On the homology of infinite graphs with ends

Dissertation
zur Erlangung des Doktorgrades
der Fakultät für Mathematik, Informatik
und Naturwissenschaften
der Universität Hamburg

vorgelegt
im Department Mathematik
von

Philipp Sprüssel
aus Hamburg

Hamburg
2010

Als Dissertation angenommen vom Department
Mathematik der Universität Hamburg

auf Grund der Gutachten von Prof. Dr. Reinhard Diestel
und Prof. Dr. Birgit Richter

Hamburg, den 13.01.2010

Prof. Dr. Reiner Lauterbach
Leiter des Departments Mathematik

Contents

Acknowledgment

Firstly, I would like to thank my advisor Reinhard Diestel for all his support, both in mathematical and organizational respect. I am also very grateful to all members of our research group, especially Agelos Georgakopoulos for uncountably many discussions, for frisbee lessons, and for being the best office mate ever, Henning Bruhn-Fujimoto for letting me share in his know-how in mathematical matters, and Matthias Hamann for just being there the moment I needed to talk about an algebraic problem. I thank Ron Aharoni for Erdős 2 and for hosting me for two months, Eli Berger, Matt DeVos, and Imre Leader for collaboration during several visits, Laurent Bartholdi for helpful discussions on algebraic matters and for drawing our attention to the Čech homology, and the Studienstiftung des deutschen Volkes for supporting me with scholarships throughout the years. My special thanks go to my family and to Robin for being part of it for almost fourteen years.

Chapter 1

Introduction and overview

This thesis is about infinite graphs. The main question that inspired the research presented in this thesis is how the homology of a locally finite graph interacts with its combinatorial structure. Graphs are simplicial 1-complexes and therefore one traditionally, and conveniently, considers their simplicial homology. The coefficients are usually taken from a field such as $\mathbb{F}_2$, $\mathbb{R}$ or $\mathbb{C}$, which makes the first simplicial homology group into a vector space called the *cycle space* of G. We will use this terminology also in the case that the coefficients are taken from $\mathbb{Z}$.

For reasons to become apparent soon let us denote the first homology group of G as $\mathcal{C}_{\text{fin}} = \mathcal{C}_{\text{fin}}(G)$. For the moment it will suffice to take our coefficients from $\mathbb{F}_2$ and interpret the elements of $\mathcal{C}_{\text{fin}}$ as sets of edges. Moreover, let us first consider finite graphs. Apparently, the structure of $\mathcal{C}_{\text{fin}}$ as a group tells us little as such, since it is always a direct sum of $\mathbb{F}_2$'s depending only on the number of vertices and edges of G. But there are a number of classical theorems—often called 'cycle space theorems'—showing that the interaction of $\mathcal{C}_{\text{fin}}$ with the combinatorial structure of G can tell us more about commonly investigated graph properties. Examples are:

- Tutte's theorem that the induced non-separating cycles in a 3-connected graph generate the cycle space,
- MacLane's theorem that a graph is planar if and only if its cycle space has a 'simple' basis (i.e. each edge lies in at most two elements of the basis),
- Whitney's theorem that a graph is planar if and only if it has an abstract dual (where 'dual' is defined in terms of the cycle space),

and many more. The cycle space has thus become one of the standard aspects of finite graphs used in their structural analysis.

For infinite graphs the interaction between $\mathcal{C}_{\text{fin}}$ and the combinatorial structure of G is not nearly that powerful: Most of the cycle space theorems, including the ones cited above, fail for infinite graphs, even for *locally* *finite* ones, graphs in which each vertex has only finitely many neighbours. However, for locally finite graphs this can be remedied by defining the cycle space slightly differently.

Let G be a locally finite graph. In combinatorial terms, the cycle space $\mathcal{C}_{\text{fin}}$ of G is defined by considering the (finite) edge sets of (combinatorial) cycles in G and letting the cycle space consist of all finite sums of those edge sets. Diestel and Kühn [16, 17] introduced the *topological cycle space* $\mathcal{C} = \mathcal{C}(G)$, which is defined as follows. Let $|G|$ be the Freudenthal compactification [28] of G and consider the (possibly infinite) edge sets of circles in $|G|$, homeomorphic images of the unit circle S^1 in $|G|$. These edge sets are called *circuits*. The sums of circuits—where infinite sums are allowed as long as they are *thin*, i.e. every edge of the graph lies in only finitely many summands—then form the topological cycle space.

Since [16, 17] first appeared, the topological cycle space has shown to be a surprisingly successful approach. Indeed, all the standard cycle space theorems have been shown to extend to locally finite graphs for the topological cycle space, see [2, 3, 4, 5, 7, 8, 9, 13, 18, 30, 31, 32, 40], or [14, 15] for an overview.

Given the success of $\mathcal{C}$ for graphs, it seems desirable to recast its definition in homological terms that make no reference to the one-dimensional character of $|G|$ (e.g., to circles), to obtain a homology theory for similar but more general spaces (such as non-compact CW complexes of any dimension) that implements the ideas and advantages of $\mathcal{C}$ more generally. To obtain more general results, we shall always choose the coefficients from $\mathbb{Z}$; the results for $\mathbb{F}_2$ will follow from those for $\mathbb{Z}$ by taking them modulo 2. The *oriented cycle space* $\vec{\mathcal{C}}(G)$ is defined analogously to the topological cycle space: Let $\vec{\mathcal{E}}(G)$ be the *oriented edge space* of G, the group of all integer-valued functions on the set of oriented edges of G that take inverse values on inverse orientations of edges, and let $\vec{\mathcal{C}}$ be the subgroup of $\vec{\mathcal{E}}$ that consists of thin sums of *oriented circuits*, functions that are 1 on the edges of a circle (with cyclic orientation) and 0 elsewhere.

For such an extendable translation of our combinatorial definition of $\vec{\mathcal{C}}$ into algebraic terms, simplicial homology is easily seen not to be the right approach: while $|G|$ is not a simplicial complex, the simplicial homology of G itself (without ends) yields the classical (oriented) cycle space $\vec{\mathcal{C}}_{\text{fin}}$. One way of extending simplicial homology to more general spaces is Čech homology; and indeed we will show that its first group applied to $|G|$ is isomorphic to $\vec{\mathcal{C}}$. But there the usefulness of Čech homology for graphs ends: since its groups are constructed as limits rather than directly from chains and cycles, they do not interact with the combinatorial structure of G in the way we expect and know it from $\vec{\mathcal{C}}$.

The next candidate for the desired description of $\vec{\mathcal{C}}$ in terms of homology is singular homology. Indeed, $\vec{\mathcal{C}}$ is built from circles in $|G|$ and circles are singular 1-cycles, which generate the first singular homology group $H_1(|G|)$ of $|G|$, so both groups are built from similar elements. On the face of it, it is not clear whether $\vec{\mathcal{C}}$ might in fact be isomorphic, even canonically, to $H_1(|G|)$. However, it will turn out that it is not: we shall prove that $\vec{\mathcal{C}}$ is always a natural quotient of $H_1(|G|)$, but this quotient is proper unless G is essentially finite. This may seem surprising, since $\vec{\mathcal{C}}$ is defined via (thin) infinite sums while all sums in the definition of $H_1(|G|)$ are finite, which suggests that $\vec{\mathcal{C}}$ might be larger than $H_1(|G|)$.

Our approach for the comparison of $\vec{\mathcal{C}}$ and $H_1(|G|)$ will be to define a homomorphism from $Z_1(|G|)$ to $\vec{\mathcal{E}}$ that determines how often the (oriented) edges of G are traversed by the simplices of a 1-cycle z and that maps z to this function on the oriented edges. It will turn out that this homomorphism vanishes on boundaries and that its image is precisely $\vec{\mathcal{C}}$. Hence it defines an epimorphism $f\colon H_1(|G|) \to \vec{\mathcal{C}}(G)$. Finally, we will show that f is not normally injective, by constructing loops that traverse every edge equally often in either direction but that are not null-homologous; an example is given in Figure 1.1. Thus, $\vec{\mathcal{C}}$ (and also $\mathcal{C}$) is a genuinely new object, also from a topological point of view.

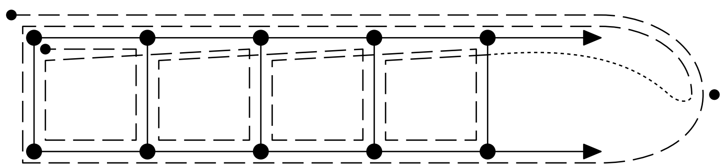

Figure 1.1: A loop that is not null-homologous but whose homology class is mapped to the zero element in the oriented edge space.

For our proof that loops like the 'curling simplex' shown in Figure 1.1 are not null-homologous we shall need a better understanding of the fundamental group of $|G|$. This will enable us to define an invariant on 1-chains in $|G|$ that can distinguish 1-cycles like the curling simplex from boundaries of singular 2-chains, hence completing the proof that f need not be injective. The fundamental group of a finite graph G is easy to describe: it is the free group on the (oriented) *chords* of a spanning tree of G, the edges of G that are not edges of the spanning tree. For the Freudenthal compactification of infinite graphs, the situation is different, since a loop in $|G|$ can traverse infinitely many chords while the elements of a free group are always finite sums of its generators.

One of the main aims of this thesis will thus be to develop a combinatorial description of the fundamental group of the space $|G|$ for an arbitrary connected locally finite graph G. We shall describe $\pi_1(|G|)$, as for finite G, in terms of reduced words in the oriented chords of a spanning tree. However, when G is infinite this does not work for arbitrary spanning trees but only for *topological spanning trees*. Moreover, we will have to allow infinite words of any countable order type, and reduction by cancelling adjacent inverse letters or sequences of letters does not suffice. However, the kind of reduction we need can be described in terms of word reductions in the free groups F_I on all the finite subsets I of chords, which enables us to embed the group F_∞ of infinite reduced words in the inverse limit of those F_I, and handle it in this form. On the other hand, mapping a loop in $|G|$ to the sequence of chords it traverses, and then reducing that sequence (or word), turns out to be well defined on homotopy classes and hence defines an embedding of $\pi_1(|G|)$ as a subgroup in F_∞.

Our combinatorial characterization of $\pi_1(|G|)$ re-proves the description of the fundamental group of the *Hawaiian Earring* by Higman [35] and Cannon and Conner [10]. The Hawaiian Earring $\bigcup_{n\in\mathbb{N}}\{x \in \mathbb{R}^2 \mid \|x-(0,1/n)\| = 1/n\}$ is homotopy equivalent to the Freudenthal compactification of any graph G that has precisely one *non-trivial* end, an end to which there converges a sequence of chords of some topological spanning tree of G. Our characterization of $\pi_1(|G|)$ will hence yield that the fundamental group of the Hawaiian Earring is precisely F_∞, as shown by Higman [35] and later, with a characterization in terms of words similar to ours, by Cannon and Conner [10].

The last aim of this thesis will then be to define a variant of singular homology that captures, for locally finite graphs G and dimension 1, precisely the topological cycle space of G. Our hope with this plan is to stimulate further work in two directions. One is that its new topological guise should make the cycle space accessible to topological methods that might generate some windfall for the study of graphs. And conversely, that as the approach that gave rise to $\mathcal{C}$ is made accessible to more general spaces and higher dimensions, its proven usefulness for graphs might find some more general topological analogues. It is therefore natural to require the spaces for which we shall define our homology theory to have some properties that makes the approach of $\mathcal{C}$ applicable to them: Analogously to G and its Freudenthal compactification $|G|$ that gave rise to $\mathcal{C}$, we shall consider locally compact Hausdorff spaces X with a fixed Hausdorff compactification $\hat{X}$.

The construction of our homology theory will be done in two steps: First, we shall define an ad-hoc homology that satisfies $H_1(G) = \vec{\mathcal{C}}(G)$. This homology will not satisfy the axioms for homology, but it will serve as an introduction of the main ideas of how to capture $\vec{\mathcal{C}}$ by a homology. In the second step, we shall then define a homology theory for locally compact Hausdorff spaces with compactification that will satisfy all the Eilenberg–Steenrod axioms [26]. The proof that this homology theory satisfies $H_1(G) = \vec{\mathcal{C}}(G)$ will depend on the work done in the first step.

A main feature of our homology will be that it treats the compactification points, or *ends*,[1] differently from other points: All simplices will be 'based in' X, i.e. its 0-faces will live in X; our chains, which we allow to be infinite, will have to be locally finite in X but not at ends. On the face of it, the second requirement looks similar to the definition of the 'locally finite homology' given in [36], but this homoloy does not consider ends and hence yields different groups. Indeed, the locally finite homology would not succeed in capturing the topological cycle space: It allows for cycles like the sum $\sum_{n\in\mathbb{Z}} \sigma_n$ in $\mathbb{R}$, where $\sigma_n(x) = n + x$ for $x \in [0, 1]$, which does not correspond to an element of the cycle space (note that $\mathbb{R}$ is homeomorphic to the double-ray and thus 'cycle space' is defined in this case).

This thesis is organized as follows. After going through the main definitions and notation Chapter 2 we shall develop our combinatorial characterization of $\pi_1(|G|)$ in Chapter 3. The Čech homology will be discussed briefly in Chapter 5, and in Chapter 4 we define the homomorphism $f\colon H_1(|G|) \to \vec{\mathcal{C}}(G)$ and show that it is surjective, but not normally injective. In Chapter 6 we will then define our new homology theory for locally compact spaces and show that it satisfies the axioms for homology and coincides with the topological cycle space for graphs and dimension 1.

Parts of this thesis have been submitted for publication: the combinatorial characterization of the fundamental group of $|G|$ [19], the comparison of the topological cycle space and singular homology, Čech homology, and the ad-hoc homology from Section 6.2 [20], and the definition and analysis of the new homology theory defined in Section 6.3 [21]. Moreover, there are extended versions [23] and [22] of [19] and [20], respectively, which include proofs of easy facts that have been omitted in the original papers.

[1] With a slight abuse of notation, we will call the point in $\hat{X} \setminus X$ ends, although we do not require that they are ends in the sense of Freudenthal [28].

Chapter 2

Definitions and basic facts

For graphs we use the terminology of [14], for topology that of Hatcher [34]. We shall need a couple of facts and definitions from graph theory and algebraic topology, which we will introduce in the following sections, separated by fields. But first, let us look at some basic topological concepts.

2.1 Topology

All graphs in this thesis may have multiple edges but no loops. This said, we shall from now on use the terms *path* and *loop* topologically, for continuous but not necessarily injective maps $\sigma\colon [0,1] \to X$, where X is any Hausdorff space. If σ is a loop, it is *based at* the point $\sigma(0) = \sigma(1)$. We write σ^- for the path $s \mapsto \sigma(1-s)$. If a path is injective, we call its image an *arc* in X. If a loop σ is 'internally injective' (i.e. $\sigma(x) = \sigma(y)$ implies $x = y$ or $\{x, y\} = \{0, 1\}$), then its image is homeomorphic to the unit circle S^1 in $\mathbb{C}$. In this case we call σ a *circle path* and its image a *circle* in X.

The following fact can be found in [33, p. 208].

Lemma 2.1. *The image of a topological path with distinct endpoints x, y in a Hausdorff space X contains an arc in X between x and y.*

Given a set $\{X_k \mid i \in I\}$ of topological spaces, we write $X = \bigsqcup X_k$ for their disjoint union endowed with the disjoint union topology.

The homology we shall define in Chapter 6 will crucially rely on the concept of topological dimension. There are many different ways to define topological dimension, most of them have proved to be very useful when applied to certain types of spaces; the definition we use in this thesis is usually called *Lebesgue covering dimension*, other ones are, for instance, the large and the small *inductive dimension*. For an introduction and (much) more about topological dimension, see [27].

Let X be a normal space[1] and let $k \geq -1$ be an integer. Then X has dimension at most k if every open cover $\mathcal{U}$ of X has a refinement $\mathcal{U}'$ such that every $x \in X$ lies in at most $k+1$ sets from $\mathcal{U}'$. If n is the smallest number with this property, then we say that X has *dimension* n. An example for an n-dimensional space is the Euclidian space $\mathbb{R}^n$—in fact, every dimension theory is supposed to assign dimension n to $\mathbb{R}^n$. More generally, it is not hard to see that every subset of $\mathbb{R}^n$ containing an open ball has dimension n; in particular, the standard n-simplex Δ^n in $\mathbb{R}^{n+1}$, being homeomorphic to such a subset of $\mathbb{R}^n$, has dimension n.

The following is an immediate consequence of the definition.

Lemma 2.2. *Let X be a normal space. The following claims hold:*

(i) *X has dimension -1 if and only if X is empty.*

(ii) *If X has dimension 0, then it is totally disconnected.*

For the proof that the homology we shall define in Chapter 6 satisfies the axioms for homology—which will be stated shortly—we will need to consider Cartesian products of spaces. It would be natural to assume that the dimension of a product does not exceed the sum of dimensions of its factors, but this is not generally true. However, it is true if both spaces are compact [27, Theorem 3.2.13]:

Lemma 2.3. *Let X, Y be compact Hausdorff spaces. Then the product space $X \times Y$ has dimension at most $\dim(X) + \dim(Y)$.*

2.2 Graph theory

Let G be a locally finite connected graph, fixed throughout this section. For graphs, ends and the Freudenthal compactification are usually defined in a combinatorial way, as follows. A 1-way infinite (graph-theoretical) path in G is a *ray*. Two rays are *equivalent* if no finite set of vertices separates them in G, and the resulting equivalence classes are the *ends* of G; write $\Omega = \Omega(G)$ for the set of ends of G.

