gamma}_1(1) - \widetilde{\gamma}_1(0)) - (\widetilde{\gamma}_2(1) - \widetilde{\gamma}_2(0)) &= (\widetilde{\gamma}_2(1) - \widetilde{\gamma}_1(1)) - (\widetilde{\gamma}_2(0) - \widetilde{\gamma}_1(0)) \\ &= 2ik\pi - 2ik\pi \\ &= 0,\end{aligned} \tag{B.18}$$

from which we conclude that the winding number $n_\gamma(0)$ is indeed indeed independent of our choice of lift. □

Corollary B.3.4. *Winding numbers are path homotopy invariant, in the sense that if we have some pair $\gamma_1, \gamma_2\colon [0,1] \to \mathbb{C}$ of path-homotopic paths, then for every point $z \in \mathbb{C} - (\operatorname{supp}\gamma_1 \cup \operatorname{supp}\gamma_2)$, we have that $n_{\gamma_1}(z) = n_{\gamma_2}(z)$*

Proof. It will suffice for us to prove the claim for winding numbers about the origin. We know from Definition B.3.1 that we can find a pair of path-homotopic lifts $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$ of γ_1 and γ_2 respectively through the exponential map, and in particular these lifts must share their end points. The result now follows from Definition B.3.3 of the winding number. □

We conclude this section with a pair of lemmas describing the behaviour of winding numbers. Of particular importance in the sequel is Lemma B.3.6, which gives some properties of the winding numbers of loops which will be essential in our proof of Theorem 1.4.1.

Lemma B.3.5. *The winding number of any path varies continuously. That is, given any path $\gamma\colon [0,1] \to \mathbb{C}$, the function $n_\gamma\colon \mathbb{C} - \operatorname{supp}\gamma \to \mathbb{C}$ is continuous.*

Proof. Consider any sequence of points $(z_n)_{n\in\mathbb{N}}$ in $\mathbb{C} - \operatorname{supp}\gamma$ with a limit $z \in \mathbb{C} - \operatorname{supp}\gamma$. It will suffice for us to prove that $n_\gamma(z_n) \to n_\gamma(z)$.

Adopting the notational convention of Definition B.3.3, let $\widetilde{\gamma}$ denote any lift of $\gamma - z$ through the exponential map, and $\widetilde{\gamma}_n$ any lift of $\gamma - z_n$ through the same, for each $n \in \mathbb{N}$. Unravelling our definition of winding number, we see that

$$\begin{aligned}\exp(2i\pi(n_\gamma(z) - n_\gamma(z_n))) &= \exp((\widetilde{\gamma} - \widetilde{\gamma}_n)(1) - (\widetilde{\gamma} - \widetilde{\gamma}_n)(0)) \\ &= \frac{(\exp\circ\widetilde{\gamma})(1)}{(\exp\circ\widetilde{\gamma})(0)} \cdot \frac{(\exp\circ\widetilde{\gamma}_n)(0)}{(\exp\circ\widetilde{\gamma}_n)(1)} \\ &= \frac{\gamma(1) - z}{\gamma(0) - z} \cdot \frac{\gamma(0) - z_n}{\gamma(1) - z_n}\end{aligned} \tag{B.19}$$

Now, (B.19) in the limit $n \to \infty$. Having assumed that $z_n \to z$, and that neither the sequence $(z_n)_{n\in\mathbb{N}}$ nor the point z lie in the support of γ, we see that

$$\lim_{n\to\infty} \exp(2i\pi(n_\gamma(z) - n_\gamma(z_n))) = 1, \tag{B.20}$$

which can be true only if

$$\lim_{n\to\infty}(n_\gamma(z) - n_\gamma(z_n)) = 0. \tag{B.21}$$

This completes the proof. □

Lemma B.3.6. *Let $\sigma\colon [0,1] \to \mathbb{C}$ denote some loop. Then,*

- *For each $z \in \mathbb{C} - \operatorname{supp}\sigma$, the winding number $n_\sigma(z)$ is an integer;*
- *The winding number n_σ is constant on each path component of $\mathbb{C} - \operatorname{supp}\sigma$, and*
- *If z is a point of the unbounded component of $\mathbb{C} - \operatorname{supp}\sigma$, then $n_\sigma(z) = 0$.*

Proof. For the first claim, let $\widetilde{\sigma}$ denote any lift of $\sigma - z$ through the exponential map, where $z \notin \operatorname{supp}\sigma$, and where we use the notational convention of Definition B.3.3. By hypothesis, $\sigma(0) = \sigma(1)$, so that

$$\begin{aligned}\exp(2i\pi n_\sigma(z)) &= \exp(\widetilde{\sigma}(1) - \widetilde{\sigma}(0)) \\ &= \frac{\sigma(1)}{\sigma(0)} \\ &= 1.\end{aligned} \tag{B.22}$$

It follows that $n_\sigma(z) \in \mathbb{Z}$.

Now, consider any pair of points $z, w \in \mathbb{C} - \operatorname{supp}\sigma$, and denote by $\gamma\colon [0,1] \to \mathbb{C} - \operatorname{supp}\sigma$ some path between them. Combining the above with Lemma B.3.5, we see that the composition $n_\sigma \circ \gamma\colon [0,1] \to \mathbb{Z}$ is continuous–but since $\mathbb{Z}$ is totally disconnected, this can only be true if n_σ is constant on $\operatorname{supp}\gamma$. From this, we can deduce that n_σ is constant on each path component of $\mathbb{C} - \operatorname{supp}\sigma$.

For the final claim, it will be enough to exhibit at least one point of the unbounded component of $\mathbb{C} - \operatorname{supp}\sigma$ at which the winding number n_σ vanishes. Knowing that $\operatorname{supp}\sigma$ is compact, and in particular bounded, we can find some $R > 0$ such that $\operatorname{supp}\sigma \subset B(0;R)$; selecting any point $z \in \mathbb{C} - B(0;R)$, we propose that $n_\sigma(z) = 0$.

Indeed, we know that $\operatorname{supp}(\sigma - z) \subset B(-z;R)$, and our choice of z ensures that $0 \notin B(-z;R)$. Knowing that $B(-z;R)$ is simply connected, we find that the path $\sigma - z$ is path-homotopic in $\mathbb{C} - \{0\}$ to the constant path at $\sigma(0) - z$. The lifts of such a path through the exponential map are nothing more than the constant paths at each point of the fibre $\exp^{-1}(\sigma(0) - z)$, so that an invocation of Corollary B.3.4 allows us to conclude that $n_\sigma(z) = 0$. □

B.4 Winding Number and Homology

The following lemma, adapted from [BG91], will prove instrumental in our proof of Theorem 1.4.1. In particular, it will allow us to prove that a homomorphism defined on a group of cycles descends to the quotient, yielding an isomorphism defined on a homology group.

Definition B.4.1. *Let $c = \sum_{i=1}^n \alpha_i\gamma_i$ denote some singular 1-chain with support in $\mathbb{R}^2$. For each point $x \in \operatorname{supp} c$, we define the* winding number of c about x *to be the sum*

$$n_c(x) = \sum_{i=1}^n \alpha_i n_{\gamma_i}(x).$$

Lemma B.4.2. *Let c denote any singular 1-cycle with support in some open subspace $U \subset \mathbb{R}^2$. The following are equivalent:*

- *There exists some singular 2-chain d with support in U such that $c = \partial d$, and*
- *The winding number n_c vanishes on all of $\mathbb{R}^2 - U$.*

Proof. If the 1-cycle c is the boundary of some singular 2-chain with support in U, then every point of $\mathbb{R}^2 - U$ is *a fortiori* a point of the unbounded component of $\mathbb{R}^2 - \operatorname{supp} c$. It follows from Lemma B.3.6 that $n_c(\mathbb{R}^2 - U) = \{0\}$, establishing one implication.

Conversely, suppose that $n_c(\mathbb{R}^2 - U) = \{0\}$ for the singular 1-chain $c = \sum_{i=1}^N \alpha_i \gamma_i$; we wish to construct some singular 2-chain d, supported in U, such that $c = \partial d$. We start by observing that, in light of Corollary B.2.6 and Corollary B.3.4, we lose no generality in assuming that each of the paths $\{\gamma_i\}_{i=1}^N$ is either vertical or horizontal.

Let $\{x_i\}_{i=1}^{n_x}$ denote the set of x-coordinates of end points of the paths $\{\gamma_i\}_{i=1}^N$, and $\{y_i\}_{i=1}^{n_y}$ the set of y-coordinates of the same.[3] Order these sets such that

$$x_1 < x_2 < \cdots < x_{n_x} \text{ and } y_1 < y_2 < \cdots < y_{n_y}, \tag{B.23}$$

and define, for each $i \in \{1, 2, \ldots, n_x - 1\}$ and for each $j \in \{1, 2, \ldots, n_y - 1\}$, the rectangle

$$R_{ij} = [x_i, x_{i+1}] \times [y_j, y_{j+1}] \subset \mathbb{R}^2, \tag{B.24}$$

the centre of which we denote by r_{ij}. (Figure B.2.) Notice that each rectangle R_{ij} can be realized as the support of a singular 2-chain

$$S_{ij} = \tau_{ij} + \tau'_{ij}, \tag{B.25}$$

where the singular 2-simplices τ_{ij} and τ'_{ij} are defined as in Figure B.3.

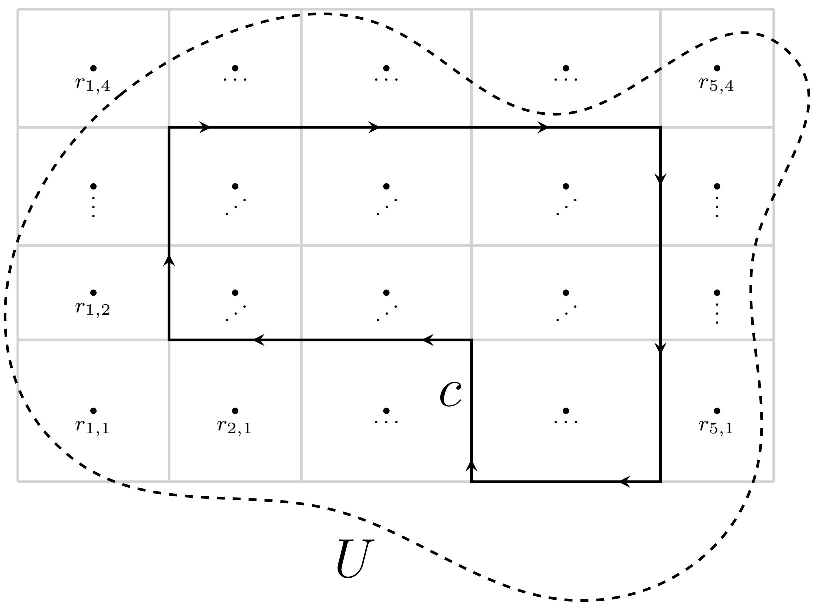

Figure B.2

We propose if $n_c(r_{ij}) \neq 0$, then $R_{ij} \subset U$, and prove this by contradiction. Indeed, suppose that there exists some point $a \in R_{ij} - U$ for some rectangle R_{ij}; if we can prove that there exists some path in $\mathbb{R}^2 - \operatorname{supp} c$ which connects r_{ij} to a, then Lemma B.3.5 will tell us that $n_c(r_{ij})$ vanishes, as hoped.

We denote by $\gamma\colon [0,1] \to R_{ij}$ the straight-line path from r_{ij} to a, which we know by convexity to be supported in R_{ij}. We chose a not to lie in U, which means that $a \notin \operatorname{supp} c$; knowing also that $\gamma([0,1[) \subset \operatorname{int} R_{ij}$, we conclude that $\operatorname{supp} \gamma$ and $\operatorname{supp} c$ are disjoint. It follows that r_{ij} and a belong to the same path component of $\mathbb{R}^2 - \operatorname{supp} \sigma$, justifying our proposal.

[3] Notice that, in general, n_x and n_y are unequal, and either may be strictly less than N: in general, several of the paths $\{\gamma_i\}_{i=1}^N$ will have end points which share an x- or y-coordinate.