The Freudenthal compactification $|G|$ of G can now be defined as follows: its point set is the union of G and $\Omega(G)$, its basic open sets are the basic open sets of G itself (as a 1-complex) and the sets $\hat{C}(S,\omega)$ defined for every end ω and every finite set S of vertices, as follows. Let $C(S,\omega) =: C$ be the unique component of $G - S$ in which ω *lives* (i.e., in which every ray of ω has a *tail*, or subray), and let $\hat{C}(S,\omega)$ be the union of C with the

[1]Lebesgue covering dimension has proved to have very useful properties especially for normal spaces, but it can also be defined for general spaces. However, we will only need the definition for normal spaces.

set of all the ends of G that live in C and the (finitely many) open edges between S and C.[2] It is not hard to see that $|G|$ is indeed the Freudenthal compactification of G.

Note that the boundary of $\hat{C}(S, \omega)$ in $|G|$ is a subset of S, that every ray converges to the end containing it, and that the set of ends is totally disconnected. Since every basic open neighbourhood of a vertex or a point on an edge contains at most one vertex, we have at once

Lemma 2.4. *Let U be a basic open set in $|G|$. Then every sequence of vertices in U converges to a point in U.*

The ends of G provide a new way to connect points in $|G|$: If F is a *cut* (i.e. the set of edges between the sides of a bipartition of $V(G)$), then there is no arc in G from one side of the partition to the other that avoids all edges in F. In $|G|$ there can be such an arc, but only when F contains infinitely many edges [14, Lemma 8.5.5]:

Lemma 2.5. *Let F be a finite cut and let X, Y be the sides of the corresponding partition of $V(G)$. Then every arc from X to Y meets an inner point of some edge in F.*

We shall frequently use the following non-trivial lemma.

Lemma 2.6 ([18])**.** *For a locally finite graph G, every closed, connected subspace of $|G|$ is arc-connected.*

A *standard subspace* of $|G|$ is a closed connected subspace of $|G|$ that contains every edge of which it contains an inner point. Note that by Lemma 2.6 every standard subspace is arc-connected. A *topological spanning tree* of G is a standard subspace of $|G|$ that contains every vertex—and hence also every end—of G and that contains no circle. It is easy to see that every locally finite connnected graph has a topological spanning tree; for instance, the closure of a normal spanning tree is a topological spanning tree. An edge of G that does not lie in T is a *chord* of T.

Normal spanning trees can be used to show that $|G|$ is a metric space: If T is a normal spanning tree, let the edges in T have length $\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots$, according to their height in T, and let every chord of T have length the sum of lengths of the tree-edges it spans. These edge-lengths are easily seen to induce a metric of $|G|$. We thus have

Theorem 2.7. *$|G|$ is a compact metric space.*

By Theorem 2.7, $|G|$ is normal. We can thus consider the topological dimension of $|G|$.

[2]The definition given in [14] is slightly different, but equivalent to the simpler definition given here when G is locally finite. Generalizations are studied in [37, 41].

Lemma 2.8. $|G|$ *is one-dimensional.*

Proof. Since $|G|$ is compact, it suffices to consider finite open covers. So let $\mathcal{U}$ be a finite open cover of $|G|$. For every end ω of G consider a neighbourhood $\hat{C}(S_\omega, \omega)$ that is contained in some set in $\mathcal{U}$. Since $\Omega(G)$, being a closed subspace of the compact Hausdorff space $|G|$, is compact, finitely many neighbourhoods $\hat{C}(S_\omega, \omega)$ suffice to cover it. Let S be the union of the sets S_ω beloging to those finitely many neighbourhoods. Then the open sets $\hat{C}(S, \omega)$ cover $\Omega(G)$, they are pairwise disjoint, and each of them is contained in some set from $\mathcal{U}$. The part of $|G|$ that is not yet covered is a finite graph and hence we can easily extend our choice of open sets $\hat{C}(S, \omega)$ to a cover of all of $|G|$ in which every point of $|G|$ is contained in at most two sets. □

An edge $e = uv$ of G has two *directions*, (u, v) and (v, u). A triple (e, u, v) consisting of an edge together with one of its two directions is an *oriented edge*. The two oriented edges corresponding to e are its two *orientations*, denoted by $\vec{e}$ and $\overleftarrow{e}$. Thus, $\{\vec{e}, \overleftarrow{e}\} = \{(e, u, v), (e, v, u)\}$, but we cannot generally say which is which. However, from the definition of G as a 1-complex we have a fixed homeomorphism $\theta_e \colon [0, 1] \to e$. We call $(\theta_e(0), \theta_e(1))$ the *natural direction* of e, and $(e, \theta_e(0), \theta_e(1))$ its *natural orientation*.

Let $\sigma \colon [0, 1] \to |G|$ be a path in $|G|$. Given an edge $e = uv$ of G, if $[s, t]$ is a subinterval of $[0, 1]$ such that

$$\{\sigma(s), \sigma(t)\} = \{u, v\} \qquad \text{and} \qquad \sigma\big((s, t)\big) = \mathring{e} := \theta_e\big((0, 1)\big),$$

we say that σ *traverses* e on $[s, t]$. It does so *in the direction of* $\big(\sigma(s), \sigma(t)\big)$, or *traverses* $\vec{e} = (e, \sigma(s), \sigma(t))$. We then call its restriction to $[s, t]$ a *pass of σ through* e, or $\vec{e}$, *from $\sigma(s)$ to $\sigma(t)$*.

Using that $[0, 1]$ is compact and $|G|$ is Hausdorff, one easily shows that a path in $|G|$ contains at most finitely many passes through any given edge:

Lemma 2.9. *A path in $|G|$ traverses each edge only finitely often.*

Proof. Let σ be a path in $|G|$, and let $e = uv$ be an edge such that σ contains infinitely many passes $\sigma \restriction [s_n, t_n]$ through e, $n = 1, 2, \dots$. Passing to a subsequence if necessary, we may assume that the sequence $s_1, s_2, \dots$ converges, say to $s \in [0, 1]$. Then the sequence of the corresponding t_n also converges to s: given $\varepsilon > 0$, choose m large enough that for all $n > m$ both $|s_n - s| < \varepsilon/2$ and $t_n - s_n < \varepsilon/2$ (using that the lengths of the intervals $[s_n, t_n]$ converge to 0, which they clearly do); then $|t_n - s| < \varepsilon$ for all $n > m$. But now σ fails to be continuous at s, because $\{\sigma(s_n), \sigma(t_n)\} = \{u, v\}$ for each n but each point in $|G|$ has a neighbourhood avoiding u or v. □

Let $\vec{\mathcal{E}} = \vec{\mathcal{E}}(G)$ denote the set of all integer-valued functions φ on the set $\vec{E}$ of all oriented edges of G that satisfy $\varphi(\overleftarrow{e}) = -\varphi(\vec{e})$ for all $\vec{e} \in \vec{E}$. This is an abelian group under pointwise addition. A family $(\varphi_i \mid i \in I)$ of elements of $\vec{\mathcal{E}}$ is *thin* if for every $\vec{e} \in \vec{E}$ we have $\varphi_i(\vec{e}) \neq 0$ for only finitely many i. Then $\varphi = \sum_{i \in I} \varphi_i$ is a well-defined element of $\vec{\mathcal{E}}$: it maps each $\vec{e} \in \vec{E}$ to the (finite) sum of those $\varphi_i(\vec{e})$ that are non-zero. We shall call a function $\varphi \in \vec{\mathcal{E}}$ obtained in this way the *thin sum* of those φ_i.

We can now define our oriented version of the topological cycle space of G. When α is a circle path in $|G|$ based at a vertex, we call the function $\varphi_\alpha \colon \vec{E} \to \mathbb{Z}$ defined by

$$\varphi_\alpha \colon \vec{e} \mapsto \begin{cases} 1 & \text{if } \alpha \text{ traverses } \vec{e} \\ -1 & \text{if } \alpha \text{ traverses } \overleftarrow{e} \\ 0 & \text{otherwise.} \end{cases}$$

an *oriented circuit* in G, and write $\vec{\mathcal{C}} = \vec{\mathcal{C}}(G)$ for the subgroup of $\vec{\mathcal{E}}$ formed by all thin sums of oriented circuits.

We remark that $\vec{\mathcal{C}}$ is closed also under infinite thin sums [16, Cor. 5.2], but this is neither obvious nor generally true for thin spans of subsets of $\vec{\mathcal{E}}$ [6, Sec. 3]. We remark further that composing the functions in $\vec{\mathcal{C}}$ with the canonical homomorphism $\mathbb{Z} \to \mathbb{Z}_2$ yields the usual *topological cycle space* $\mathcal{C}(G)$ of G as studied in [3, 4, 5, 8, 9, 16, 17, 18, 30, 32, 40], the $\mathbb{F}_2$ vector space of subsets of E obtained as thin sums of (unoriented) circuits. The *combinatorial cycle space* consisting of all finite sums of finite circuits will be denoted by $\mathcal{C}_{\text{fin}}$, or by $\vec{\mathcal{C}}_{\text{fin}}$ in the oriented case.

The topological cycle space $\mathcal{C}(G)$ can be characterized as the set of those subsets of E that meet every *finite* cut of G in an even number of edges [16, Thm. 7.1], [14, Thm. 8.5.8]. The characterization has an oriented analogue:

Theorem 2.10. *An element φ of $\vec{\mathcal{E}}$ lies in $\vec{\mathcal{C}}$ if and only if $\sum_{\vec{e} \in \vec{F}} \varphi(\vec{e}) = 0$ for every finite oriented cut $\vec{F}$ of G.*

The proof of Theorem 2.10 is not completely trivial. But it adapts readily from the unoriented proof given e.g. in [14], which we leave to the reader to check if desired.

Given a topological spanning tree T of G and a chord e of T, we write C_e for the *fundamental circuit* of e, the unique circuit in $T+e$. The *oriented fundamental circuit* $\vec{C}_e$ is the unique oriented circuit in $T+e$ with positive value on the natural orientation of e.

Theorem 2.11. *The fundamental circuits of a topological spanning tree generate the topological cycle space. The oriented fundamental circuits generate the oriented cycle space.*

The first statement of Theorem 2.11 is given in [18, Theorem 6.1]. Again, the proof of the oriented statement adapts readily from the proof of the unoriented version.

We close this section with a standard tool in infinite graph theory, König's infinity lemma (eg. [14, Lemma 8.1.2]).

Lemma 2.12. *Let $V_1, V_2, \ldots$ be non-empty finite sets and let G be a graph on $\bigcup_n V_n$ such that every vertex in each V_{n+1} has at least one neigbour in V_n. Then G contains a ray whose nth vertex lies in V_n for every n.*

2.3 Algebra and algebraic topology

For our study of the fundamental group of $|G|$ we shall need two basic theorems: The Nielsen–Schreier theorem and the Seifert–van Kampen theorem.

Theorem 2.13 ([39]). *Every subgroup of a free group is free.*

We shall only need a weaker version of the Seifert–van Kampen theorem, the theorem in its whole strength it can be found eg. in [34, Theorem 1.20].

Theorem 2.14. *Let G be a locally finite graph and assume that a standard subspace H of $|G|$ is the union of standard subspaces H_1, H_2 with $H_1 \cap H_2 = \{x\}$, where x is a vertex. Then $\pi_1(H) \simeq \pi_1(H_1) * \pi_1(H_2)$.*

Note that Theorem 2.14 indeed follows from the Seifert–van Kampen theorem: Let U be an open star of radius $\varepsilon < 1$ around x; then $H_1' := H_1 \cup U$ and $H_2' := H_2 \cup U$ are open sets with intersection U, and clearly $\pi_1(H_i')$ is canonically isomorphic to $\pi_1(H_i)$. As U is simply connected, the Seifert–van Kampen theorem yields $\pi_1(H) \simeq \pi_1(H_1') * \pi_1(H_2') \simeq \pi_1(H_1) * \pi_1(H_2)$.

A 1-cycle that can be written as a sum of 1-simplices no two of which share their first point is an *elementary cycle*. Every 1-cycle is easily seen to be a sum of elementary 1-cycles, a decomposition which is not normally unique. When we prove statements about $H_1(|G|)$, it will often suffice to consider elementary 1-cycles.

The following lemma enables us to subdivide or concatenate the simplices in a 1-cycle while keeping it in its homology class.

Lemma 2.15. *Let σ be a singular 1-simplex in $|G|$, and let $s \in (0,1)$. Write σ' and σ'' for the 1-simplices obtained from the restrictions of σ to $[0,s]$ and to $[s,1]$ by reparametrizing linearly. Then $\sigma' + \sigma'' - \sigma$ is the boundary of a 2-simplex with image $\operatorname{Im}\sigma$.* □

When σ is a summand in a cycle $z \in Z_1$, we shall say that the equivalent cycle z' obtained by replacing σ with $\sigma' + \sigma''$ in the sum arises by *subdividing* σ (at s or at $\sigma(s)$). A frequent application of Lemma 2.15 is the following:

Corollary 2.16. *Every non-zero element of* $H_1(|G|)$ *is represented by a sum of loops each based at a vertex.*

Proof. Pick a cycle representing a given homology class, and decompose it into elementary cycles. Use Lemma 2.15 to concatenate their simplices into a single loop. If such a loop α passes through a vertex, we can subdivide it there and suppress its original boundary point, obtaining a homologous loop based at that vertex. If α does not pass through a vertex, then $\operatorname{Im}\alpha \subseteq \mathring{e}$ for some edge e (since non-trivial sets of ends are never connected), so α is null-homotopic and $[\alpha] = 0$. □

A *homology theory* assigns to every space X and every subspace A of X a sequence $\big(H_n(X,A)\big)_{n\in\mathbb{Z}}$ of abelian groups,[3] and to every continuous map $f\colon X \to Y$ with $f(A) \subset B$ for subspaces A of X and B of Y (which we indicate by writing $f\colon (X,A) \to (Y,B)$) a sequence of homomorphisms $f_*\colon H_n(X,A) \to H_n(Y,B)$ so that $(fg)_* = f_*g_*$ for compositions of maps and $\mathbb{1}_* = \mathbb{1}$ for the identity maps. We abbreviate $H_n(X,\emptyset)$ to $H_n(X)$. Finally, the following *axioms for homology* have to be satisfied:

Homotopy equivalence: If continuous maps $f, g\colon (X,A) \to (Y,B)$ are homotopic, then $f_* = g_*$.

The Long Exact Sequence of a Pair: For every pair (X,A) there are *boundary homomorphisms* $\partial\colon H_n(X,A) \to H_{n-1}(A)$ such that

$$\cdots \xrightarrow{\partial} H_n(A) \xrightarrow{\iota_*} H_n(X) \xrightarrow{\pi_*} H_n(X,A) \xrightarrow{\partial} H_{n-1}(A) \xrightarrow{\iota_*} H_{n-1}(X) \xrightarrow{\pi_*} \cdots$$

is an exact sequence, where ι denotes the inclusion $(A,\emptyset) \to (X,\emptyset)$ and π denotes the inclusion $(X,\emptyset) \to (X,A)$. These boundary homomorphisms are *natural*, i.e. given a continuous map $f\colon (X,A) \to (Y,B)$ the diagrams

$$\begin{array}{ccc} H_n(X,A) & \xrightarrow{\partial} & H_{n-1}(A) \\ \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} \\ H_n(Y,B) & \xrightarrow{\partial} & H_{n-1}(B) \end{array}$$

commute.

[3]Usually (in particular, for all homology theories considered in this thesis) $H_n(X,A)$ is the trivial group for $n < 0$, but this is not part of the requirements on homology theories.

Excision: Given subspaces A, B of X whose interiors cover X, the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms $H_n(B, A \cap B) \to H_n(X, A)$ for all n.

Disjoint unions: If $X = \coprod_\alpha X_\alpha$ with inclusions $\iota_\alpha \colon X_\alpha \hookrightarrow X$, the direct sum map $\bigoplus_\alpha (\iota_\alpha)_* \colon \bigoplus_\alpha H_n(X_\alpha, A_\alpha) \to H_n(X, A)$, where $A = \coprod_\alpha A_\alpha$, is an isomorphism.

The axioms above are the ones stated in [34, p. 160–161]. The original Eilenberg-Steenrod axioms [26] contain an additional axiom, called the 'dimension axiom', stating that the homology groups of a single point are nonzero only in dimension zero. However, this axiom is not always regarded as an essential part of the requirements for a homology theory. An example for a homology theory that does not satisfy the dimension axiom is bordism theory; in this case the groups of a single point are non-trivial in infinitely many dimensions. We omit the dimension axiom, but note that the homology theory we construct will trivially satisfy it.

The groups $H_n(X, A)$ above are called *relative homology groups*; specializations $H_n(X) = H_n(X, \emptyset)$ are *absolute homology groups*.

Let us note that not every homology theory satisfies all axioms without modification; it is often the case that a homology theory is defined for certain types of (pairs of) spaces or certain groups of coefficients. The Čech homology for instance, which we will discuss in Chapter 5, does satisfy all the axioms under the additional assumption that both X and A are compact and the coefficients are taken from a module over a ring or from a compact topological group. Otherwise, the exactness axiom is not generally satisfied.

A *cohomology theory* has to satisfy axioms dual to those for a homology theory. Thus, a cohomology theory assigns to every space X, every subspace A of X, and every abelian group G a sequence $\big(H^n(X, A; G)\big)_{n\in\mathbb{Z}}$ of abelian groups, and to every continuous map $f\colon (X, A) \to (Y, B)$ a sequence of homomorphisms $f^*\colon H^n(Y, B; G) \to H^n(X, A; G)$ so that $(fg)^* = g^* f^*$ and $\mathbb{1}^* = \mathbb{1}$. The axioms are the following:

Homotopy equivalence: If continuous maps $f, g\colon (X, A) \to (Y, B)$ are homotopic, then $f^* = g^*\colon H^n(Y, B; G) \to H^n(X, A; G)$.