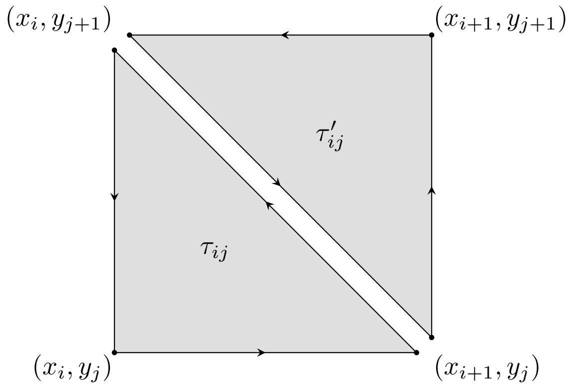

Figure B.3

The upshot of this is that we can define a singular 2-chain

$$d = \sum_{n_c(r_{ij}) \neq 0} n_c(r_{ij}) S_{ij}, \tag{B.26}$$

and we propose that our original 1-chain c is precisely the boundary of d.

Let us consider the formal difference $c - \partial d$. Our choice of c, together with our definition of d, allows us to write this in the form

$$c - \partial d = \sum_{k=1}^{M} \alpha_k \nu_k, \tag{B.27}$$

where each of the paths $\{\nu_k\}_{k=1}^{M}$ is either vertical or horizontal. Moreover, no generality is lost in assuming that, whenever $k \neq l$, the intersection $(\operatorname{supp} \nu_k) \cap (\operatorname{supp} \nu_l)$ is either empty, or contains one or both end points of the paths ν_k and ν_l. Our approach from here shall be to prove that the coefficients $\{\alpha_k\}_{k=1}^{M}$ all vanish, which will imply that $c = \partial d$, so completing the proof.

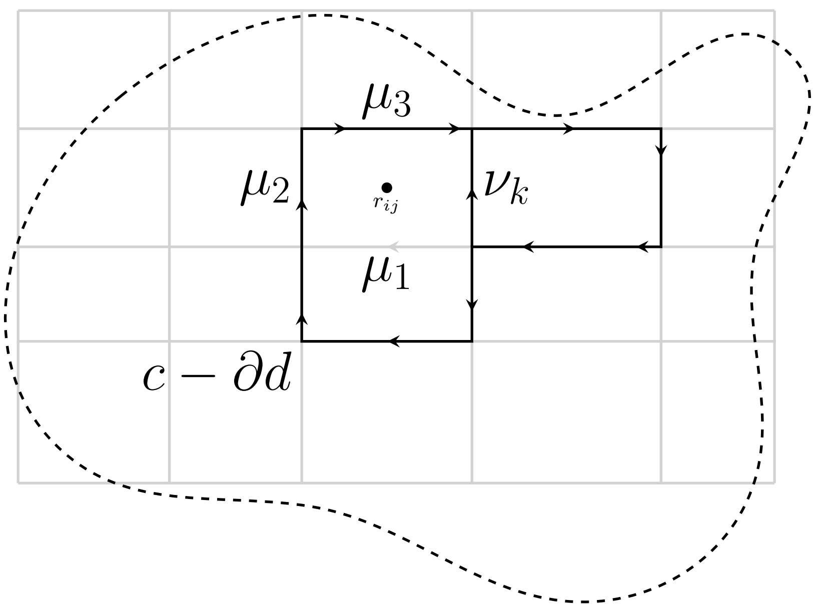

Figure B.4

Suppose that, for some $k \in \{1, 2, \ldots, M\}$, the support of the path ν_k comprises the right-hand edge of some rectangle R_{ij}, and the left-hand edge of the rectangle $R_{i+1,j}$, as depicted schematically in

Figure B.4. Adopting the notation of Figure B.4, we know that

$$\partial R_{ij} = \mu_1 + \mu_2 + \mu_3 \pm \nu_k, \tag{B.28}$$

with the sign of the final term depending on the orientation of ν_k, and from this we can immediately deduce that the linear combination $c - \partial d \mp \alpha_k \partial R_{ij}$ contains no term involving ν_k. In particular, the support of $c - \partial d \mp \alpha_k \partial R_{ij}$ does not intersect the straight-line path from r_{ij} to $r_{i+1,j}$; recalling Lemma B.3.6 and Definition B.4.1, we arrive at the equality

$$n_{c-\partial d}(r_{ij}) \mp \alpha_k n_{\partial R_{ij}}(r_{ij}) = n_{c-\partial d}(r_{i+1,j}) \mp \alpha_k n_{\partial R_{ij}}(r_{i+1,j}). \tag{B.29}$$

The winding number of ∂R_{ij} about r_{ij} and $r_{i+1,j}$ is readily evaluated, allowing us to rewrite (B.29) in the form

$$\mp \alpha_k = n_{c-\partial d}(r_{i+1,j}) - n_{c-\partial d}(r_{ij}). \tag{B.30}$$

However, our definition of the chain d in (B.26) ensures that the right-hand side of (B.30) vanishes, whence we conclude that $\alpha_k = 0$ in this case. This reasoning, appropriately modified, reveals that α_k also vanishes when $\operatorname{supp} \nu_k$ comprises the top edge of some rectangle R_{ij} and the bottom edge of some rectangle $R_{i,j+1}$.

Now, suppose instead that all points of $\operatorname{supp} \nu_k$ have x-coordinate x_0, so that $\operatorname{supp} \nu_k$ is the left-hand edge of some rectangle R_{0j}; again, we know that the linear combination $c - \partial d \mp \alpha_k \partial R_{0j}$ has no term involving ν_k. If we choose some point $z \in U$ with x-coordinate strictly less than x_0 and y-coordinate equal to that of r_{0j}, then it follows that the straight-line path from r_{0j} to z does not meet $\operatorname{supp}(c - \partial d \mp \alpha_k \partial R_{0j})$, and so we deduce (once more via Lemma B.3.6 and Definition B.4.1) that

$$\begin{aligned} n_{c-\partial d}(r_{0j}) \mp \partial \alpha_k n_{\partial R_{0j}}(r_{0j}) &= \mp \alpha_k \\ &= n_{c-\partial d}(z) \mp \partial \alpha_k n_{\partial R_{0j}}(z). \end{aligned} \tag{B.31}$$

However, our choice of z ensures that this point belongs to the unbounded components of both $\mathbb{R}^2 - \operatorname{supp}(c - \partial d)$ and $\mathbb{R}^2 - \operatorname{supp}(\partial R_{0j})$, so that a final invocation of Lemma B.3.6 tells us that $\alpha_k = 0$ in this case also.