The Long Exact Sequence of a Pair: For every pair (X, A) there are *coboundary homomorphisms* $\delta\colon H^n(A; G) \to H^{n+1}(X, A; G)$ such that

$$\cdots \xrightarrow{\delta} H^n(X, A; G) \xrightarrow{\pi^*} H^n(X; G) \xrightarrow{\iota^*} H^n(A; G) \xrightarrow{\delta} H^{n+1}(X, A; G) \xrightarrow{\pi^*} H^{n+1}(X; G) \xrightarrow{\iota^*} \cdots$$

is an exact sequence, where ι denotes the inclusion $(A, \emptyset) \to (X, \emptyset)$ and π denotes the inclusion $(X, \emptyset) \to (X, A)$. These coboundary homomorphisms are *natural*, i.e. given a standard map $f : (X, A) \to (Y, B)$ the diagrams

$$\begin{array}{ccc} H^n(B; G) & \xrightarrow{\delta} & H^{n+1}(Y, B; G) \\ \downarrow{\scriptstyle f^*} & & \downarrow{\scriptstyle f^*} \\ H^n(A; G) & \xrightarrow{\delta} & H^{n+1}(X, A; G) \end{array}$$

commute.

Excision: Given subspaces A, B of X whose interiors cover X, the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms $H^n(X, A; G) \to H^n(B, A \cap B; G)$ for all n.

Disjoint unions: If $X = \bigsqcup_\alpha X_\alpha$ with inclusions $\iota_\alpha \colon X_\alpha \hookrightarrow X$, the direct product map $\prod_\alpha (\iota_\alpha)_* : H^n(X, A; G) \to \prod_\alpha H^n(X_\alpha, A_\alpha; G)$, where $A = \bigsqcup_\alpha A_\alpha$, is an isomorphism.

We shall need another basic algebraic lemma, the Five-Lemma (eg. [34, p. 129]).

Lemma 2.17. *Let*

$$\begin{array}{ccccccccc} A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & D & \longrightarrow & E \\ \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \gamma} & & \downarrow{\scriptstyle \delta} & & \downarrow{\scriptstyle \varepsilon} \\ A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & D' & \longrightarrow & E' \end{array}$$

be a commutative diagram in which both horizontal sequences are exact. If α, β, δ, and ε are isomorphisms, then so is γ.

2.4 The groups of a finite graph

In this section let G be a finite graph and let T be a fixed spanning tree of G. We sketch the well known fact that the fundamental group and the first homology group of G can be expressed in terms of the chords of T.

Let us first consider the fundamental group of G. Every loop based at a vertex can be characterized by its passed through the chords of T, hence every such loop induces a word with letters the (oriented) chords of T. Now if two loops σ and τ induce words w_σ and w_τ, and if w_σ and w_τ reduce to the same word—where a reduction is a sequence of deleting inverse pairs of adjacent letters (such as $\vec{e}\,\overleftarrow{e}$ or $\overleftarrow{e}\,\vec{e}$)—, then σ and τ are homotopic.

This homotopy can be realized by recursively retracting subpaths that pass through a chord and back (and possibly travel along T in between the passes, but not along other chords) in the order given by the reductions of the words. Therefore, $\pi_1(G)$ is canonically isomorphic to the group of reduced words, i.e. the free group with generators the (oriented) chords of T.

The first homology group $H_1(G)$ can be obtained similarly. Using the Seifert–van Kampen theorem and the fact that the fundamental group of the unit circle S^1 is canonically isomorphic to $\mathbb{Z}$ (or, alternatively, by the equivalence of simplicial and singular homology), on shows that $H_1(G)$ is canonically isomorphic to the free abelian group with generators the (oriented) chords of T. Hence $H_1(G) \simeq \vec{\mathcal{C}}(G)$. It is now clear that $H_1(G)$ is the abelianization of $\pi_1(G)$.

Let us remark that the first cohomology group H^1 of G yields the same result as H_1, provided that the coefficients in H_1 and H^1 are taken from the same group.

All facts metioned in this section also remain true for locally finite graphs which have a topological spanning tree with only finitely many chords.

Chapter 3

The fundamental group of $|G|$

3.1 Introduction

In this chapter we give a combinatorial characterization of the fundamental group of the Freudenthal compactification $|G|$ of a locally finite graph G.

When G is finite, we saw in Section 2.4 that $\pi_1(|G|) = \pi_1(G)$ is the free group on the set of (arbitrarily oriented) chords of a spanning tree of G. We shall see that when G is infinite and there are infinitely many chords, then $\pi_1(|G|)$ is not a free group. However, we show that it embeds canonically as a subgroup in an inverse limit F^* of free groups: those on the finite sets of (oriented) chords of any *topological* spanning tree T of G.

More precisely, we characterize $\pi_1(|G|)$ in terms of subgroup embeddings

$$\pi_1(|G|) \to F_\infty \to F^*,$$

where F_∞ is a group formed by 'reduced' infinite words of chords of T.[1] These words arise as the traces of loops in $|G|$, so in general they will have arbitrary countable order types. Unlike for finite graphs, many natural homotopies between such loops do not proceed by retracting passes through chords one by one. (We give a simple example in Section 3.3.) Nevertheless, we show that to generate the homotopy classes of loops in $|G|$ from suitable representatives we only need homotopies that do retract passes through chords one at a time, in some linear order. As a consequence, we are again able to define reduction of words as a linear sequence of steps each cancelling one pair of letters, although the order in which the steps are performed may now have any countable order type (such as that of the rationals).

The fact that our sequences of reduction steps are not well-ordered will make it difficult or impossible to handle reductions in terms of their

[1] Covering space theory does not apply since, trivial exceptions aside, $|G|$ is not semi-locally simply connected at ends.

definition. However we show that reduction of infinite words can be characterized in terms of the reductions they induce on all their finite subwords. A formalization of this observation yields the embedding $F_\infty \to F^*$.

An end of G is *trivial* if it has a contractible neighbourhood. If every end of G is trivial, then $|G|$ is homotopy equivalent to a finite graph. If G has exactly one non-trivial end, then $|G|$ is homotopy equivalent to the Hawaiian Earring. Its fundamental group was studied by Higman [35] and by Cannon and Conner [10]. Our characterization of $\pi_1(|G|)$ coincides with their combinatorial description of this group when G has only one non-trivial end.

The characterization of $\pi_1(|G|)$ is not only an important result in its own right—the fundamental group of such a standard space as $|G|$ ought to be understood—, but it also has a substantial application: in the proof that $\vec{\mathcal{C}}(G)$ is usually a proper quotient of $H_1(|G|)$ in Chapter 4, Theorem 3.14 below is the cornerstone of the proof.

This chapter is organized as follows. We begin with a section collecting together the remaining definitions and known background that we need. In Section 3.3 we introduce our group F_∞ of infinite words, and show how it embeds in the inverse limit of the free groups on its finite subsets of letters. In Section 3.4 we embed $\pi_1(|G|)$ in F_∞, leaving the proof of the main lemma to Section 3.5.

3.2 Terminology and basic facts

All homotopies between paths that we consider are relative to the first and last point of their domain, usually $\{0,1\}$. In many cases we shall construct homotopies between paths using that certain subpaths are homotopic. For instance, if the restrictions of two given paths to $[0,\frac{1}{2}]$ are homotopic, as well as their restrictions to $[\frac{1}{2},1]$, then performing both 'subpath homotopies' at the same time yields a homotopy between the original paths. The same is clearly true for any finite number of subpath homotopies—but not for infinitely many, as one cannot guarantee continuity at accumulation points of the subintervals when combining the subpath homotopies. However, if the subpath homotopies behave 'nicely', then the following lemma enables us to combine them.

Lemma 3.1. *Let α,β be paths in a topological space X. Assume that there are disjoint subintervals $(a_0,b_0),(a_1,b_1),\dots$ of $[0,1]$ such that α and β conincide on $[0,1]\setminus\bigcup_n(a_n,b_n)$, while each segment $\alpha\restriction[a_n,b_n]$ is homotopic in $\alpha\big([a_n,b_n]\big)\cup\beta\big([a_n,b_n]\big)$ to $\beta\restriction[a_n,b_n]$. Then α and β are homotopic.*

Proof. Write $D:=\bigcup_n(a_n,b_n)$. For every $n\in\mathbb{N}$ let $F^n=\big(f^n_t\big)_{t\in[0,1]}$ be a homotopy in $\alpha\big([a_n,b_n]\big)\cup\beta\big([a_n,b_n]\big)$ between $\alpha\restriction[a_n,b_n]$ and $\beta\restriction[a_n,b_n]$. We define the desired homotopy $F = \big(f_t\big)_{t\in[0,1]}$ between α and β as

$$f_t(x) := \begin{cases} f_t^n(x) & \text{if } x \in (a_n, b_n), \\ \alpha(x) = \beta(x) & \text{if } x \in [0,1] \setminus D. \end{cases}$$

Clearly, $f_0 = \alpha$ and $f_1 = \beta$. It remains to prove that F is continuous.

Let $x, t \in [0,1]$ and a neighbourhood U of $F(x,t)$ in X be given. We find an $\varepsilon > 0$ so that $F\big((x-\varepsilon, x], (t-\varepsilon, t+\varepsilon)\big) \subset U$; the case $F\big([x, x+\varepsilon), (t-\varepsilon, t+\varepsilon)\big) \subset U$ is analogous. If $x \in (a_n, b_n]$ for some n, then as F^n is continuous, there is an ε with $F\big((x-\varepsilon, x], (t-\varepsilon, t+\varepsilon)\big) \subset U$.

We may thus assume that $x \notin \bigcup_n (a_n, b_n]$. This implies that no interval $(x-\varepsilon, x)$ is contained in D, as otherwise we would have $(x-\varepsilon, x) \subseteq (a_n, b_n)$ for some n and hence $x \in (a_n, b_n]$. Now by assumption $x \in [0,1] \setminus D$, and hence $F(x,t) = \alpha(x) = \beta(x)$. Pick $\varepsilon > 0$ with $x - \varepsilon \in [0,1] \setminus D$ small enough that both α and β map $[x-\varepsilon, x]$ into U. We claim that $F\big((x-\varepsilon, x], (t-\varepsilon, x+\varepsilon)\big) \subset U$. Indeed, for every $x' \in (x-\varepsilon, x] \setminus D$ and every $t' \in (t-\varepsilon, t+\varepsilon)$ we have $F(x', t') = \alpha(x') = \beta(x') \in U$. On the other hand, for every $x' \in (x-\varepsilon, x] \cap D$ and $t' \in (t-\varepsilon, t+\varepsilon)$ we have $x' \in (a_n, b_n)$ for some n. As x and $x - \varepsilon$ lie in $[0,1] \setminus D$, we have $(a_n, b_n) \subset (x-\varepsilon, x)$ and hence

$$F(x', t') = F^n(x', t') \in \alpha\big([a_n, b_n]\big) \cup \beta\big([a_n, b_n]\big) \subseteq \alpha\big([x-\varepsilon, x]\big) \cup \beta\big([x-\varepsilon, x]\big) \subseteq U$$

by the choice of ε. □

Many topological spaces that are not normally associated with graphs can be expressed as a graph with ends, or as a subspace thereof. The Hawaiian Earring, for example, is homeomorphic to the subspace of the infinite grid that consists of all the vertical double rays and its end. Since the subspaces of graphs with ends form a richer class than the spaces of graphs with ends themselves, we prove all our results not just for $|G|$ but more generally for subspaces H of $|G|$. However, the reader will lose little by thinking of H as the entire space $|G|$. All subspaces we shall consider will be standard subspaces of G.

A *topological tree in* $|G|$ is a standard subspace of $|G|$ that contains no circle. Note that the subgraph that such a space induces in G need not be connected: its connectedness may hinge on the ends it contains. By Lemma 2.6 every topological tree is arc-connected. Chords of topological trees are defined the same way as for ordinary trees or topological spanning trees.

Lemma 3.2. *Topological trees in $|G|$ are locally arc-connected.*

Proof. Let T be a topological tree in $|G|$. Let D be any open subset of T, and $x \in D$. We have to find an arc-connected open neighbourhood of x in T inside D. This is trivial if x is a vertex or an inner point of an edge, so we assume that x is an end. Then D may be chosen of the form $D = \hat{C}(S, x) \cap T$, for some finite set $S \subseteq V(G)$. Since $G - S$ has only finitely many components, $T \setminus S$ is a finite union of open sets of this form, so D is open and closed in $T \setminus S$.

Similarly, $T \setminus S$ has only finitely many arc-components, and hence only finitely many components. Each of them is closed and open in $T \setminus S$, and open even in T. One component, C_x say, contains x. Then $C_x \subseteq D$, since D is open and closed in $T \setminus S$. To complete the proof, we show that C_x is arc-connected.

Suppose not. As C_x is the union of some of the finitely many arc-components of $T \setminus S$, it has only finitely many arc-components. Not all of them can be closed in C_x, since C_x is connected. Let C be an arc-component of C_x that is not closed in C_x. Then its closure $\overline{C}$ in T meets $C_x \setminus C$, and clearly it does so only in ends of G.

Since the components of $T \setminus S$ other than C_x are open in T, we have $\overline{C} \subseteq C_x \cup S$. As C is connected and T is closed in $|G|$, we know that $\overline{C}$ is connected and closed in $|G|$, and hence arc-connected by Lemma 2.6. Let A be an arc in $\overline{C}$ from a point in C to one in $C_x \setminus C$. As S is finite and $\overline{C} \setminus (S \cup C)$ only consists of ends, we can choose A so that it avoids S: If A does not avoid S, then after its last vertex in S it traverses an edge. As the inner points of this edge lie in $\overline{C}$, they also lie in C. Hence the final segment of A starting at the centre of this edge is an arc as desired. But then $A \subseteq C_x$, contradicting the definition of C as an arc-component of C_x. □

Between any two of its points, x and y say, a topological tree T in $|G|$ contains a unique arc, which we denote by xTy. These arcs are 'short' also in terms of the topology on T induced by $|G|$:

Lemma 3.3. *If a sequence $(z_i)_{i \in \mathbb{N}}$ of points in T converges to a point z, then every neighbourhood of z contains all but finitely many of the arcs $z_i T z_{i+1}$.*

Proof. Since the arcs $z_i T z_{i+1}$ are unique, Lemma 3.2 implies that they lie in arbitrarily small neighbourhoods of z. □

We shall need topological trees in $|G|$ as spanning trees for our analysis of $\pi_1(|G|)$: arbitrary graph-theoretical spanning trees of G can have non-trivial loops in their closures, which would leave no trace of chords and thus be invisible to our intended representation of homotopy classes by words of such chords.

Clearly, a topological tree T in $|G|$ is a topological spanning tree of G if and only if it contains $V(G)$ (and hence also $\Omega(G)$). Similarly, given a standard subspace H of $|G|$, we call a topological tree T in $|G|$ a *topological spanning tree of* H if $T \subseteq H$ and T contains every vertex or end of G that lies in H.

As mentioned in Section 2.2, topological spanning trees exist in all locally finite connected graphs. We now prove that they also exist in all the relevant subspaces. In fact, we need a slight technical strengthening of this:

Lemma 3.4. *Let* $T \subseteq H$ *be standard subspaces of* $|G|$*. If* T *is a topological tree, it can be extended to a topological spanning tree of* H*.*

Proof. By Theorem 2.7, $|G|$ is a compact Hausdorff space. Let $\mathcal{S}$ be the set of standard subspaces of $|G|$ such that $T \subseteq S \subseteq H$ and S contains all the vertices and ends of G that lie in H. Since H is closed in $|G|$, every $S \in \mathcal{S}$ is closed not only in H but also in $|G|$, and therefore compact. Since the intersection of a nested chain of compact connected Hausdorff spaces is connected [42, p. 203], $\mathcal{S}$ has a minimal element T' by Zorn's Lemma. By Lemma 2.6, T' is arc-connected, and it contains no circle: if it did, we could delete an edge to obtain a smaller element of $\mathcal{S}$. (Since $V(G) \cup \Omega(G)$ is totally disconnected, every circle in $|G|$ contains an edge.) Hence T' is a topological tree in $|G|$, and by definition of $\mathcal{S}$ a topological spanning tree of H containing T. □

Like graph-theoretical trees, topological trees in $|G|$ are contractible. Again, we shall need a slightly technical strengthening of this. Call a homotopy $F(x,t)$ *time-injective* if for every x the map $t \mapsto F(x,t)$ is either constant or injective.

Lemma 3.5. *For every point* x *in a topological tree* T *in* $|G|$ *there is a time-injective deformation retraction of* T *onto* x*.*

Proof. Similar to the metric of $|G|$ defined in Section 2.2 the space T is metrizable as follows. Choose an enumeration of the edges in T and give the nth edge length 2^{-n}. Define the distance $d(y,z)$ between points y,z in T as the sum of lengths of the edges (and partial edges) in yTz; note that if $y \neq z$ then yTz meets the interior of at least one edge. Then clearly d is a metric with $d(y,z) \leq 1$ for all $y,z \in T$, and it is easy to check that it induces the given topology on T: The neighbourhoods of vertices and inner points of edges are trivially identical in both topologies. For an end ω, every basic open neighbourhood $\hat{C}(S,\omega) \cap T$ contains the open ε-ball (with respect to d) around ω, where ε is the minimum of lengths of the edges incident with S. On the other hand, for a given $\varepsilon > 0$, choose a finite vertex set S large enough so that the sum of lengths of all edges incident with $V(G) \setminus S$ is less than ε. Let $U \subseteq \hat{C}(S,\omega) \cap T$ be an arc-connected neighbourhood of ω in T, which exists by Lemma 3.2. Then by the choice of S all points in U have mutual distance less than ε, thus U is contained in the ε-ball around ω.

By construction, if $z \in yTy'$ for some $y, y' \in T$, then we have $d(y, y') = d(y, z) + d(z, y')$. We now define the time-injective homotopy F in T from the identity on T to the map $T \to \{x\}$ as follows: For every $y \in T$ and $t \in [0, 1]$ let $F(y, t)$ be the unique point on xTy at distance $(1 - t) \cdot d(x, y)$ from x.

For the proof that F is continuous, we show that $d\big(F(y, t), F(y', t)\big) \leq d(y, y')$ for every $y, y' \in T$ and $t \in [0, 1]$; then for every $\varepsilon > 0$ and every $y, y' \in T$ with $d(y, y') < \varepsilon/2$ and $t, t' \in [0, 1]$ with $|t - t'| < \varepsilon/2$ we have

$$\begin{aligned} d\big(F(y,t), F(y',t')\big) &\leq d\big(F(y,t), F(y',t)\big) + d\big(F(y',t), F(y',t')\big) \\ &\leq d(y,y') + |t - t'| \cdot d(x, y') \\ &< \varepsilon/2 + (\varepsilon/2) \cdot 1 = \varepsilon. \end{aligned}$$

As xTy and xTy' are closed, there is a last point z on xTy that also lies in xTy'; this point satisfies $xTz = xTy \cap xTy'$ as the unique x–z arc xTz is contained in both xTy and xTy'. Then $yTz \cup zTy'$ is a y–z arc in T and hence $yTy' = yTz \cup zTy'$. This implies $d(y, y') = d(y, z) + d(z, y')$. If $F(y, t) \in zTy$ and $F(y', t) \in zTy'$, then

$$\begin{aligned} d\big(F(y,t), F(y',t)\big) &\leq d\big(F(y,t), z\big) + d\big(z, F(y',t)\big) \\ &\leq d(y,z) + d(z,y') = d(y,y'). \end{aligned}$$

Otherwise at least one of $F(y, t), F(y', t)$ lies in $xTz = xTy \cap xTy'$ and hence both $F(y, t)$ and $F(y', t)$ are contained in xTy or in xTy'. In particular, one of $F(y, t), F(y', t)$ lies on the arc between the other and x. Then

$$\begin{aligned} d\big(F(y,t), F(y',t)\big) &= \big|d\big(x, F(y,t)\big) - d\big(x, F(y',t)\big)\big| \\ &= (1-t) \cdot |d(x,y) - d(x,y')| \leq d(y,y'). \end{aligned}$$

□

Corollary 3.6. *If a topological spanning tree T of a standard subspace H of $|G|$ has only finitely many chords, then H is homotopy equivalent to a finite graph H'. Moreover, H' has a spanning tree T' such that the homotopy equivalence of H and H' maps T to T' and vice versa and maps chords of T to chords of T' and vice versa.*

Proof. Let $T_0 \subseteq T$ be the union of all arcs vTw where v and w are end-vertices of chords of T. Then T_0 is arc-connected and hence a topological tree. The closure $\overline{C}$ of each component C of $T \setminus T_0$ meets T_0 in a single point: It meets T_0 since each point in C sends an arc in T to T_0, and since $\overline{C}$ is arc-connected by Lemma 2.6, it cannot meet T_0 in two points x, y as the union of the x–y arcs in T_0 and $\overline{C}$ would form a circle in T.

By Lemma 3.5 we can retract each such component C to $\overline{C} \cap T_0$, showing that H is homotopy equivalent to its subspace H_0 defined as the union of T_0 and all the chords of T. Now H_0 need not be a finite graph: The arcs that formed T_0 may well travel through ends of G. But as T_0 is the union of finitely many arcs in T, it is homeomorphic to a finite tree T'. Then H_0 is homeomorphic—in particular homotopy equivalent—to the union H' of T' and all the chords of T_0, where each chord of T' connects the images of the endvertices of the corresponding chord of T_0. Thus, H' and T' are as desired. □

Given a standard subspace H of $|G|$, let us call an end ω of G *trivial in H* if $\omega \in H$ and ω has a contractible neighbourhood in H. For instance, all the ends of G are trivial in any topological spanning tree of G, by Lemma 3.5. Trivial ends in larger subspaces can also be made visible by topological spanning trees:

Lemma 3.7. *Let T be a topological spanning tree of a standard subspace H of $|G|$. An end $\omega \in H$ of G is trivial in H if and only if ω has a neighbourhood in H that contains no chord of T.*

Proof. Suppose first that ω has a neighbourhood in H containing no chord of T. This neighbourhood U can be chosen of the form $\hat{C}(S, \omega) \cap H$, since these form a neighbourhood basis of ω, and so that the S–C edges in H are no chords of T either. Then U, indeed its closure $\overline{U}$ in H, contains no inner point of any chord of T, i.e., $\overline{U} \subseteq T$. By Lemma 3.2, there is an arc-connected neighbourhood $U' \subseteq U$ of ω in H, and we may clearly choose U' to be a standard subspace. Its closure T' in H lies in $\overline{U} \subseteq T$, is closed in $|G|$, and therefore is arc-connected by Lemma 2.6. So T' is a topological tree in $|G|$, and contractible by Lemma 3.5.

Conversely, suppose that ω has a contractible neighbourhood U in H; this cannot contain a circle. By Lemma 3.2, the end ω has an arc-connected open neighbourhood T' in T inside U. Since T carries the subspace topology from H, this has the form $T' = U' \cap T$ for an open subset $U' \subseteq U$ of H. This U' is a neighbourhood of ω in H that contains no chord of T: for any such chord it would also contain an arc in $T' \subseteq U$ between its vertices, to form a circle in U that does not exist. □

3.3 Infinite words and inverse limits

In the this section and the next, we will give a combinatorial description of $\pi_1(|G|)$—indeed of $\pi_1(H)$ for any standard subspace H of $|G|$, when G is any connected locally finite graph. Our description will involve infinite words and their reductions in a 'continuous' setting, and embedding the group they form as a subgroup of a limit of finitely generated free groups.

Such things have been studied also by Eda [24], Cannon and Conner [10], and by Chiswell and Müller [11].

In Section 2.4 we saw that when G is finite, $\pi_1(|G|)$ is the free group F on the set of chords of any fixed spanning tree. The standard description of F is given in terms of reduced words of those oriented chords, where reduction is performed by cancelling adjacent inverse pairs of letters such as $\vec{e}_i\overleftarrow{e}_i$ or $\overleftarrow{e}_i\vec{e}_i$. The map assigning to a path in $|G|$ the sequence of chords it traverses defines the canonical group isomorphism between $\pi_1(|G|)$ and F; in particular, reducing the words obtained from homotopic paths yields the same reduced word.

Our description of $\pi_1(|G|)$ when G is infinite will be similar in spirit, but more complex. We shall start not with an arbitrary spanning tree but with a topological spanning tree of $|G|$. Then every path in $|G|$ defines as its 'trace' an infinite word in the oriented chords of that tree, as before. However, these words can have any countable order type, and it is no longer clear how to define the reduction of words in a way that captures homotopy of paths.

Consider the following example. Let G be the infinite ladder, with a topological spanning tree T consisting of one side of the ladder, all its rungs, and its unique end ω (Figure 3.1). The path running along the bottom side of the ladder and back is a null-homotopic loop. Since it traces the chords $\vec{e}_0, \vec{e}_1, \dots$ all the way to ω and then returns the same way, the infinite word $\vec{e}_0\vec{e}_1 \dots \overleftarrow{e}_1\overleftarrow{e}_0$ should reduce to the empty word. But it contains no cancelling pair of letters, such as $\vec{e}_i\overleftarrow{e}_i$ or $\overleftarrow{e}_i\vec{e}_i$.

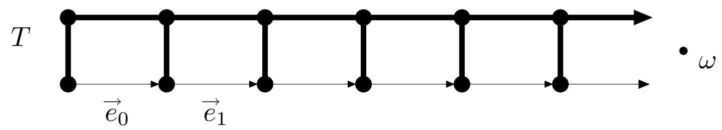

Figure 3.1: The infinite ladder and its topological spanning tree T (bold edges)

This simple example suggests that some transfinite equivalent of cancelling pairs of letters, such as cancelling inverse pairs of infinite sequences of letters, might lead to a suitable notion of reduction. However, in graphs with infinitely many ends one can have null-homotopic loops whose trace of chords contains no cancelling pair of subsequences whatsoever:

Example 3.8. We construct a locally finite graph G and a null-homotopic loop σ in $|G|$ whose trace of chords contains no cancelling pair of subsequences, of any order type.

Let T be the binary tree with root r. Write V_n for the set of vertices at distance n in T from r, and let T_n be the subtree of T induced by

$V_0 \cup \cdots \cup V_n$. Our first aim will be to construct a loop σ in $|T|$ that traverses every edge of T once in each direction. We shall obtain σ as a limit of similar loops σ_n in $T_n \subseteq |T|$.

We start with the constant map $\sigma_0\colon [0,1] \to T_0 = \{r\}$. Assume inductively that $\sigma_n\colon [0,1] \to T_n$ is a loop traversing every edge of T_n exactly once in each direction. Assume further that σ_n pauses every time it visits a vertex in V_n (i.e., a leaf of T_n), remaining stationary at that vertex for some time. More precisely, we assume for every vertex $v \in V_n$ that $\sigma_n^{-1}(v)$ is a non-trivial closed interval. Let us call the restriction of σ_n to such an interval a *pass* of σ_n through v.

Let σ_{n+1} be obtained from σ_n by replacing, for each vertex v in V_n, the pass of σ_n through v by a topological path that first travels from v to its first neighbour in V_{n+1} and back, and then to its other neighbour in V_{n+1} and back, pausing at each of those neighbours for some non-trivial time interval. Outside the passes of σ_n through leaves of T_n, let σ_{n+1} agree with σ_n.

Let us now define σ. Let $s \in [0,1]$ be given. If its values $\sigma_n(s)$ coincide for all large enough n, let $\sigma(s) := \sigma_n(s)$ for these n. If not, then $s_n := \sigma_n(s) \in V_n$ for every n, and $s_0 s_1 s_2 \dots$ is a ray in T; let σ map s to the end of G to which this ray belongs. This map σ is easily seen to be continuous, and by Lemma 3.5 it is null-homotopic. It is also easy to check that no sequence of passes of σ through the edges of T is followed immediately by the inverse of this sequence.

The edges of T are not chords of a topological spanning tree, but this can be achieved by changing the graph: just double every edge.[2] The new edges together with all vertices and ends then form a topological spanning tree in the resulting graph G, whose chords are the original edges of our tree T, and σ is still a (null-homotopic) loop in $|G|$.

Example 3.8 shows that there is no hope of capturing homotopies of loops in terms of word reduction defined recursively by cancelling pairs of inverse subwords, finite or infinite. We shall therefore define the reduction of infinite words differently, though only slightly. We shall still cancel inverse letters in pairs, even only one at a time, and these reduction 'steps' will be ordered linearly (rather unlike the simultaneous dissolution of all the chords by the homotopy in the example). However, the reduction steps will not be well-ordered.

This definition of reduction is less straightforward, but it has an important property: as for finite G, it will be purely combinatorial in terms of letters, their inverses, and their linear order, making no reference to the interpretation of those letters as chords and their relative positions under

[2]And subdivide the new edges once, in case you prefer to obtain a simple graph instead of a graph with multiple edges.

the topology of $|G|$.

Another problem, however, is more serious: since the reduction steps are not well-ordered, it will be difficult to handle reductions—e.g. to prove that every word reduces to a unique reduced word, or that word reduction captures the homotopy of loops, i.e. that traces of homotopic loops can always be reduced to the same word. The key to solving these problems will lie in the observation that the property of being reduced can be characterized in terms of all the finite subwords of a given word. We shall formalize this observation by way of an embedding of our group F_∞ of infinite words in the inverse limit F^* of the free groups on the finite subsets of letters.

The remainder of this section is devoted to carrying out this programme. In Section 3.4 we shall then study how $\pi_1(|G|)$ embeds as a subgroup in F_∞ when its letters are interpreted as oriented chords of a topological spanning tree of G. We shall prove that, as in the finite case, the map assigning to a loop in $|G|$ its trace of chords and reducing that trace is well defined on homotopy classes, giving us injective homomorphisms

$$\pi_1(|G|) \to F_\infty \to F^*.$$

By determining their precise images we shall complete our combinatorial characterization of $\pi_1(|G|)$—and likewise of $\pi_1(H)$ for subspaces H of $|G|$.

Let $\vec{A} = \{\vec{e}_0, \vec{e}_1, \dots\}$ and $\{\overleftarrow{e}_0, \overleftarrow{e}_1, \dots\}$ be disjoint countable sets. Let us call the elements of

$$A := \{\vec{e}_0, \vec{e}_1, \dots\} \cup \{\overleftarrow{e}_0, \overleftarrow{e}_1, \dots\}$$

letters, and say that $\vec{e}_i$ and $\overleftarrow{e}_i$ are *inverse* to each other. A *word* in A is a map $w\colon S \to A$ from a totally ordered countable set S, the set of *positions* of (the letters used by) w, such that $w^{-1}(a)$ is finite for every $a \in A$. The only property of S relevant to us is its order type, so two words $w\colon S \to A$ and $w'\colon S' \to A$ will be considered the same if there is an order-preserving bijection $\varphi\colon S \to S'$ such that $w = w' \circ \varphi$. If S is finite, then w is a *finite* word; otherwise it is *infinite*. The *concatenation* $w_1 w_2$ of two words is defined in the obvious way: we assume that their sets S_1, S_2 of positions are disjoint, put S_1 before S_2 in $S_1 \cup S_2$, and let $w_1 w_2$ be the combined map $w_1 \cup w_2$. For $I \subseteq \mathbb{N}$ we let

$$A_I := \{\vec{e}_i \mid i \in I\} \cup \{\overleftarrow{e}_i \mid i \in I\},$$

and write $w \restriction I$ as shorthand for the restriction $w \restriction w^{-1}(A_I)$. Note that if I is finite then so is the word $w \restriction I$, since $w^{-1}(a)$ is finite for every a.

An *interval* of S is a subset $S' \subseteq S$ closed under betweenness, i.e., such that whenever $s' < s < s''$ with $s', s'' \in S'$ then also $s \in S'$. The most frequently used intervals are those of the form $[s', s'']_S := \{s \in S \mid$

$s' \leq s \leq s''\}$ and $(s', s'')_S := \{s \in S \mid s' < s < s''\}$. If $(s', s'')_S = \emptyset$, we call s', s'' *adjacent* in S.

A *reduction* of a finite or infinite word $w\colon S \to A$ is a totally ordered set R of disjoint 2-element subsets of S such that the two elements of each $p \in R$ are adjacent in $S \setminus \bigcup\{q \in R \mid q < p\}$ and are mapped by w to inverse letters $\vec{e}_i, \overleftarrow{e}_i$. We say that w *reduces to* the word $w \restriction (S \setminus \bigcup R)$. If w has no nonempty reduction, we call it *reduced*.

Informally, one may think of the ordering on R as expressing time. A reduction of a finite word thus recursively deletes cancelling pairs of (positions of) inverse letters; this agrees with the usual definition of reduction in free groups. When w is infinite, cancellation no longer happens 'recursively in time', because R need not be well ordered.

As is well known, every finite word w reduces to a unique reduced word, which we denote as $r(w)$. Note that $r(w)$ is unique only as an abstract word, not as a restriction of w: if $w = \vec{e}_0 \overleftarrow{e}_0 \vec{e}_0$ then $r(w) = \vec{e}_0$, but this letter $\vec{e}_0$ may have either the first or the third position in w. The set of reduced finite words forms a group, with multiplication defined as $(w_1, w_2) \mapsto r(w_1 w_2)$, and identity the empty word. This is the free group with free generators $\vec{e}_0, \vec{e}_1, \ldots$ and inverses $\overleftarrow{e}_0, \overleftarrow{e}_1, \ldots$. For finite $I \subseteq \mathbb{N}$, the subgroup

$$F_I := \{w \mid \operatorname{Im} w \subseteq A_I\}$$

is the free group on $\{\vec{e}_i \mid i \in I\}$.

Consider a word w, finite or infinite, and $I \subseteq \mathbb{N}$. The definitions of reduction and restriction immediately imply the following:

If R is a reduction of w, then $\big\{\{s, s'\} \in R \mid w(s) \in A_I\big\}$, with the ordering induced from R, is a reduction of $w \restriction I$. (3.1)

In particular:

Any result of first reducing and then restricting a word can also be obtained by first restricting and then reducing it. (3.2)

By (3.2), the homomorphisms $F_J \to F_I$, $I \subseteq J$, defined by mapping $w \in F_J$ to $r(w \restriction I) \in F_I$ now make the family of all F_I with finite I an inverse system. Let us write

$$F^* := F^*(\vec{A}) := \varprojlim F_I$$

for the corresponding inverse limit of the F_I. By our assumption that I runs through all the finite subsets of some countable set, and that F_I can be viewed as the free group on I, this defines F^* uniquely as an abstract group.

Our next aim is to show that also every infinite word reduces to a unique reduced word. We shall then be able to extend the map $w \mapsto r(w)$, defined so far only for finite words w, to infinite words w. The operation $(w_1, w_2) \mapsto r(w_1 w_2)$ will then make the set of reduced (finite or infinite) words a group, our desired group F_∞.

Existence is immediate:

Lemma 3.9. *Every word reduces to some reduced word.*

Proof. Let $w\colon S \to A$ be any word. By Zorn's Lemma there is a maximal reduction R of w. Since R is maximal, the word $w \restriction (S \setminus \bigcup R)$ is reduced. $\square$

To prove uniqueness, we begin with a characterization of the reduced words in terms of reductions of their finite subwords. Let $w\colon S \to A$ be any word. If w is finite, call a position $s \in S$ *permanent* in w if it is not deleted in any reduction, i.e., if $s \in S \setminus \bigcup R$ for every reduction R of w. If w is infinite, call a position $s \in S$ *permanent* in w if there exists a finite $I \subseteq \mathbb{N}$ such that $w(s) \in A_I$ and s is permanent in $w \restriction I$. By (3.2), a permanent position of $w \restriction I$ is also permanent in $w \restriction J$ for all finite $J \supseteq I$. The converse, however, need not hold: it may happen that $\{s, s'\}$ is a pair ('of cancelling positions') in a reduction of $w \restriction I$ but $w \restriction J$ has a letter from $A_J \setminus A_I$ whose position lies between s and s', so that s and s' are permanent in $w \restriction J$.