Suitably adapted reasoning further indicates that α_k also vanishes when every point of $\operatorname{supp} \nu_k$ has x-coordinate x_{n_k}, y-coordinate y_0, or y-coordinate y_{n_y}. Thus, we deduce that $c = \partial d$, which completes the proof. □

B.5 Alexander Duality

We need four results before we are ready to prove Theorem 1.4.1, but we quote three without proof in order to obviate the need for a lengthy detour into the realms of algebraic topology. The first of these is the standard result that the second homology group of any open subspace of the plane must be trivial, and is proved in [Vic94]; the second two, due to Hurewicz [Hur35] and Johansson [Joh31], concern themselves with the fundamental groups of surfaces, with modern treatments provided in [Lee00] and [Sti93] respectively.

The fourth result, by contrast, involves the topology of the plane, and we prove it explicitly. Heuristically, this lemma states that if we are given some compact subspace of the plane with only finitely many components, then we can always 'draw a loop separating any one of those components from all of the others'.

Lemma B.5.1. *If $U \subset \mathbb{R}^2$ denotes some open subspace, then the second singular homology group $H_2(U)$ is the trivial group.*

Lemma B.5.2 (Hurewicz)**.** *If X denotes some path connected topological space with some distinguished base point $x \in X$, then the first singular homology group $H_1(X)$ is isomorphic to the Abelianization of the fundamental group $\pi_1(X, x)$.* □

Lemma B.5.3 (Johansson)**.** *The fundamental group of any non-compact 2-manifold is free.* □

Lemma B.5.4. *Let $C \subset \mathbb{R}^2$ denote some compact subspace with finitely many components, and A one of these components. There exists a Jordan curve $J \subset \mathbb{R}^2 - C$ such that A lies in the bounded complementary domain of J in $\mathbb{R}^2$, while every other component of C lies in the unbounded complementary domain of the same.*

Proof. If C is connected, then our claim follows immediately from the knowledge that a compact subspace of the plane is necessarily bounded.

If C instead fails to be connected, then let $\{A, B_1, B_2, \ldots, B_n\}$ denote the set of components of C. For each $i \in \{1, 2, \ldots, n\}$, we define

$$\delta_i = \inf\{\|a - b\| \mid a \in A \text{ and } b \in B_i\} \tag{B.32}$$

to be the distance from A to B_i, and choose

$$\delta = \min\{\delta_1, \delta_2, \ldots, \delta_n\} \tag{B.33}$$

to be the least of these distances. Notice that $\delta > 0$.

Choosing some $\varepsilon \in \left]0, \frac{\delta}{2\sqrt{2}}\right[$, we tile the plane with the squares

$$S_{ij} = [i - \varepsilon, i + \varepsilon] \times [j - \varepsilon, j + \varepsilon] \text{ where } i, j \in \mathbb{Z}, \tag{B.34}$$

roughly sketched in Figure B.5a, and define the subspace

$$X = \bigcup\{S_{ij} \mid i, j \in \mathbb{Z} \text{ and } S_{ij} \cap A \neq \emptyset\} \subset \mathbb{R}^2, \tag{B.35}$$

as indicated in Figure B.5b. Letting X' denote the union of X with all of its bounded complementary domains[4] in the plane, we propose that $J = \operatorname{fr} X'$ (Figure B.5c) is precisely the Jordan curve we seek.

Certainly, J is a Jordan curve.[5] Moreover, we lose no generality in assuming that J fails to meet A, for if this is not the case, then we can simply replace ε in (B.34) by an appropriate value in $\left]\varepsilon, \frac{\delta}{2\sqrt{2}}\right[$. Our construction of J ensures that A lies in the bounded component of $\mathbb{R}^2 - J$, while every other component of C must lie in the unbounded component of the same, and so we conclude that the Jordan curve J has the desired properties. □

With what has been established up to this point, we can complete our proof of Theorem 1.4.1 in two steps. First, we prove an appropriate analogue of Theorem 1.4.1 in the plane with Lemma B.5.5; then, we transfer this to the sphere via a Mayer-Vietoris argument, so completing the proof.

Lemma B.5.5. *Let $C \subset \mathbb{R}^2$ denote some compact subspace with $n \in \mathbb{N}$ connected components. The first singular homology group $H_1(\mathbb{R}^2 - C)$ is isomorphic to $\mathbb{Z}^n$.*

Proof. Let $(G, +)$ denote the group of locally constant $\mathbb{Z}$-valued functions on C, where the group operation acts by pointwise addition. We readily observe that $G \cong \mathbb{Z}^n$, for if we let $\{C_1, C_2, \ldots, C_n\}$ denote the set of components of C, then G is freely generated by the functions

$$\{f_1, f_2, \ldots, f_n\}, \text{ where } f_i(x) = \begin{cases} 1 & \text{if } x \in C_i \text{ and} \\ 0 & \text{otherwise.} \end{cases} \tag{B.36}$$

Our strategy shall be to define a group homomorphism $Z_1(\mathbb{R}^2 - C) \to G$, where $Z_1(\mathbb{R}^2 - C)$ denotes the group of singular 1-cycles supported in $\mathbb{R}^2 - C$, and then to show that this homomorphism descends to an isomorphism $H_1(\mathbb{R}^2 - C) \cong G$.

[4] Sometimes, such a union is referred to as the *filling* of X.

[5] There is a subtlety which we have elided here. It is in principle possible that some pair of squares comprising X' intersect only at a vertex, which would prevent $\operatorname{fr} X'$ from being a Jordan curve. However, this can be rectified by subdividing each of the offending squares into nine congruent smaller squares, and using the boundary of what results.