Lemma 3.10. *A word is reduced if and only if all its positions are permanent.*

Proof. The assertion is clear for finite words, so let $w\colon S \to A$ be an infinite word. Suppose first that all positions of w are permanent. Let R be any reduction of w; we will show that $R = \emptyset$. Let s be any position of w. As s is permanent, there is a finite $I \subseteq \mathbb{N}$ such that $w(s) \in A_I$ and s is not deleted in any reduction of $w \restriction I$. By (3.1), the pairs in R whose elements map to A_I form a reduction of $w \restriction I$, so s does not lie in such a pair. As s was arbitrary, this proves that $R = \emptyset$.

Now suppose that w has a non-permanent position s. We shall construct a non-trivial reduction of w. For all $n \in \mathbb{N}$ put $S_n := \{s \in S \mid w(s) \in A_{\{0,\dots,n\}}\}$; recall that these are finite sets. Write w_n for the finite word $w \restriction I$ with $I = \{0, \dots, n\}$, the restriction of w to S_n. For any reduction R of w_{n+1}, the set $R^- := \big\{\{t, t'\} \in R \mid t, t' \in S_n\big\}$ with the induced ordering is a reduction of w_n, by (3.1).

Pick $N \in \mathbb{N}$ large enough so that $s \in S_N$. Since s is not permanent in w, every w_n with $n \geq N$ has a reduction in which s is deleted. As w_n has only finitely many reductions, Lemma 2.12 gives us an infinite sequence $R_N, R_{N+1}, \dots$ in which each R_n is a reduction of w_n deleting s, and $R_n = R^-_{n+1}$ for every $n \geq N$. Inductively, this implies:

$$\text{\textit{For all } } m \leq n, \text{\textit{ we have }} R_m = \big\{\{t,t'\} \in R_n \mid t,t' \in S_m\big\}, \text{\textit{ and the ordering of }} R_m \text{\textit{ on this set agrees with that induced by }} R_n. \tag{3.3}$$

Let $s' \in S$ be such that $\{s,s'\} \in R_n$ for some n; then $\{s,s'\} \in R_n$ for every $n \geq N$, by (3.3).

Our sequence (R_n) divides the positions of w into two types. Call a position t of w *essential* if there exists an $n \geq N$ such that $t \in S_n$ and t remains undeleted in R_n; otherwise call t *inessential*. Consider the set

$$R := \bigcup_{m \geq N} \bigcap_{n \geq m} R_n$$

of all pairs of positions of w that are eventually in R_n. Let R be endowed with the ordering $p < q$ induced by all the orderings of R_n with n large enough that $p, q \in R_n$; these orderings are compatible by (3.3). Note that R is non-empty, since it contains $\{s,s'\}$. We shall prove that R is a reduction of w.

We have to show that the elements of each $p \in R$, say $p = \{t_1, t_2\}$ with $t_1 < t_2$, are adjacent in $S \setminus \bigcup\{q \in R \mid q < p\}$. Suppose not, and pick $t \in (t_1, t_2)_S \setminus \bigcup\{q \in R \mid q < p\}$. If t is essential, then t is a position of w_n remaining undeleted in R_n for all large enough n. But then $\{t_1, t_2\} \notin R_n$ for all these n, contradicting the fact that $\{t_1, t_2\} \in R$. Hence t is inessential. Then t is deleted in every R_n with n large enough. By (3.3), the pair $\{t,t'\} \in R_n$ deleting t is the same for all these n, so $\{t,t'\} =: p' \in R$. By the choice of t, this implies $p' \not< p$. For n large enough that $p, p' \in R_n$, this contradicts the fact that t_1, t_2 are adjacent in $S_n \setminus \bigcup\{q \in R_n, q < p\}$, which they are since R_n is a reduction of w_n. □

Note that a word can consist entirely of non-permanent positions and still reduce to a non-empty word: the word $\vec{e}_0 \overleftarrow{e}_0 \vec{e}_0$ is again an example.

Lemma 3.10 offers an easy way to check whether an infinite word is reduced. In general, it can be hard to prove that a given word w has no nontrivial reduction, since this need not have a 'first' cancellation, see Example 3.8.[3] By Lemma 3.10 it suffices to check whether every position becomes permanent in some large enough but finite $w \restriction I$.

Similarly, it can be hard to prove that two words reduce to the same word. The following lemma provides an easier way to do this, in terms of only the finite restrictions of the two words:

[3]On the other hand, a reduction R of the trace of the path in Example 3.8 to the empty word is not hard to find: Clearly, the elements of R have to be all pairs $\{\vec{e}, \overleftarrow{e}\}$ with e an edge of T. The ordering on R can be any ordering in which a pair $\{\vec{e}, \overleftarrow{e}\}$ appears before a pair $\{\vec{e}', \overleftarrow{e}'\}$ whenever e' lies on the path from e to the root of T.

However, one can construct paths with even more complicated traces, see Theorem 3.14 (i).

Lemma 3.11. *Two words w, w' can be reduced to the same (abstract) word if and only if $r(w \restriction I) = r(w' \restriction I)$ for every finite $I \subseteq \mathbb{N}$.*

Proof. The forward implication follows easily from (3.2). Conversely, suppose that $r(w \restriction I) = r(w' \restriction I)$ for every finite $I \subseteq \mathbb{N}$. By Lemma 3.9, w and w' can be reduced to reduced words v and v', respectively. Our aim is to show that $v = v'$, that is to say, to find an order-preserving bijection $\varphi \colon S \to S'$ between the domains S of v and S' of v' such that $v = v' \circ \varphi$. For every finite I, our assumption and the forward implication of the lemma yield

$$r(v \restriction I) = r(w \restriction I) = r(w' \restriction I) = r(v' \restriction I).$$

Hence for every possible domain $S_I \subseteq S$ of $r(v \restriction I)$ and every possible domain $S'_I \subseteq S'$ of $r(v' \restriction I)$ there exists an order isomorphism $S_I \to S'_I$ that commutes with v and v'. For every I, there are only finitely many such maps $S_I \to S'_I$, since there are only finitely many such sets S_I and S'_I. And for $I \subseteq J$, every such map $S_J \to S'_J$ induces such a map $S_I \to S'_I$ with $S_I \subseteq S_J$ and $S'_I \subseteq S'_J$, by (3.2). Hence by Lemma 2.12 there exists a sequence $\varphi_0 \subseteq \varphi_1 \subseteq \cdots$ of such maps $\varphi_n \colon S_{\{0,\dots,n\}} \to S'_{\{0,\dots,n\}}$, whose union φ maps all of S onto S', since by Lemma 3.10 every position of v and of v' is permanent. □

With Lemma 3.11 we are now able to prove:

Lemma 3.12. *Every word reduces to a unique reduced word.*

Proof. By Lemma 3.9, every word w reduces to some reduced word w'. Suppose there is another reduced word w'' to which w can be reduced. By the easy direction of Lemma 3.11, we have

$$r(w' \restriction I) = r(w \restriction I) = r(w'' \restriction I)$$

for every finite $I \subseteq \mathbb{N}$. By the other direction of Lemma 3.11, this implies that w' and w'' can be reduced to the same word. Since w' reduces only to w' and w'' reduces only to w'', this must be the word $w' = w''$. □

As in the case of finite words, we denote the unique reduced word that a word w reduces to by $r(w)$. The set of reduced words now forms a group

$$F_\infty = F_\infty(\vec{A}),$$

with multiplication defined as $(w_1, w_2) \mapsto r(w_1 w_2)$, identity the empty word, and inverses w^- of $w \colon S \to A$ defined as the map on the same S, but with the inverse ordering, satisfying $\{w(s), w^-(s)\} = \{\vec{e}_i, \overleftarrow{e}_i\}$ for some i for every $s \in S$. (Thus, w^- is w taken backwards, replacing each letter with its inverse.) Note that the proof of associativity requires an application of Lemma 3.12.

As noted earlier, infinite words have been studied by Cannon and Conner [10] and—in a more general setting allowing the letters to live in any group (not necessarily the same for each letter)—by Eda [24]. For reference, let us note that F_∞ equals $BF(\aleph_0)$ in the notation of Cannon & Conner and $\times_{\mathbb{N}}\mathbb{Z} = \times^{\sigma}_{\mathbb{N}}\mathbb{Z}$ in Eda's notation.

As indicated earlier, we claim that F_∞ embeds canonically in the inverse limit F^* of the groups F_I. Indeed, by (3.2), the maps $h_I\colon w \mapsto r(w \restriction I)$ are homomorphisms $F_\infty \to F_I$ that commute with the homomorphisms $F_J \to F_I$ from the inverse system, so they define a homomorphism

$$h\colon F_\infty \to F^*$$

satisfying $\pi_I \circ h = h_I$ for all I (where π_I is the projection $F^* \to F_I$ given by the definition of the inverse limit). To show that h is injective, consider an element w of its kernel. For every I, we have

$$r(w \restriction I) = h_I(w) = \pi_I\big(h(w)\big) = \pi_I(\mathbb{1}) = \emptyset,$$

where $\mathbb{1}$ denotes the identity in F^* and $\emptyset$ that of F_I, the empty word. Thus, w is a reduced word which has no permanent positions. By Lemma 3.10, this means that w is the empty word in F_∞. Thus, h is a group embedding of F_∞ in F^*, as claimed.

We remark that h is not surjective. Indeed, while every letter occurs only finitely often in any given word, there are elements of F^* whose projections to the F_I contain some fixed letter—or even every letter—unboundedly often; such an element will not lie in the image of h. (For example, if w_i is the word $\vec{e}_0\,\vec{e}_1 \cdots \vec{e}_i\overleftarrow{e}_{i-1} \cdots \overleftarrow{e}_1\overleftarrow{e}_0 \in F_{\{1,\dots,i\}}$, then the words $w_0 w_1 \cdots w_i$ for each i define an element of F^* that contains each letter infinitely often.) However, these are clearly the only elements of F^* that h misses: the subgroup $h(F_\infty)$ of F^* consists of precisely those elements (w_I) of F^* that are *bounded* in the sense that for every letter $\vec{e} \in A$ there exists a $k \in \mathbb{N}$ such that $|w_I^{-1}(\vec{e}\,)| \leq k$ for all I.

Theorem 3.14 (ii) below summarizes what we have shown so far.

3.4 Embedding $\mathcal{C}(G)$ in F_∞

Let G be a locally finite connected graph. Let H be a standard subspace of $|G|$, and let T be a fixed topological spanning tree of H. If T has only finitely many chords, then H is homotopy equivalent to a finite graph by Corollary 3.6, and all we shall prove below will be known. We therefore assume that T has infinitely many chords. Enumerate these as $e_0, e_1, \dots$, let $\vec{A} := \{\vec{e}_0, \vec{e}_1, \dots\}$ be the set of their natural orientations, and put

$$A := \{\vec{e}_0, \vec{e}_1, \dots\} \cup \{\overleftarrow{e}_0, \overleftarrow{e}_1, \dots\}.$$

Let

$$F_\infty = F_\infty(\vec{A})$$

be the group of infinite reduced words with letters in A, as defined in Section 3.3.

Unless otherwise mentioned, the endpoints of all paths considered from now on will be vertices or ends. When we speak of 'the passes' of a given path σ, without referring to any particular edges, we shall mean the passes of σ through chords of T.

Every path σ in H defines a word w_σ by its passes through the chords of T. Formally, we take as S the set of the domains $[a, b]$ of passes of σ, ordered naturally as internally disjoint subsets of $[0, 1]$, and let w_σ map every $[a, b] \in S$ to the directed chord that σ traverses on $[a, b]$. We call w_σ the *trace of* σ. Our aim is to show that $\langle\alpha\rangle \mapsto r(w_\alpha)$ defines a group embedding $\pi_1(H) \to F_\infty$.

For a proof that $\langle\alpha\rangle \mapsto r(w_\alpha)$ is well defined, consider homotopic loops $\alpha \sim \beta$ in H. We wish to show that $r(w_\alpha) = r(w_\beta)$. By Lemma 3.11 it suffices to show that $r(w_\alpha \restriction I) = r(w_\beta \restriction I)$ for every finite $I \subset \mathbb{N}$. Consider the space obtained from H by attaching a disc to H for every $j \notin I$, by an injective attachment map from the boundary of the disc onto the fundamental circle of e_j, the unique circle in $T + e_j$. This space deformation-retracts onto $T \cup \bigcup\{e_i \mid i \in I\}$, and hence is homotopy equivalent by Corollary 3.6 to a finite graph W_I having a spanning tree T_I with $|I|$ chords, one for each e_i. Composing α and β with the map $H \to W_I$ from this homotopy equivalence yields homotopic loops α' and β' in W_I, whose traces in F_I are $w_{\alpha'} = w_\alpha \restriction I$ and $w_{\beta'} = w_\beta \restriction I$. Since the map $\langle\gamma\rangle \mapsto r(w_\gamma)$ is known to be well defined for finite graphs, we deduce that

$$r(w_\alpha \restriction I) = r(w_{\alpha'}) = r(w_{\beta'}) = r(w_\beta \restriction I).$$

This completes the proof that $\langle\alpha\rangle \mapsto r(w_\alpha)$ is well defined. By (3.2), it is a homomorphism. For injectivity, we shall prove in Section 3.5 the following extension to paths that need not be loops:

Lemma 3.13. *Paths σ, τ in H with the same endpoints are homotopic in H if (and only if) their traces reduce to the same word.*

In Section 5.2 we shall show that Lemma 3.13 also follows from a result of Eda and Kawamura [25]. Nevertheless, we will show in Section 3.5 that the homotopy between σ and τ can be chosen so that it contracts pairs of passes, one at a time, like known from finite graphs. Our proof of Lemma 3.13 will thus show more: that either there is a 'natural' homotopy between σ and τ, one that essentially uses the combinatorial structure of the graph, or no homotopy at all.

We remark that the map $\langle\alpha\rangle \mapsto r(w_\alpha)$ will not normally be surjective. For example, $\vec{e}_0\vec{e}_1\cdots$ will always be a reduced word, but no loop in $|G|$ can pass through these chords in precisely this order if they do not converge to an end. Hence if two ends are non-trivial in H, then by Lemma 3.7 the sequence $\vec{e}_0, \vec{e}_1, \ldots$ of chords of T in H does not converge and therefore the reduced word $\vec{e}_0\vec{e}_1\cdots$ lies outside the image of our map $\langle\alpha\rangle \mapsto r(w_\alpha)$.

In order to describe the image of this map precisely, let us call a subword $w' := w \restriction S'$ of a word $w\colon S \to A$ *monotonic* if S' is infinite and can be written as $S' = \{s_0, s_1, \ldots\}$ so that either $s_0 < s_1 < \cdots$ or $s_0 > s_1 > \cdots$. Let us say that w' *converges* (in $|G|$) if there exists an end to which every sequence $x_0, x_1, \ldots$ with $x_n \in w(s_n)$ for all n converges. If w is the trace of a path in H, then by the continuity of this path all the monotonic subwords of w—and hence those of $r(w)$—converge.

We can now summarize our combinatorial description of $\pi_1(H)$ as follows.

Theorem 3.14. *Let G be a locally finite connected graph, and let H be a standard subspace of $|G|$. Let T be a topological spanning tree of H, and let $e_0, e_1, \ldots$ be its chords.*

(i) *The map $\langle\alpha\rangle \mapsto r(w_\alpha)$ is an injective homomorphism from $\pi_1(H)$ to the group F_∞ of reduced finite or infinite words in $\{\vec{e}_0, \vec{e}_1, \ldots\} \cup \{\overleftarrow{e}_0, \overleftarrow{e}_1, \ldots\}$. Its image consists of those reduced words whose monotonic subwords all converge in $|G|$.*

(ii) *The homomorphisms $w \mapsto r(w \restriction I)$ from F_∞ to F_I embed F_∞ as a subgroup in $\varprojlim F_I$. It consists of those elements of $\varprojlim F_I$ whose projections $r(w \restriction I)$ use each letter only boundedly often. (The bound may depend on the letter.)*

Proof. (i) We already saw that $\langle\alpha\rangle \mapsto r(w_\alpha)$ is a homomorphism, and injectivity follows from Lemma 3.13 (which will be proved in Section 3.5). We have also seen that for every loop α in H all the monotonic subwords of $r(w_\alpha)$ converge in $|G|$. It remains to show the converse: that if all the monotonic subwords of a reduced word w converge, then there is a loop α in H such that $w = r(w_\alpha)$.

We prove the following more general fact: If w is a word (not necessarily reduced) whose monotonic subwords all converge, then w is the trace of a loop in H. So let $w\colon S \to A$ be such a word. Enumerate S as $s_0, s_1, \ldots$.[4] We will inductively choose disjoint closed intervals $I_n \subseteq [0,1]$ ordered correspondingly, i.e. so that I_m precedes I_n in $[0,1]$ whenever $s_m < s_n$. For each n, we will let α_n be an order-preserving homeomorphism from I_n to the oriented chord $w(s_n)$. We will then extend the union of all the α_n to a loop $\alpha\colon [0,1] \to H$.

[4]Note that the enumeration is not related to the order of S.

In order that such a continuous extension α exist, we have to take some precautions when we choose the I_n. For example, suppose that the chords $w(s_0), w(s_2), \dots$ converge to one end, while the chords $w(s_1), w(s_3), \dots$ converge to another end. If S is ordered as $s_0 < s_2 < \cdots | \cdots < s_3 < s_1$ (note that every monotonic subword of w converges), there may be a point $x \in [0,1]$ such that every interval around x contains all but finitely many of the intervals I_n. In this case, any extension of $\bigcup_n \alpha_n$ will fail to be continuous at x. In order to prevent this, we shall first formalize such critical situations in terms of partitions of S, then prove that there are only countably many of them, reserve open intervals as padding around potentially critical points such as x, and finally choose the I_n so as to avoid these intervals.

Consider a partition P of S into non-empty parts $S^-(P)$ and $S^+(P)$, such that $s^- < s^+$ whenever $s^- \in S^-(P)$ and $s^+ \in S^+(P)$. Note that for all cofinal sequences in $S^-(P)$, finite or infinite, the final vertices of the corresponding chords of T converge to a common point $z^-(P) \in H$: otherwise there would be a monotonic subword of w that does not converge in $|G|$. Likewise, there is a point $z^+(P) \in H$ such that for all coinitial sequences in $S^+(P)$ the first vertices of the corresponding chords converge to $z^+(P)$. We call P *critical* if $z^-(P) \neq z^+(P)$.

Let us show that there are only countably many critical partitions. Suppose not, and note that there are only countably many critical partitions P for which $S^-(P)$ has a greatest element, since S is countable, and likewise $S^+(P)$ has a least element only countably often. Hence there are uncountably many P for which neither $S^-(P)$ has a greatest element nor $S^+(P)$ has a least element. For each such P, both $z^-(P)$ and $z^+(P)$ are ends. Let $W(P)$ be a finite set of vertices that separates $z^-(P)$ from $z^+(P)$ in $|G|$. As G is countable, it contains only countably many finite vertex sets, and hence there is a set W such that $W = W(P)$ for uncountably many P. Pick a sequence $P_0, P_1, \dots$ of critical partitions with $W(P_i) = W$, and so that either $S^-(P_0) \subsetneq S^-(P_1) \subsetneq \cdots$ or $S^+(P_0) \subsetneq S^+(P_1) \subsetneq \cdots$. We assume that $S^-(P_0) \subsetneq S^-(P_1) \subsetneq \cdots$, the other case being analogous. To obtain a contradiction, let us use this sequence to construct a non-convergent monotonic subword of w.

Choose $s_0^- \in S^-(P_0)$ and $s_0^+ \in S^+(P_0) \cap S^-(P_1)$ so that W separates $w(s_0^-)$ from $w(s_0^+)$ in G; this is possible, since W separates $z^-(P_0)$ from $z^+(P_0)$. Then for $i = 1, 2, \dots$ in turn choose $s_i^- \in S^-(P_i)$ and $s_i^+ \in S^+(P_i) \cap S^-(P_{i+1})$ so that $s_i^- > s_{i-1}^+$ and W separates $w(s_i^-)$ from $w(s_i^+)$ in G. Then $w \restriction \{s_0^-, s_0^+, s_1^-, s_1^+, \dots\}$ is a monotonic subword of w that does not converge in $|G|$, since W separates all the pairs $w(s_i^-), w(s_i^+)$. This completes the proof that there are only countably many critical partitions; enumerate them as $P_0, P_1, \dots$.

We now construct α. Inductively choose disjoint, closed, non-trivial intervals $I_n, J_n \subseteq [0,1]$ so that I_m precedes I_n on $[0,1]$ whenever $s_m < s_n$, and so that I_m precedes J_n if and only if $s_m \in S^-(P_n)$. For each n, let α_n be an order-preserving homeomorphism from I_n to the oriented chord $w(s_n)$. Extend the union of all the α_n to a loop α, as follows. Consider the connected components I of $[0,1] \setminus \bigcup I_n$. These are again intervals, possibly trivial; we call them *connecting intervals*. We shall first define α on the boundary points $a \le b$ of each I, and then extend it continuously to map $\overline{I}$ onto the arc $\alpha(a)T\alpha(b)$. In order to make α continuous, we will have to make sure when we choose $\alpha(a)$ and $\alpha(b)$ that the following is satisfied for $x = a$ or $x = b$:

For every boundary point x of a connecting interval, and every neighbourhood U of $\alpha(x)$, there is an $\varepsilon > 0$ such that $\alpha(y) \in U$ whenever $y \in (x-\varepsilon, x+\varepsilon)$ lies in an interval I_n. (3.4)

Suppose first that I is not an initial or final segment of $[0,1]$, i.e., that the boundary points a, b of I satisfy $0 < a \le b < 1$. Now the sets

$$S^- := \{s_n \mid I_n \text{ precedes } I\} \text{ and } S^+ := \{s_n \mid I \text{ precedes } I_n\}$$

form a partition P of S into non-empty parts. If P is critical, then $P = P_m$ for some m and hence $J_m \subseteq I$ by the choice of I as a component of $[0,1] \setminus \bigcup I_n$. In particular, I is non-trivial in this case. Let α map a to $z^-(P)$ and b to $z^+(P)$. (Note that if one of these points, a say, is a boundary point of some I_n, then $\alpha_n(a) = z^-(P)$, so α coincides with α_n on I_n.) If P is not critical, we have $z^-(P) = z^+(P)$ and let α map a and b to this point.[5] In both these cases, (3.4) follows easily from the definition of $z^-(P)$ and $z^+(P)$.

Now consider the case $a = 0$. If S has a least element s_n, then b is a boundary point of I_n on which α is already defined. Otherwise, S has a coinitial sequence $\cdots < s_{n_1} < s_{n_0}$, in which case the restriction of w to this sequence is a monotonic subword of w; this converges in $|G|$ to some end ω, and we put $\alpha(b) = \omega$. Note that ω is independent of the choice of the coinitial sequence of S, as otherwise there would be a monotonic subword of w that does not converge. In both cases, our choice of $\alpha(b)$ satisfies (3.4). As $\alpha(a)$, if $a \neq b$, we choose any vertex or end in H, satisfying (3.4) trivially. In the case of $b = 1$, we proceed analogously.

For each I, we now extend α to $\overline{I} \to \alpha(a)T\alpha(b)$ as planned. In particular, if $\alpha(a) = \alpha(b)$ we let α map all of I to that point.

By construction, α is continuous on the interior of each interval I_n and of each connecting interval. Hence it remains to check that α is continuous at boundary points of such intervals. (Recall that every point in $[0,1]$ either lies in some I_n or in an connecting interval.) Suppose first that x is the

[5] In particular, if I is trivial, the image of $a = b$ is well defined.

boundary point of a connecting interval. Let $x_0 < x_1 < \cdots$ be a sequence of points in $[0,1]$ converging to x. If all but finitely many of these lie in intervals I_n, their images under α converge to $\alpha(x)$ by (3.4). If not, we may assume that each x_i lies in some connecting interval $I^i = [y_i, z_i]$. At most one of these intervals I contains infinitely many x_j (because then it contains x), whose values then converge to $\alpha(x)$ by the continuity of $\alpha \restriction I$. Disregarding these x_j, we may thus assume that each I^i contains no x_j with $j \neq i$.

Let us show that the sequence $\alpha(y_0), \alpha(z_0), \alpha(y_1), \alpha(z_1), \ldots$ converges to $\alpha(x)$. Let a basic open neighbourhood U of $\alpha(x)$ be given, and let ε be as provided by (3.4). For all but finitely many i we have $y_i, z_i \in (x-\varepsilon, x+\varepsilon)$, and we claim that $\alpha(y_i), \alpha(z_i) \in U$ for such i. By definition of y_i (the case of z_i is analogous), there is a sequence of boundary points of intervals I_n that converges to y_i, and we may choose this sequence in $(x-\varepsilon, x+\varepsilon)$. By (3.4), α maps these points to vertices converging to $\alpha(y_i)$. By our choice of ε, these vertices lie in U. As every sequence of vertices in U converges to a point in U, by Lemma 2.4, we obtain $\alpha(y_i) \in U$ as desired. We now apply Lemma 3.3 to the sequence $\alpha(y_0), \alpha(z_0), \alpha(y_1), \alpha(z_1), \ldots$: By the lemma, the entire arcs $\alpha(I^i) = \alpha(y_i) T \alpha(z_i)$ converge to $\alpha(x)$. In particular, $\alpha(x_i) \to \alpha(x)$ as desired.

Suppose now that x is not the boundary point of a connecting interval but of an interval I_n, say $I_n = [x, y]$. Then the sets $S^- := \{s_m \mid I_m \text{ precedes } x\}$ and $S^+ := \{s_m \mid x \text{ precedes } I_m\}$ form a partition P of S, with $s_n = \min S^+$. If $S^- = \emptyset$, then $x = 0$ and continuity at x is trivial. Otherwise, $S^-(P)$ has a greatest element or P is critical. In both cases, x would be the boundary point of a connecting interval, a contradiction.

Note that α as defined here does not need to be a loop. But we can turn it into a loop without changing its trace, by appending to it a path in T from $\alpha(1)$ to $\alpha(0)$.

(ii) This was proved at the end of Section 3.3. □

Corollary 3.15. *Let $H \subseteq H'$ be standard subspaces of $|G|$. Then $\pi_1(H)$ is a subgroup of $\pi_1(H')$.*

Proof. Let T be a topological spanning tree of H. By Lemma 3.4, T extends to a topological spanning tree T' of H'. Since every chord of T is also a chord of T', Theorem 3.14 (i) implies the assertion. □

Let us call a standard subspace H of $|G|$ *non-trivial* if it contains an end (of G) that is non-trivial in H, otherwise H is *trivial*. The fundamental groups of all such spaces contain and are contained in the abstract group F_∞ defined in Section 3.3; recall that this group is independent of G, as long as $|G|$ itself is non-trivial. Indeed:

Corollary 3.16. *For every non-trivial standard subspace H of any $|G|$ there are subgroup embeddings $F_\infty \hookrightarrow \pi_1(H) \hookrightarrow F_\infty$.*

Proof. Theorem 3.14 (i) says that $\pi_1(H)$ is a subgroup of F_∞. Conversely, let T be a topological spanning tree of H. Since H is non-trivial, T has a sequence of chords in H that converge to an end $\omega \in H$ of G (Lemma 3.7). The union H' of T with all these chords is a standard subspace of $|G|$ contained in H, so $\pi_1(H') \leq \pi_1(H)$ by Corollary 3.15. Since T is a topological spanning tree of H' all whose (infinite subwords of words of) chords converge in $|G|$, Theorem 3.14 (i) implies that $\pi_1(H')$ is isomorphic to F_∞. □

Corollary 3.17. *The fundamental group $\pi_1(H)$ of a standard subspace H of $|G|$ is free if and only if H is trivial. In particular, $\pi_1(|G|)$ is free if and only if every end has a contractible neighbourhood in $|G|$.*

Proof. Let T be a topological spanning tree of H. If H is trivial, then T has only finitely many chords in H: otherwise some of them would converge to a end, which would be non-trivial in H by Lemma 3.7. By Corollary 3.6, H is homotopy equivalent to a finite graph whose fundamental group is the free group on this set of chords.

Conversely, let us assume that H is non-trivial and show that $\pi_1(H)$ is not free. By Corollary 3.16, $\pi_1(H)$ contains F_∞ as a subgroup, so by Theorem 2.13 it suffices to prove that F_∞ is not free. As pointed out in [10], this was shown by Higman [35]. □

When H is non-trivial, then $\pi_1(H)$ is uncountable, so any representation needs uncountably many generators. We believe that one also needs uncountably many relations, but have no proof of this.

Theorem 3.14 provides a reasonably complete solution to our original graph-theoretical problem, which asked for a canonical combinatorial description of the fundamental group of $|G|$ for given G. However, it does not answer the group-theoretical question of how interesting or varied the groups occurring as $\pi_1(|G|)$ or $\pi_1(H)$ are.

We close with some evidence that this question may indeed be interesting. For all we know so far, F_∞ might be the only abstract group ever occuring as $\pi_1(H)$ for non-trivial H. However, this is far from the truth:

Theorem 3.18. *The fundamental group $\pi_1(H)$ of a standard subspace H of $|G|$ is isomorphic to F_∞ if and only if H contains precisely one end of G that is non-trivial in H.*

For the proof of Theorem 3.18 we need a lemma about F_∞. By Theorem 4.1 of Conner and Eda [12], for any homomorphism from F_∞ to a free product $A * B$ there are finitely generated subgroups A' of A and B' of B such that the image of the homomorphism is contained in $A' * B$ or in $A * B'$. Hence,

Lemma 3.19. *If groups A, B are not finitely generated then there is no epimorphism $F_\infty \to A * B$.* $\square$

Proof of Theorem 3.18. Suppose first that H contains precisely one end of G that is non-trivial in H. Then every sequence of chords in H in which each chord appears only finitely often converges to this end. The group embedding $\pi_1(H) \to F_\infty$ from Theorem 3.14 (i) is therefore surjective.

Now suppose that $\pi_1(H)$ is isomorphic to F_∞. Then H is non-trivial (see the proof of Corollary 3.17), so H contains an end of G that is non-trivial in H. It remains to show that there cannot be two such ends, ω_1 and ω_2 say. Let T be a topological spanning tree of H, and pick a finite set S of vertices separating ω_1 from ω_2. Consider the graph $G' = G/S$ obtained from G by deleting any edges between vertices in S and identifying all the vertices in S to a new vertex v_S. Since rays in G have a tail in G' and vice versa, there is an obvious bicontinuous bijection between the ends of G and those of G', and we shall not distinguish them notationally. Let H/S and T/S be the quotient spaces of H and T defined analogously.

Our next aim is to construct a standard subspace H' of $|G'|$ and a topological spanning tree T' of H' such that the chords of T' are precisely the chords of T other than those between vertices in S. Clearly T/S is a path-connected subspace of H/S that contains all its vertices and ends, but T/S can contain circles. However, all these contain v_S, so by deleting some edges at v_S we can make T/S into a topological spanning tree T' of H/S. As we do not want T' to have chords in H/S that are not chords of T, we remove the same edges at v_S also from H/S, to obtain a standard subspace H' of $|G'|$ of which T' is still a topological spanning tree.

It is easy to see that the chords of T' are precisely the chords of T that do not join two vertices in S. We enumerate the chords of T as $e_0, e_1, \dots$ so that the chords of T' are precisely $e_n, e_{n+1}, \dots$ for some $n \in \mathbb{N}$.

The vertex v_S separates the ends ω_1 and ω_2 in G', so $|G'| \setminus \{v_S\}$ is the disjoint union of open sets O_1, O_2 with $\omega_1 \in O_1$ and $\omega_2 \in O_2$. For $i = 1, 2$ let $H_i := \{v_S\} \cup (H' \cap O_i)$. Both H_i are non-trivial, and hence their fundamental groups are not finitely generated. As $H_1 \cap H_2 = \{v_S\}$, Theorem 2.14 yields that $\pi_1(H') \simeq \pi_1(H_1) * \pi_1(H_2)$.

We now define an epimorphism $f\colon \pi_1(H) \to \pi_1(H')$; as $\pi_1(H) \simeq F_\infty$ by assumption, this will induce an epimorphism $F_\infty \to \pi_1(H_1) * \pi_1(H_2)$ contradicting Lemma 3.19. By Theorem 3.14 (i), every element a of $\pi_1(H)$ corresponds to a reduced word w_a in $F_\infty(\{\vec{e}_0, \vec{e}_1, \dots\})$ all of whose monotonic subwords converge in $|G|$. Consider the word $r(w_a \restriction \{n, n+1, \dots\})$. This word corresponds to an element of $\pi_1(H')$: its monotonic subwords are subwords of w_a, so they converge in $|G|$ and hence also in $|G'|$. Hence by (3.2) and Theorem 3.14 (i) the map $w_a \mapsto r(w_a \restriction \{n, n+1, \dots\})$ induces a homomorphism $f\colon \pi_1(H) \to \pi_1(H')$. For the same reason f is surjective: by Theorem 3.14 (i), every element of $\pi_1(H')$ corresponds to a reduced word w in $\{\vec{e}_n, \vec{e}_{n+1}, \dots\} \cup \{\overleftarrow{e}_n, \overleftarrow{e}_{n+1}, \dots\}$ all of whose monotonic subwords converge in $|G'|$, and hence also in $|G|$. Therefore $w = w_a$ for some $a \in \pi_1(H)$, by Theorem 3.14 (i). $\square$

3.5 Proof of the main lemma

We conclude our proof of Theorem 3.14 with the proof of Lemma 3.13. In this proof we shall need another lemma:

Lemma 3.20. *Let σ be a path in H and let $w_\sigma\colon S \to A$ be its trace. Then for every $S' \subseteq S$ there is a path τ in H with the same endpoints as σ and such that $w_\tau = w_\sigma \restriction S'$. Moreover, τ can be chosen so that $\tau \restriction [a,b] = \sigma \restriction [a,b]$ for every domain $[a,b] \in S'$ of a pass of τ.*

Proof. Note that the first statement follows from the more general fact that we proved in Theorem 3.14 (i): As w_σ is the trace of a path, all its monotonic subwords converge in H. Hence also all monotonic subwords of $w_\sigma \restriction S'$ converge in H, and thus it is the trace of a path τ. To ensure that τ coincides with σ on every domain of a pass of τ, we will construct τ explicitely.

For every $x \in [0,1]$ that lies in an interval in S' we define $\tau(x) := \sigma(x)$. Further, we put $\tau(0) := \sigma(0)$ and $\tau(1) := \sigma(1)$. Then for every $x \in [0,1]$ with $\tau(x)$ still undefined there is a unique maximal interval $[a,b]$ that contains x and is disjoint from (s,t) for every $[s,t] \in S'$. It is easy to see that $\sigma(a)$ and $\sigma(b)$ are vertices or ends, and hence lie in T. If $a \neq b$, we call $[a,b]$ a *non-traversing interval* and define τ on $[a,b]$ as a path from $\sigma(a)$ to $\sigma(b)$ whose image is precisely the arc in T between these two points. If $a = b$, then we let $\tau(a) := \sigma(a)$. Clearly $w_\tau = w_\sigma \restriction S'$ and $\tau \restriction [a,b] = \sigma \restriction [a,b]$ for each $[a,b] \in S'$; it remains to show that τ continuous.