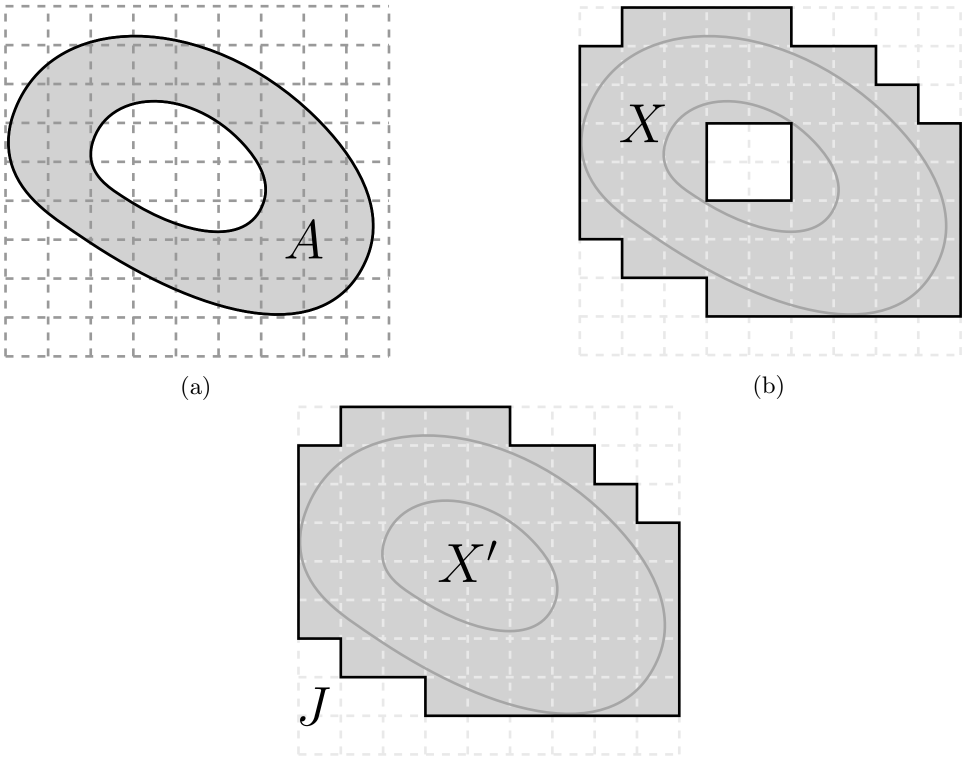

(a) (b) (c)

Figure B.5

To this end, we define a function $\varphi\colon Z_1(\mathbb{R}^2 - C) \to G$ by

$$\varphi(c)(x) = n_c(x), \tag{B.37}$$

essentially acting to map a chain $c \in Z_1(\mathbb{R}^2 - C)$ to its winding number. Recalling Lemma B.3.6 reassures is that such a function is at least well defined. That φ is also a homomorphism of groups is a direct consequence of Definition B.4.1; explicitly, given some pair of chains $c_1, c_2 \in Z_1(\mathbb{R}^2 - C)$, we find that

$$\begin{aligned} \varphi(c_1 + c_2)(x) &= n_{c_1+c_2}(x) \\ &= n_{c_1}(x) + n_{c_2}(x) \\ &= (\varphi(c_1) + \varphi(c_2))(x), \end{aligned} \tag{B.38}$$

as needed.

Our next task is to show that φ is constant on homology classes, so that it descends to a homomorphism $\widetilde{\varphi}\colon H_1(\mathbb{R}^2 - C) \to G$. Indeed, if we select some arbitrary pair of homologous chains $c_1, c_2 \in Z_1(\mathbb{R}^2 - C)$, we know from Lemma B.4.2 that the winding number $n_{c_1 - c_2}$ must vanish on C. It follows that $\varphi(c_1 - c_2) \in G$ must be the constant function $C \to \{0\}$, so that $\varphi(c_1) = \varphi(c_2)$.

Finally, it falls to us to demonstrate that our homomorphism $\widetilde{\varphi}\colon H_1(\mathbb{R}^2 - C) \to G$ is in fact an isomorphism. To see that $\widetilde{\varphi}$ is injective, suppose that $\widetilde{\varphi}([c])$ is the constant function $C \to \{0\}$ for some homology class $[c] \in H_1(\mathbb{R}^2 - C)$; then, $\widetilde{\varphi}(c)$ is also the constant function $C \to \{0\}$. In other words, the winding number n_c vanishes on C, and an appeal to Lemma B.4.2 allows us to deduce that c is a boundary. It follows that $\widetilde{\varphi}([c])$ is the identity of G if and only if $[c]$ is the identity of $H_1(\mathbb{R}^2 - C)$, establishing injectivity.

For surjectivity, fix any $i \in \{1, 2, \ldots, n\}$ and consider the corresponding component C_i of C. Using Lemma B.5.4, we can find some Jordan curve $J \subset \mathbb{R}^2 - C$ such that C_i is the only component of C contained within the bounded component of $\mathbb{R}^2 - J$. If $\sigma\colon [0,1] \to J$ denotes some parametrization of J, then Lemma B.3.6 asserts that $\widetilde{\varphi}([\sigma]) = f_i$, where f_i denotes the generator defined in (B.36). Thus, the image of $\widetilde{\varphi}$ contains every generator of G, and hence surjectivity follows. □

Proof of Theorem 1.4.1. Distinguishing some point $\infty \in S^2 - C$, we define two subspaces of $S^2 - C$; namely, some open ε-ball $A \subset S^2 - C$ centred at ∞, and $B \subset S^2 - C$, the complement in $S^2 - C$ of the open $\frac{\varepsilon}{2}$ ball centred at ∞. Ultimately, we wish to examine the Mayer-Vietoris sequence of the triad $\big(S^2 - C, A, B\big)$, but before doing so, we determine the homotopy types of A, B and $A \cap B$.

Clearly, A is contractible. If we denote by $\Pi\colon S^2 - \{\infty\} \to \mathbb{R}^2$ the stereographic projection with projection point ∞, then $\Pi(B) \subset \mathbb{R}^2$ is the complement of $\Pi(C)$ in some closed disc; such a subspace is a deformation retract of $\mathbb{R}^2 - \Pi(C)$, and therefore is of the same homotopy type. Finally, the same stereographic projection reveals that $A \cap B$ is homeomorphic to an annulus in the plane, so that this intersection has the homotopy type of the circle.

With these observations in hand, the tail of the Mayer-Vietoris sequence of the triad $\big(S^2 - C, A, B\big)$ reads

$$\begin{array}{l} \cdots \longrightarrow H_2\big(S^2 - C\big) \longrightarrow \mathbb{Z} \longrightarrow H_1\big(\mathbb{R}^2 - \Pi(C)\big) \\ \longrightarrow H_1\big(S^2 - C\big) \longrightarrow \mathbb{Z} \longrightarrow \mathbb{Z} \oplus H_0\big(\mathbb{R}^2 - \Pi(C)\big) \longrightarrow H_0\big(S^2 - C\big) \longrightarrow 0. \end{array} \tag{B.39}$$

Now, we can see by stereographically projecting with any point of C as our projection point that $S^2 - C$ is a homeomorph of some open subspace of $\mathbb{R}^2$, so that its second singular homology group is trivial (Lemma B.5.1); meanwhile, Lemma B.5.5 asserts that $H_1\big(\mathbb{R}^2 - \Pi(C)\big) \cong \mathbb{Z}^n$. Yet another stereographic projection about ∞ allows us to see that the sets of path components of $S^2 - C$ and $\mathbb{R}^2 - \Pi(C)$ are in bijection, and we let α denote the (potentially infinite) cardinality of these sets.

In light of all of this, the exact sequence (B.39) becomes

$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathbb{Z}^n \longrightarrow H_1\big(S^2 - C\big) \longrightarrow \mathbb{Z} \longrightarrow \mathbb{Z}^{\alpha+1} \longrightarrow \mathbb{Z}^{\alpha} \longrightarrow 0, \tag{B.40}$$

and insisting that the alternating sum of ranks in this truncated sequence vanish allows us to deduce that the homology group $H_1\big(S^2 - C\big)$ has rank $n - 1$.

From here, we can deduce that $H_1\big(S^2 - C\big) \cong \mathbb{Z}^{n-1}$ via an appeal to Lemma B.5.3. The complement $S^2 - C$ inherits the local connectedness of S^2, so that it is topologically the disjoint union of each of its connected components, and each such component is a non-compact surface with free fundamental group. Combining the standard result that the Abelianization of any free group is a free Abelian group with Lemma B.5.2, we find that $H_1\big(S^2 - C\big)$ is a direct sum of free Abelian groups, and in particular is itself free Abelian. Every free Abelian group of rank $n - 1$ is isomorphic to $\mathbb{Z}^{n-1}$, and so the proof is complete. □

Bibliography

[AB06] C.D. Aliprantis and K.C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, Berlin, 2006.

[ABČ15] A. Anušić, H. Bruin, and J. Činč. The Core Ingram Conjecture for non-recurrent critical points, 2015, 1512.07073.

[Ale15] J.W. Alexander. A proof of the invariance of certain constants of analysis situs. *Transactions of the American Mathematical Society*, 16:148–154, 1915.

[Ale23] J.W. Alexander. On the deformation of an n-cell. *Proceedings of the National Academy of Sciences of the United States of America*, 9(12):406–407, 1923.

[Ale24] J.W. Alexander. An Example of a Simply Connected Surface Bounding a Region which is not Simply Connected. *Proceedings of the National Academy of Sciences of the United States of America*, 10(1):8–10, 1924.

[Ale27] P.S. Alexandroff. Über stetige Abbildungen kompakter Räume. *Mathematische Annalen*, 96:555–571, 1927.

[AU24] P.S. Alexandroff and P. Urysohn. Zur Theorie der topologischen Räume. *Mathematische Annalen*, 92:258–266, 1924.

[Ayr27] W.L. Ayres. Concerning continuous curves and correspondences. *Annals of Mathematics*, 28:396–418, 1927.

[BBD96] M. Barge, K. Brucks, and B. Diamond. Self-Similarity in Inverse Limit Spaces of the Tent Family. *Proceedings of the American Mathematical Society*, 124(11):3563–3570, 1996.

[BBŠ12] M. Barge, H. Bruin, and S. Štimac. The Ingram conjecture. *Geometry and Topology*, 16(4):2481–2516, 2012.

[BdCH17] P. Boyland, A. de Carvalho, and T. Hall. Natural extensions of unimodal maps: virtual sphere homeomorphisms and prime ends of basin boundaries, 2017, 1704.06624.

[Bel79] D.P. Bellamy. A tree-like continuum without the fixed point property. *Houston Journal of Mathematics*, 6:1–13, 1979.

[BG91] C.A. Berenstein and R. Gay. *Complex Variables: An Introduction*. Springer-Verlag, New York, 1991.

[Bin69] R.H. Bing. The elusive fixed point property. *The American Mathematical Monthly*, 76(2):119–132, 1969.

[Can83] G. Cantor. Über unendliche lineare Punktmannigfaltigkeiten. *Mathematische Annalen*, 21(5):545–591, 1883.

[Can95] G. Cantor. Beiträge zur Begründung der transfiniten Mengenlehre. *Mathematische Annalen*, 46:481–512, 1895.

[Can78] J.W. Cannon. The recognition problem: What is a topological manifold? *Bulletin of the American Mathematical Society*, 84(5):832–866, 1978.

[Car13] C. Carathéodory. Über die gegenseitige Beziehung der Ränder bei der konformen Abbildung des Inneren einer Jordanschen Kurve auf einen Kreis. *Mathematische Annalen*, 73:305–320, 1913.

[Cla16] P.L. Clark. Convergence. `http://math.uga.edu/~pete/convergence.pdf`, 2016. Accessed 2nd August 2020.

[Dav86] R.J. Daverman. *Decompositions of Manifolds*. Academic Press, Inc., Orlando, 1986.

[DH82] A. Douady and J.H. Hubbard. Exploring the Mandelbrot set: The Orsay Notes. `http://pi.math.cornell.edu/~hubbard/OrsayEnglish.pdf`, 1982. Accessed 2nd August 2020.

[Die05] R. Diestel. *Graph Theory*. Springer, New York, 2005.