Continuity is clear at inner points of intervals in S' or of non-traversing intervals, so let x be any other point. By definition, $\tau(x) = \sigma(x)$. It is easy to see that $\sigma(x)$ is a vertex or an end, so $\tau(x) \in T$. Continuity of τ at x is now follows easily: Given a neighbourhood U of $\tau(x)$, Lemma 3.2 gives us an arc-connected neighbourhood $U' \subseteq U$. As σ is continuous, there is an open interval I around x which σ maps to U'. Now also $\tau(I) \subseteq U'$ unless I meets non-traversing intervals that are not contained in I. In this case we can use the continuity of τ on those intervals to find an interval $I' \subseteq I$ which τ maps to U'. (Note that there are at most two such intervals.) $\square$

We are now ready for the proof of Lemma 3.13. The proof that $\alpha \sim \beta$ implies $r(w_\alpha) = r(w_\beta)$ was already shown for the case that α and β are loops. The general case follows, since σ and τ can be made into loops by appending a path in T joining their endpoints, which does not change their traces. It remains to prove the converse: we assume that $r(w_\sigma) = r(w_\tau)$, and show that $\sigma \sim \tau$.

Our aim is to construct a homotopy $F = (f_t)_{t\in[0,1]}$ of paths f_t in H with $f_0 = \sigma$ and $f_1 = \tau$. We first assume that τ does not traverse any chords; the general statement will then follow from this case. Our proof of this case will consist of the following four parts. We begin with some simplifications of the problem, straightening σ and τ to homotopic but less complicated paths. We then pair up the passes of σ through chords, with a view to cancel such pairs $(\vec{e}, \overleftarrow{e})$ by a local homotopy $f_t \to f_{t'}$ that retracts a small segment of f_t running through $\vec{e}$ and back through $\overleftarrow{e}$ without traversing any other chords. These pairs of passes have to be nested in the right way, and we will have to determine a (time) order in which to cancel them. Defining the pairing and the ordering of the pairs will be the second part of the proof. The ordering will have to list inner pairs before outer pairs, but among all the linear orders doing this we have to find a suitable one: since the partial order of the nestings of pairs can have limits, so will the linear order of times t at which to cancel the pairs. At these limits we may encounter discontinuities in our homotopy. But we shall be able to choose the order of cancellations so that this happens only countably often. In the third step we handle with those discontinuities and smooth them out by inserting further local homotopies $f_t \to f_{t'}$ at those countably many times t. In the last part of the proof we finally show that the homotopy F thus defined is indeed continuous.

3.5.1 Straightening σ and τ

Although τ does not, by assumption, traverse chords, its image might still contain inner points of chords. Let us call a path α in H *straight* if it does not have this property, i.e. if $\alpha(x) \in T$ for every x not contained in the domain of a pass. Applying Lemma 3.1 to local homotopies retracting any segments of α visiting a chord $e = uv$ to a constant map with image u or v, we can make any path in H straight without touching its passes:

Every path α in H is homotopic in $\operatorname{Im} \alpha$ *to a straight path α' that has the same passes as α.* (3.5)

By 'having the same passes' we mean not only that α, α' have the same trace but that for every interval $[a, b] \subseteq [0, 1]$ either both segments $\alpha \restriction [a, b]$ and $\alpha' \restriction [a, b]$ are a pass of their respective path or neither is, and if both are then $\alpha \restriction [a, b] = \alpha' \restriction [a, b]$.

By (3.5), σ and τ are homotopic to straight paths σ' and τ' such that σ' has the same passes as σ while τ' has the same passes as τ (namely, none). In particular, $w_{\sigma'} = w_\sigma$ and $w_{\tau'} = w_\tau = \emptyset$, so

$$r(w_{\sigma'}) = r(w_\sigma) = r(w_\tau) = r(w_{\tau'})$$

by assumption. We may therefore assume that σ and τ are straight, otherwise replace them by σ' and τ'.

Let us start with the simplest case: assume that σ also traverses no chord. Then σ and τ are homotopic. Indeed:

> *Let α and β be paths in T with identical endpoints $x, y \in V(G) \cup \Omega(G)$. Then there is a homotopy between α and β in $\operatorname{Im}\alpha \cup \operatorname{Im}\beta$. If β is constant, this homotopy can be chosen time-injective.* (3.6)

To prove (3.6), we construct homotopies of α and β to an x–y path in xTy. We shall assemble these two homotopies from homotopies between the segments $\gamma\colon [a,b] \to T$ of α or β that are maximal with $\gamma\big((a,b)\big) \subseteq T \setminus xTy$, and constant maps, defined as follows. Since $T \setminus xTy$ is open in T, the maximality of $[a,b]$ implies that $\gamma(a), \gamma(b) \in xTy$. Let us show that $\gamma(a) = \gamma(b)$. If not, then these two points are joined by two arcs that have disjoint interiors: one in xTy, and another in $\operatorname{Im}\gamma$ (Lemma 2.1). These two arcs would form a circle in T, which does not exist since T is a topological spanning tree of H. By Lemma 3.5, there is a time-injective homotopy in $\operatorname{Im}\gamma$ from γ to the constant map $[a,b] \to \{\gamma(a)\}$ $(= \{\gamma(b)\})$. By Lemma 3.1 we can combine all these homotopies, one for every γ, and obtain time-injective homotopies (in $\operatorname{Im}\alpha$ and $\operatorname{Im}\beta$) from α and β to x–y paths in xTy. If β is constant (with image $x = y$), the first of these is the desired time-injective homotopy from α to β. Otherwise we note that since $xTy \simeq [0,1]$, the latter two paths are homotopic in xTy, and we can combine our three homotopies to the desired homotopy between α and β.

Using Lemma 3.1 to apply (3.6) to segments between passes, we obtain the following generalization:

> *If paths α and β in H with identical endpoints have the same passes, then there is a homotopy between them in $\operatorname{Im}\alpha \cup \operatorname{Im}\beta$.* (3.7)

Let us assume now that σ traverses chords. Then $[0,1]$ is the disjoint union of the following intervals: the interiors of domains of passes, and the components of the rest of $[0,1]$. Every such component is a closed interval. By (3.7), we may assume that

> *If $(a,b) \subset [0,1]$ is maximal with the property that it avoids all domains of passes of σ, then σ maps $[a,b]$ onto the arc $\sigma(a)T\sigma(b)$. In particular, if $\sigma(a) = \sigma(b)$ then σ is constant on $[a,b]$.* (3.8)

This completes part one of the proof.

3.5.2 Ordering pairs of passes

Our next step towards the construction of the desired homotopy F between σ and τ will base upon the fact that w_σ reduces to $w_\tau = \emptyset$. So let us consider a reduction R of w_σ to the empty word. By definition, R is a totally ordered set of disjoint pairs of positions of w_σ such that the elements of each $p \in R$ are adjacent in $S \setminus \bigcup\{q \in R \mid q < p\}$ and are mapped by w_σ to inverse letters $\vec{e}_i, \overleftarrow{e}_i$. The positions of w_σ, in this case, are the domains of the passes of σ. Let us write

$$\mathcal{P} := \big\{(\pi, \rho) \mid \exists (s,t) \in R : (\pi = \sigma \restriction s \text{ and } \rho = \sigma \restriction t)\big\}$$

for the set of pairs of passes corresponding to R, where the order of (s,t) is that from S (so the interval s precedes the interval t in $[0,1]$). Note that $\mathcal{P}$ is countable, because σ has only countably many passes. Our homotopy F will remove the passes of σ in pairs as specified by $\mathcal{P}$, one at a time. The order in which this is done will not necessarily be the ordering which R induces on $\mathcal{P}$, so let us for now think of $\mathcal{P}$ as an unordered set.

Let us fix some more notation for later use. Let α be a path in H, and let $p = (\pi, \rho)$ be a pair of passes of α through the same chord $e =: e(p)$, but in opposite directions.[6] Let $[a, a^-]$ be the domain of π and $[b^-, b]$ that of ρ, and assume that $a < a^- \le b^- < b$. (Note that this assumption is satisfied by every element of $\mathcal{P}$.) Then $\pi(a) = \rho(b) =: z$ and $\pi(a^-) = \rho(b^-) =: z^-$ are the two vertices of e. If $\alpha \restriction [a^-, b^-]$ traverses no chord, and if β is the path obtained from α by replacing its segment $\alpha \restriction [a, b]$ with the constant map $[a, b] \to \{z\}$, we say that β is obtained from α by *cancelling the pair p of passes*. Since a, a^-, b, b^-, z, and z^- depend only on the pair $p = (\pi, \rho)$ but not on the rest of α, we denote them by $a(p)$, $a^-(p)$, $b(p)$, $b^-(p)$, $z(p)$, and $z^-(p)$. Thus:

For all $p \in \mathcal{P}$, we have $\sigma\big(a(p)\big) = \sigma\big(b(p)\big) = z(p)$ and $\sigma\big(a^-(p)\big) = \sigma\big(b^-(p)\big) = z^-(p)$. These points are vertices and hence lie in T. (3.9)

We now wish to determine the order in which our homotopy F will cancel the pairs in $\mathcal{P}$. This order will have to satisfy an obvious necessary condition imposed by the relative position of the passes in these pairs. Indeed, our definition of $\mathcal{P}$ implies that, given two pairs $p, p' \in \mathcal{P}$, the intervals $\big(a(p), b(p)\big)$ and $\big(a(p'), b(p')\big)$ are either disjoint or nested; accordingly, we call p and p' *parallel* or *nested*. If p and p' are nested and $\big[a(p), b(p)\big]$ contains $\big[a(p'), b(p')\big]$, we say that p *surrounds* p' and write $p \ge p'$. This is

[6] While p will always be an element of $\mathcal{P}$, the path α will in general not be σ but a path which has less passes than σ, i.e. every pass of α will be a pass of σ but not vice versa.

clearly a partial ordering. In fact, $\leq$ is the inverse of a tree order:

Whenever a pair is surrounded by two other pairs, these latter pairs are nested. (3.10)

This partial ordering $\leq$ on $\mathcal{P}$ will have to be respected by any order in which our homotopy F can cancel the pairs in $\mathcal{P}$: if p surrounds p', then p' has to be cancelled before p.

Our next aim, therefore, is to extend $\leq$ to a total ordering $\preceq$ on $\mathcal{P}$. We may think of $\preceq$ informally as the 'order in which F will cancel the pairs', but should remember that the order type of $\preceq$ may be arbitrarily complicated for a countable linear ordering. Note that, in order for our desired homotopy $F(x,t)$ to be continuous, it will not be enough to take as $\preceq$ the linear extension of $\leq$ defined by the reduction R: While the given order on R has to respect $\leq$, its relations between pairs that are incomparable in $\leq$ can be chosen almost arbitrarily. In order to be able to handle the discontinuities in Section 3.5.3 we shall need a more structured ordering.

To define $\preceq$, we start by enumerating the elements of $\mathcal{P}$ arbitrarily. Then, recursively for $i = 0, 1, \ldots$, let M_i be a maximal chain in $\mathcal{P} \setminus (M_0 \cup \ldots \cup M_{i-1})$ containing the first pair from this set in the enumeration of $\mathcal{P}$. (If there is no pair left we terminate the recursion, so all the chains M_i are non-empty.) Given two pairs p, p' (not necessarily distinct), let $p \preceq p'$ if either

- p, p' lie in the same M_i and $p \leq p'$; or
- $p \in M_j$ and $p' \in M_i$ with $i < j$.

Thus, $\preceq$ puts later chains below earlier chains. When $p \preceq p'$ we say that p *precedes* p', and p' *succeeds* p. By (3.10) and the maximality of the chains M_i, each of the sets $M^n := \bigcup_{i=0}^{n} M_i$ is closed upwards in $\leq$: if $p' \geq p \in M^n$ then also $p' \in M^n$. In words:

For all $i < j$, no pair in M_j surrounds a pair in M_i. (3.11)

By definition of $\preceq$, this implies that $\preceq$ is indeed a linear extension of $\leq$, i.e. that $p \preceq p'$ whenever $p \leq p'$: every pair precedes any pair that surrounds it.

Having partitioned the set $\mathcal{P}$ of passes of σ into pairs, and having chosen an order $\preceq$ in which we want our homotopy F to cancel them, we next wish to map our pairs p to time intervals $\big[s(p), t(p)\big] \subseteq [0,1]$ in which F can cancel p. These intervals will reflect $\preceq$ in that

$$t(p) < s(p') \text{ whenever } p \prec p'; \tag{3.12}$$

in particular, they will be disjoint for different p. While t runs through $\big[s(p), t(p)\big]$, the path f_t will change only on $\big[a(p), b(p)\big]$, so as to cancel p.

In order to state precisely what we require, we need another definition. Let α be a topological path in H, with $\alpha(a) = \alpha(b) =: z$ for some $a < b$, and let β be the path obtained from α by replacing $\alpha \restriction [a, b]$ with the constant map $[a, b] \to \{z\}$. We say that a homotopy from α to β *retracts* $\alpha \restriction [a, b]$ *to* z if it is relative to $[0, a] \cup [b, 1]$, time-injective, and maps $[a, b] \times [0, 1]$ to $\alpha\big([a, b]\big)$. If $\alpha\big([a, b]\big) \subseteq X$ for a subspace X of H, we may also say that this retraction is performed *in* X.

Our paths f_t will satisfy the following assertions:

The passes of f_t are also passes of σ; thus, $f_t \restriction [c, d] = \sigma \restriction [c, d]$ for every $t \in [0, 1]$ and every domain $[c, d]$ of a pass of f_t. (3.13)

For every $p \in \mathcal{P}$, the passes of $f_{s(p)}$ are exactly those passes of σ that are contained in pairs $p' \succeq p$; in particular, p is a pair of passes of $f_{s(p)}$. (3.14)

For every $p \in \mathcal{P}$, the path $f_{t(p)}$ is obtained from $f_{s(p)}$ by cancelling the pair p. This is achieved by a homotopy $(f_t)_{t \in [s(p), t(p)]}$ that first retracts $f_{s(p)} \restriction \big[a^-(p), b^-(p)\big]$ to $z^-(p)$ in T and then retracts the resulting path $f_t \restriction \big[a(p), b(p)\big]$ to $z(p)$ in $e(p)$. (3.15)

The first part of the homotopy in (3.15) will be obtained by applying (3.6), with $\alpha = f_{s(p)} \restriction \big[a^-(p), b^-(p)\big]$ and $\beta\colon \big[a^-(p), b^-(p)\big] \to \{z^-(p)\}$. This will turn $f_{s(p)}$ into a path f_t mapping $\big[a(p), b(p)\big]$ onto the chord $e(p)$, to which the second part of the homotopy is then applied.

Let us note for later use:

The path $f_{t(p)}$ maps $\big[a(p), b(p)\big]$ to the vertex $z(p)$. (3.16)

With this preview of how the linear order $\preceq$ of cancellations of passes will be implemented by F, we complete the second part of our proof.

3.5.3 Smoothing out the discontinuities

In the third part of the proof, we now turn to the reason why we have not chosen the intervals $\big[s(p), t(p)\big]$ explicitly yet. This is because observing the rules just outlined will not suffice to make our homotopy F continuous.

To see this, consider a bipartition $r = (P^-, P^+)$ of $\mathcal{P}$ into non-empty sets P^-, P^+ such that $p^- \prec p^+$ whenever $p^- \in P^-$ and $p^+ \in P^+$. Given $i \in \mathbb{N}$, let

$$P_i^+ := P^+ \cap M_i \quad \text{and} \quad P_i^- := P^- \cap M_i.$$

Since P^- is non-empty, P^+ meets only finitely many M_i: For if $q \in P^-$ lies in M_k, say, then by the definition of $\preceq$, P^+ cannot meet chains M_i with $i > k$. Denote the largest i with $P_i^+ \neq \emptyset$ by $i(r)$, and call it the *index* of r. Then $P_{i(r)}^+$ is an initial segment of P^+. Since the elements of M_i are nested, $\preceq$ coincides on $P_{i(r)}^+$ with $\leq$.

Let us call r *critical* if P^+ has no least element (with respect to $\preceq$). For critical partitions r, we define

$$[a^+, b^+] := \bigcap_{p \in P_{i(r)}^+} \big[a(p), b(p)\big].$$

Since $P_{i(r)}^+$ is countable, it has an (infinite) coinitial sequence $p_0^+ > p_1^+ > \cdots$. Then $\lim_n a(p_n^+) = a^+$. As σ is continuous, the vertices $z(p_n^+) = \sigma\big(a(p_n^+)\big)$ converge in H to $\sigma(a^+)$. As only finitely many of the vertices $z(p_n^+)$ can coincide by Lemma 2.9 and the fact that G is locally finite, $\sigma(a^+)$ must be an end, which we denote by

$$z(r) \in \Omega(G).$$

Call a point $x \in [a^+, b^+]$ *critical* (with respect to r) if $x \in \big(a(q), b(q)\big)$ also for some $q \in P^-$. Then

$$P_x^- := \big\{q \in P^- \mid x \in \big(a(q), b(q)\big)\big\} \neq \emptyset$$

is a $\leq$-chain, which may or may not have a greatest element. Put

$$(a_x^-, b_x^-) := \bigcup_{q \in P_x^-} \big(a(q), b(q)\big).$$

Every $q \in P_x^-$ is nested with every $p \in P_{i(r)}^+$, since x lies in both $\big(a(q), b(q)\big)$ and $\big(a(p), b(p)\big)$. As $q \preceq p$, this means that $q \leq p$, so

$$[a_x^-, b_x^-] \subseteq [a^+, b^+].$$

Let us give an example of such a path σ.