[FMOT08] R. Fokkink, J.C. Mayer, L.G. Oversteegen, and E.D. Tymchatyn. The Plane Fixed Point Problem, 2008, 0805.1184v2.

[Hah14] H. Hahn. Mengentheoretische Charakterisierung der stetigen Kurve. *Sitzungsberichte der Akademie der Wissenschaten in Wien*, 123:2433–2489, 1914.

[Hat02] A. Hatcher. Algebraic Topology. `http://pi.math.cornell.edu/~hatcher/AT/AT.pdf`, 2002. Accessed 23rd August 2020.

[Hau14] F. Hausdorff. *Grundzüge der Mengenlehre*. Verlag von Veit and Co., Leipzig, 1914.

[Hau27] F. Hausdorff. *Mengenlehre*. Verlag Walter de Gruyter and Co., Berlin, 1927.

[Hen64] G.W. Henderson. The pseudo-arc as an inverse limit with one binding map. *Duke Mathematical Journal*, 31(3):421–425, 1964.

[Hur35] W. Hurewicz. Beiträge zur Topologie der Deformationen II. Homotopie- und Homologiegruppen. *Proceedings of the Koninklijke Akademie van Wetenschappen*, 38:521–528, 1935.

[IM10] W.T. Ingram and W.S. Mahavier. *Inverse Limits: From Continua to Chaos*. Springer, New York, 2010.

[Jan13] Z. Janiszewski. O rozcinaniu plaszczyny przez kontinua. *Prace Mathematyczno-Fizyczne*, 26:11–63, 1913.

[Joh31] I. Johansson. Topologische Untersuchungen über unverzweigte Überlagerungsflächen. *Skriften Norske Videnskaps-Akademi Oslo Mathematisk-Naturvidenskabelig Klasse*, 1:1–69, 1931.

[KH95] A. Katok and B. Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, Cambridge, 1995.

[Kur22] K. Kuratowski. Sur l'operation $\overline{A}$ de l'analysis situs. *Fundamenta Mathematicae*, 3:182–199, 1922.

[KW30] K. Kuratowski and G.T. Whyburn. Sur les éléments cycliques et leurs applications. *Fundamenta Mathematicae*, 16:305–331, 1930.

[Lav24] M. Lavrentieff. Contribution à la théorie des ensembles homéomorphes. *Fundamenta Mathematicae*, 6:149–160, 1924.

[Lee00] J. Lee. *Introduction to Topological Manifolds*. Springer-Verlag, New York, 2000.

[Len05] N.J. Lennes. Curves in non-metrical analysis situs. *Bulletin of the American Mathematical Society*, 12:284, 1905.

[Mar00] D. Marker. *Model Theory: An Introduction*. Springer-Verlag, New York, 2000.

[Mas78] W.S. Massey. How to give an exposition of the Čech-Alexander-Spanier type homology theory. *The American Mathematical Monthly*, 85(2):75–83, 1978.

[Maz13a] S. Mazurkiewicz. O arytmetyzacji kontinuów. *Comptes Rendus de la Société des Sciences de Varsovie*, 6:305–311, 1913.

[Maz13b] S. Mazurkiewicz. O arytmetyzacji kontinuów II. *Comptes Rendus de la Société des Sciences de Varsovie*, 6:941–945, 1913.

[McA66] B.L. McAllister. Cyclic elements in topology: A history. *The American Mathematical Monthly*, 73(4):337–350, 1966.

[Men27] K. Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927.

[Mey09] D. Meyer. Expanding Thurston maps as quotients, 2009, 0910.2003.

[Mey14] D. Meyer. Unmating of rational maps: sufficient criteria and examples. *Frontiers in Complex Dynamics*, 51:197–233, 2014.

[Mil04] J. Milnor. Pasting Together Julia Sets: A Worked Out Example of Mating. *Experimental Mathematics*, 13(1):55–92, 2004.

[Mil06] J. Milnor. *Dynamics in One Complex Variable*. Princeton University Press, Princeton, 2006.

[Moo15] R.L. Moore. On a Set of Postulates which Suffice to Define a Number-Plane. *Transactions of the American Mathematical Society*, 16(1):27–32, 1915.

[Moo16] R.L. Moore. On the Foundations of Plane Analysis Situs. *Transactions of the American Mathematical Society*, 17(2):131–164, 1916.

[Moo20] R.L. Moore. Concerning Simple Continuous Curves. *Transactions of the American Mathematical Society*, 21(3):333–347, 1920.

[Moo25] R.L. Moore. Concerning upper semi-continuous collections of continua. *Transactions of the American Mathematical Society*, 27(4):416–428, 1925.

[Moo29] R.L. Moore. Concerning upper semi-continuous collections. *Monatschefte für Mathematik und Physik*, 36:81–88, 1929.

[Moo08] G.H. Moore. The emergence of open sets, closed sets, and limit points in analysis and topology. *Historia Mathematica*, 35:220–241, 2008.

[MS22] E.H. Moore and H.L. Smith. A General Theory of Limits. *American Journal of Mathematics*, 44(2):102–121, 1922.

[MT88] J. Milnor and W. Thurston. On Iterated Maps of the Interval. *Lecture Notes in Mathematics*, 1342:465–563, 1988.

[Mul22] A.M. Mullikin. Certain theorems relating to plane connected point sets. *Transactions of the American Mathematical Society*, 24:144–162, 1922.

[Nad92] S. Nadler. *Continuum Theory: An Introduction*. CRC Press, Boca Raton, 1992.

[Nad93] S. Nadler. Continuum Theory and Graph Theory: Disconnection Numbers. *Journal of the London Mathematical Society*, s2-47(1):167–181, 1993.

[New85] M.H.A. Newman. *Elements of the Topology of Plane Sets of Points*. Praeger, Santa Barbara, 1985.

[Nöb32] G. Nöbeling. Eine Verschärfung des n-Beinsatzes. *Fundamenta Mathematicae*, 18:23–38, 1932.

[Par66] W. Parry. Symbolic dynamics and transformations of the unit interval. *Transactions of the American Mathematical Society*, 122:368–378, 1966.

[PM13] C.L. Petersen and D. Meyer. On The Notions of Mating. *Annales de la Faculté des Sciences de Toulouse*, 21(5):839–876, 2013.

[Pug10] C.C. Pugh. *Real Mathematical Analysis*. Springer, New York, 2010.

[Ree92] M. Rees. A Partial Description of the Parameter Space of Rational Maps of Degree Two: Part 1. *Acta Mathematica*, 168:11–87, 1992.

[Rie06] F. Riesz. Die Genesis des Raumbegriffs. *Mathematische und naturwissenschaftliche Berichte aus Ungarn*, 24:309–353, 1906.

[Roi90] J. Roitman. *Introduction to Modern Set Theory.* John Wiley and Sons, New York, 1990.

[RR47] T. Radó and P. Reichelderfer. On Cyclic Transitivity. *Fundamenta Mathematicae*, 34:14–29, 1947.

[Shi00] M. Shishikura. On a theorem of M. Rees for matings of polynomials. *Mandelbrot Set: Theme and Variations*, pages 289–305, 2000.

[Sie20] W. Sierpiński. Sur une condition pour qu'un continu soit une courbe jordanienne. *Fundamenta Mathematicae*, 1:44–60, 1920.

[Sie28] W. Sierpiński. *Topologia ogólna.* Kasa imienia Józefa Mianowskiego, Warsaw, 1928.

[Sie34] W. Sierpiński. *Introduction to General Topology.* University of Toronto, Toronto, 1934.

[Spa66] E.H. Spanier. *Algebraic Topology.* Springer-Verlag, New York, 1966.

[SS78] L.A. Steen and J.A. Seebach. *Counterexamples in Topology.* Dover Publications, Mineola, 1978.

[SS03] E.M. Stein and R. Shakarchi. *Princeton Lectures in Analysis II: Complex Analysis.* Princeton University Press, Princeton, 2003.

[Sti93] J. Stillwell. *Classical Topology and Combinatorial Group Theory.* Springer-Verlag, New York, 1993.

[Tan92] L. Tan. Matings of quadratic polynomials. *Ergodic Theory and Dynamical Systems*, 12(3):589–620, 1992.

[Tho92] C. Thomassen. The Jordan-Schönflies Theorem and the Classification of Surfaces. *The American Mathematical Monthly*, 99(2):116–131, 1992.

[Tie22] H. Tietze. Über ein Beispiel von L. Vietoris zu den Hausdorffschen Umgebungsaxiomen. *Mathematische Annalen*, 87:150–152, 1922.

[Tie23] H. Tietze. Beiträge zur allgemeinen Topologie I: Axiome für verschiedene Fassungen des Umgebungsbegriffs. *Mathematische Annalen*, 88:290–312, 1923.

[Tor21] M. Torhorst. Über den Rand der einfach zusammenhängenden ebenen Gebiete. *Mathematische Zeitschrift*, 9:44–65, 1921.

[Veb04] O. Veblen. A System of Axioms for Geometry. *Transactions of the American Mathematical Society*, 5(3):343–384, 1904.

[Vic94] J.W. Vick. *Homology Theory: An Introduction to Algebraic Topology.* Springer-Verlag, New York, 1994.

[vK35] E.R. van Kampen. On some characterizations of 2-dimensional manifolds. *Duke Mathematical Journal*, 1(1):74–93, 1935.

[vMR90] J. van Mill and G.M. Reed. *Open Problems in Topology.* North-Holland, Amsterdam, 1990.

[Wal74] R.C. Walker. *The Stone-Čech Compactification.* Springer-Verlag, Berlin, 1974.

[WD79] G.T. Whyburn and E. Duda. *Dynamic Topology.* Springer-Verlag, New York, 1979.

[Wes18] T. West. *Continuum Theory and Dynamical Systems.* CRC Press, Boca Raton, 2018.

[Why27a] G.T. Whyburn. Concerning continua in the plane. *Transactions of the American Mathematical Society*, 29:369–400, 1927.

[Why27b] G.T. Whyburn. Cyclicly connected continuous curves. *Proceedings of the National Academy of Sciences of the United States of America*, 13:31–38, 1927.

[Why27c] G.T. Whyburn. Some properties of continuous curves. *Bulletin of the American Mathematical Society*, 33:305–308, 1927.

[Why28] G.T. Whyburn. Concerning the Cut Points of Continua. *Transactions of the American Mathematical Society*, 30(3):597–609, 1928.

[Why31] G.T. Whyburn. On the Cyclic Connectivity Theorem. *Bulletin of the American Mathematical Society*, 37(6):429–433, 1931.

[Why45] G.T. Whyburn. *Analytic Topology.* American Mathematical Society, Providence, 1945.

[Why48] G.T. Whyburn. On n-Arc Connectedness. *Transactions of the American Mathematical Society*, 63(3):452–456, 1948.

[Wil27] R.L. Wilder. Concerning R.L. Moore's Axioms Σ_1 for Plane Analysis Situs. *Bulletin of the American Mathematical Society*, 34(6):752–760, 1927.

[Wil49] R.L. Wilder. *Topology of Manifolds.* American Mathematical Society Colloquium Publications, Providence, 1949.

[Wil70] S. Willard. *General Topology.* Addison-Wesley, Reading, 1970.

[Wil76] R.L. Wilder. Robert Lee Moore: 1882-1974. *Bulletin of the American Mathematical Society*, 82:417–427, 1976.

[Wil78] R.L. Wilder. Evolution of the Topological Concept of 'Connected'. *The American Mathematical Monthly*, 85(9):720–726, 1978.

[Wit88] B.S. Wittner. *On the Bifurcation Loci of Rational Maps of Degree Two.* PhD thesis, Cornell University, 1988.

[YZ00] M. Yampolsky and S. Zakeri. Mating Siegel quadratic polynomials. *Journal of the American Mathematical Society*, 14(1):25–78, 2000.

[Zip30] L. Zippin. On Continuous Curves and the Jordan Curve Theorem. *American Journal of Mathematics*, 52(2):331–350, 1930.

[Zip33] L. Zippin. Independent arcs of a continuous curve. *Annals of Mathematics*, 34(1):95–113, 1933.