Example 3.21. Suppose that ω is a non-trivial end of H, and let $e_{n_0}, e_{n_1}, \ldots$ be a sequence of chords of T that converges to ω. Now let σ be a loop in H based at a vertex v that first traverses $\vec{e}_{n_2}, \vec{e}_{n_3}, \ldots$ and reaches ω at time $1/3$. On $[2/3, 1]$, we let σ go the same way backwards, i.e. $\sigma(x) := \sigma(1 - x)$ for $x \in [2/3, 1]$. Thus in $[2/3, 1]$, σ returns from ω to v, traversing $\ldots, \overleftarrow{e}_{n_3}, \overleftarrow{e}_{n_2}$. Between time $1/3$ and $2/3$, we let it traverse $\vec{e}_{n_0}, \vec{e}_{n_1}, \overleftarrow{e}_{n_1}, \overleftarrow{e}_{n_0}$, with domains of these passes $[0.35, 0.36]$, $[0.4, 0.41]$, $[0.5, 0.51]$, and $[0.64, 0.65]$, say. In $\mathcal{P}$, the unique pass through $\vec{e}_{n_i}$ forms a pair p_i with the unique pass through $\overleftarrow{e}_{n_i}$. Then in the order $\prec$, the pairs

p_0 and p_1 are smaller than alle the other pairs. Now consider the partition $r = (P^-, P^+)$ of $\mathcal{P}$ with $P^- = \{p_0, p_1\}$. Then $[a^+, b^+] = [1/3, 2/3]$, and the points $x \in (0.35, 0.65)$ are critical with respect to r. For $x \in (0.4, 0.51)$ we have $P^-_x = \{p_0, p_1\}$, for the other x in $(0.35, 0.65)$ we have $P^-_x = \{p_0\}$, but for all $x \in (0.35, 0.65)$ we have $(a^-_x, b^-_x) = (0.35, 0.65)$. The points in $(1/3, 2/3) \setminus (0.35, 0.65)$ are not critical with respect to r.

To see why such r and x are 'critical' for the construction of our homotopy F, assume now that F has been chosen so as to satisfy (3.12)–(3.15), but apart from this arbitrarily. Since $p^+_0 \succ p^+_1 \succ \cdots$ we then have $t^+_0 > t^+_1 > \cdots$, where $t^+_n := t(p^+_n)$ is the time at which F has just cancelled the pair p^+_n. Let

$$t^+ := \lim_n t^+_n = \inf_n t^+_n .$$

For every $x \in \big[a(p^+_n), b(p^+_n)\big]$, statement (3.16) yields $f_{t^+_n}(x) = z(p^+_n) = \sigma\big(a(p^+_n)\big)$. Every $x \in [a^+, b^+]$ satisfies this for all n, so for such x the continuity of F and σ imply

$$f_{t^+}(x) = \lim_n f_{t^+_n}(x) = \lim_n \sigma\big(a(p^+_n)\big) = \sigma(a^+) = z(r). \tag{3.17}$$

Now assume that x is critical. Since P^-_x is a countable $\leq$-chain, it contains a (finite or infinite) cofinal sequence $q_0 < q_1 < \cdots$. Then $\lim_n a(q_n) = a^-_x$. Let

$$t^-_x := \lim_n t(q_n) = \sup_n t(q_n).$$

As $q \prec p$ for all $q \in P^-$ and $p \in P^+$, we have $t^-_x \leq t^+$ by (3.12). As earlier, the fact that $x \in \big[a(q_n), b(q_n)\big]$ for all n implies

$$f_{t^-_x}(x) = \lim_n f_{t(q_n)}(x) \overset{(3.16)}{=} \lim_n \sigma\big(a(q_n)\big) = \sigma(a^-_x) =: z_x. \tag{3.18}$$

If P^-_x has a greatest element q_n, then z_x will be the vertex $z(q_n)$; if not, it will be an end. But this end need not be $z(r)$. And if $z_x \neq z(r)$, we shall have a problem: to avoid a contradiction between (3.17) and (3.18), we will have to ensure that $t^-_x \neq t^+$ (which does not follow from the assumptions we have made about F so far), and define $F(x, t)$ so as to move z_x to $z(r)$ in the time interval $[t^-_x, t^+]$. Let us call our critical partition *bad* if $z_x \neq z(r)$ for some critical point x, which we then also call *bad*.

In general, $\mathcal{P}$ may have uncountably many critical partitions, and a critical partition can have bad points x with infinitely many different z_x. It will be crucial for our construction of F, therefore, to prove that there can be only countably many *bad* partitions. For each of these, we shall be able to deal with all its bad points simultaneously.

For our proof that there are only countably many bad partitions, let us show first that

$P^-_{i(r)} \neq \emptyset$ for every bad partition $r = (P^-, P^+)$, indeed for every critical partition that has a critical point. (3.19)

To prove (3.19) let $i = i(r)$, let x be a critical point for r, and consider any $q \in P_x^-$. If $q \in M_i$ we are done. If not, then $q \in M_j$ for some $j > i$, because there are $p \in P_i^+$ with $q \prec p$. By the maximality of M_i, this means that there is a $p^* \in M_i$ with $p \not\geq q$. Since $q \leq p$ for all $p \in P_i^+$ (as x lies in both $\big(a(p), b(p)\big)$ and $\big(a(q), b(q)\big)$), it follows that p^* lies in $P_i^- = M_i \setminus P_i^+$, which thus is non-empty, as claimed.

By (3.19), the nested union

$$\big(a^-, b^-\big) := \bigcup_{p \in P^-_{i(r)}} \big(a(p), b(p)\big)$$

is not empty. As $q \leq p$ for all $q \in P_i^-$ and $p \in P_i^+$, we have $[a^-, b^-] \subseteq [a^+, b^+]$:

$$a^+ \leq a^- \quad \text{and} \quad b^- \leq b^+.$$

Assuming that r is bad, let us show that at least one of these inequalities is strict. Pick a bad point $x \in [a^+, b^+]$. Assume first that $x \in (a^-, b^-)$. Then x lies in $\big(a(p), b(p)\big)$ for some, and hence for all large enough, $p \in P^-_{i(r)}$. But then all these p lie in P_x^-, so $(a^-, b^-) \subseteq (a_x^-, b_x^-)$ and hence $a_x^- \leq a^-$. (In fact, we have equality, but the above inequation will be sufficient.) Recall that $a^+ \leq a_x^-$. Since x is a bad point, we have $\sigma(a_x^-) \neq \sigma(a^+)$ and therefore $a_x^- \neq a^+$. Hence $a^+ < a_x^- \leq a^-$, as desired.

On the other hand if $x \notin (a^-, b^-)$, then $x \in [a^+, a^-] \cup [b^-, b^+]$. We assume that $x \in [a^+, a^-]$ and show $a^+ < a^-$; the case of $x \in [b^-, b^+]$ is analogous, showing $b^- < b^+$. Pick $q \in P_x^-$. Then $x \in \big(a(q), b(q)\big) \subseteq \big(a(p), b(p)\big)$ for every $p \in P^+_{i(r)}$, so $a^+ \leq a(q) < x \leq a^-$ by the definition of a^+.

We have thus shown that, for every bad partition r, the set

$$D(r) := (a^+, a^-) \cup (b^-, b^+)$$

is non-empty. We next show that these sets are disjoint for distinct $r = (P^-, P^+)$ and $\tilde{r} = (\tilde{P}^-, \tilde{P}^+)$ with the same index i. As $r \neq \tilde{r}$, we may assume that there is a pair $p \in P^- \cap \tilde{P}^+$. Then $p \in M_i$, since $M_j \subseteq P^- \cap \tilde{P}^-$ for every $j > i$, while $M_j \subseteq P^+ \cap \tilde{P}^+$ for every $j < i$. But then

$$a^- \leq a(p) < \tilde{a}^+ < \tilde{b}^+ < b(p) \leq b^-$$

with the obvious notation. Since $D(r) \cap [a^-, b^-] = \emptyset$ while $D(\tilde{r}) \subseteq (\tilde{a}^+, \tilde{b}^+)$, we have $D(r) \cap D(\tilde{r}) = \emptyset$ as claimed.

As there are only countably many partition indices i, and for every bad partition r with index i the set $D(r)$ contains a rational, this completes our proof that there are only countably many bad partitions.

Denote the set of all critical and all bad partitions by $\mathcal{R}$ and $\mathcal{R}'$, respectively. Given $r \in \mathcal{R}$, write $a(r)$ and $b(r)$ for the points $a^+, b^+ \in [0,1]$ defined above. Taking limits in (3.9), we obtain:

For every $r \in \mathcal{R}$ we have $\sigma\big(a(r)\big) = \sigma\big(b(r)\big)$ $\big(= z(r) \in \Omega(G) \subseteq T\big)$. (3.20)

Our plan is to begin the construction of F by extending our linear ordering $\preceq$ to $\mathcal{P} \cup \mathcal{R}$. We shall then choose disjoint time intervals $\big[s(q), t(q)\big]$ for all $q \in \mathcal{P} \cup \mathcal{R}'$, extending (3.12) to $\mathcal{P} \cup \mathcal{R}'$. For every $r = (P^-, P^+) \in \mathcal{R}$, the choices made for $\mathcal{P}$ will define times

$$t_r^+ := \inf\{\, t(p) \mid p \in P^+ \} \quad \text{and} \quad t_r^- := \sup\{\, t(p) \mid p \in P^- \}.$$

In order to satisfy (3.17), we shall have to have $f_{t_r^+}$ map all of $\big[a(r), b(r)\big]$ to $\{z(r)\}$. For $r \in \mathcal{R} \setminus \mathcal{R}'$, this will be satisfied automatically. For bad r, this will require the insertion of a local homotopy $f_{s(r)} \to f_{t(r)}$ analogous to (3.15), where $s(r)$ and $t(r)$ are chosen so that

$$t_r^- \leq s(r) < t(r) \leq t_r^+.$$

To extend $\preceq$ from $\mathcal{P}$ to $\mathcal{P} \cup \mathcal{R}$, we simply place all the partitions $r \in \mathcal{R}$ at their natural positions in the chains $M_{i(r)}$. Indeed, let us extend our partial ordering $\leq$ on $\mathcal{P}$ to $\mathcal{P} \cup \mathcal{R}$ by letting $q \leq q'$ whenever $\big[a(q), b(q)\big] \subseteq \big[a(q'), b(q')\big]$. Then

$$\hat{M}_i := M_i \cup \{\, r \in \mathcal{R} \mid i(r) = i \,\}$$

is easily seen to be a $\leq$-chain. We now define $\preceq$ on $\mathcal{P} \cup \mathcal{R}$ as we did on $\mathcal{P}$: given q, q', we put $q \preceq q'$ if either

- q, q' lie in the same $\hat{M}_i$ and $q \leq q'$; or
- $q \in \hat{M}_j$ and $q' \in \hat{M}_i$ with $i < j$.

Let us note a couple of facts about this ordering. The fact that pairs are either nested or disjoint extends at once:

Given $q, q' \in \mathcal{P} \cup \mathcal{R}$ with $q \preceq q'$, either $q \leq q'$ or the intervals $\big(a(q), b(q)\big)$ and $\big(a(q'), b(q')\big)$ are disjoint. (3.21)

The following statements can be satisfied by suitable $p \in P^-_{i(r)}$, which we recall is non-empty if r has a critical point (see (3.19)):

For every $r \in \mathcal{R}$ that has a critical point, in particular for every $r \in \mathcal{R}'$, there is a pair $p \in \mathcal{P}$ such that $p \leq r$. Given any $r' \in \mathcal{R}$ with $r' \prec r$, this p can be chosen so that $r' \prec p \prec r$. (3.22)

It is not difficult to describe $\preceq$ directly, without reference to $\leq$:

$$\text{\textit{A partition } } (P^-, P^+) \in \mathcal{R} \text{ \textit{precedes all pairs in} } P^+ \text{ \textit{and succeeds all pairs in} } P^-. \text{ \textit{Two partitions} } r = (P^-, P^+) \text{ \textit{and} } \tilde{r} = (\tilde{P}^-, \tilde{P}^+) \text{ \textit{satisfy} } r \preceq \tilde{r} \text{ \textit{if and only if} } \tilde{P}^+ \subseteq P^+. \tag{3.23}$$

Having defined $\preceq$, we can now choose disjoint intervals $\big[s(q), t(q)\big]$ for all $q \in \mathcal{P} \cup \mathcal{R}'$, to satisfy

$$t(q) < s(q') \text{ \textit{whenever} } q \prec q'. \tag{3.24}$$

This can be done inductively, since $\mathcal{P} \cup \mathcal{R}'$ is countable.

We are finally ready to define our homotopy $F = (f_t)_{t\in[0,1]}$. We first define f_t for all $t \in [0,1]$ outside the set

$$C := \bigcup_{q\in\mathcal{P}\cup\mathcal{R}'} \big(s(q), t(q)\big).$$

Given such $t \in [0,1] \setminus C$, let

$$Q^t := \{\, q \in \mathcal{P} \cup \mathcal{R}' \mid t(q) \leq t \,\}$$

be set of all q whose time intervals took place before time t. For every $x \in [0,1]$, the set

$$Q^t_x := \{\, q \in Q^t \mid x \in \big(a(q), b(q)\big) \,\}$$

of all such q affecting x is a $\leq$-chain, by (3.21). If $Q^t_x \neq \emptyset$, we set

$$(a^t_x, b^t_x) := \bigcup_{q\in Q^t_x} \big(a(q), b(q)\big)$$

and define $f_t(x) := \sigma(a^t_x) = \sigma(b^t_x) \in T$. (Recall (3.9) and (3.20).) If $Q^t_x = \emptyset$, we call x *unchanged at time* t, define $f_t(x) := \sigma(x)$, and put $a^t_x := x =: b^t_x$. Thus,

$$\text{\textit{for all} } t \notin C \text{ \textit{and all} } x, \text{ \textit{we have} } f_t(x) = \sigma(a^t_x) = \sigma(b^t_x). \tag{3.25}$$

We need to show that these functions f_t are continuous. This will follow from the fact that σ is continous, once we have shown the following:

$$\text{\textit{For all} } t \notin C \text{ \textit{and all} } x \text{ \textit{with} } Q^t_x \neq \emptyset, \text{ \textit{the function} } f_t \text{ \textit{is constant on} } (a^t_x, b^t_x) \text{ \textit{with value} } f_t(x) \in T. \tag{3.26}$$

To prove (3.26), pick $y \in (a^t_x, b^t_x)$. We show that $Q^t_y \neq \emptyset$ and $(a^t_y, b^t_y) = (a^t_x, b^t_x)$; then $f_t(y) = f_t(x) \in T$ by definition of f_t. By the choice of y, there exists $q \in Q^t$ such that $\big(a(q), b(q)\big)$ contains both x and y, giving

$q \in Q^t_x \cap Q^t_y$. As Q^t_x is a $\le$-chain, we have $q \le q'$ for all large enough $q' \in Q^t_x$. Since $y \in \big(a(q), b(q)\big) \subseteq \big(a(q'), b(q')\big)$ implies $q' \in Q^t_y$, all large enough $q' \in Q^t_x$ lie in Q^t_y, giving $(a^t_x, b^t_x) \subseteq (a^t_y, b^t_y)$. Likewise, $(a^t_y, b^t_y) \subseteq (a^t_x, b^t_x)$ and hence $(a^t_x, b^t_x) = (a^t_y, b^t_y)$. This completes the proof of (3.26), and with it the proof that the f_t defined so far are continuous.

We have just shown that $(a^t_y, b^t_y) = (a^t_x, b^t_x)$ for all x and y with $y \in (a^t_x, b^t_x)$. Therefore, for any x, y the intervals (a^t_x, b^t_x) and (a^t_y, b^t_y) are either identical or disjoint: if (a^t_x, b^t_x) meets (a^t_y, b^t_y), in a point z say, we have $(a^t_x, b^t_x) = (a^t_z, b^t_z) = (a^t_y, b^t_y)$. An immediate consequence of this is the following:

For all $t \in [0,1] \setminus C$ and $x \in [0,1]$, the points a^t_x and b^t_x are unchanged at time t. (3.27)

It remains to define f_t for $t \in \big(s(q), t(q)\big)$ with $q \in \mathcal{P} \cup \mathcal{R}'$. Since these intervals are disjoint, $f_{s(q)}$ and $f_{t(q)}$ are already defined. For each q, our aim is to define the functions f_t with $t \in \big(s(q), t(q)\big)$ as a homotopy between $f_{s(q)}$ and $f_{t(q)}$. When $q \in \mathcal{R}'$, this homotopy should retract $f_{s(q)} \restriction \big[a(q), b(q)\big]$ to $z(q)$ in T by a direct application of (3.6). When $q \in \mathcal{P}$, our plan is to follow (3.15) and achieve the same result in two stages. We first wish to use (3.6) to retract $f_{s(q)} \restriction \big[a^-(q), b^-(q)\big]$ to $z^-(q)$ in T. This should turn $f_{s(q)} \restriction \big[a(q), b(q)\big]$ into a path consisting of the two passes of q through $e(q)$ at the beginning and end, and a constant path with image $z^-(p)$ in the middle. We then wish to retract this path to $z(q)$ in $e(q)$.

In order to apply (3.6) and implement (3.15) as just outlined, we have to verify the following prerequisites:

- that $f_{s(q)}$ maps both $a(q)$ and $b(q)$ to $z(q)$;
- that $f_{t(q)}$ maps all of $\big[a(q), b(q)\big]$ to $z(q)$;
- if $q \in \mathcal{P}$: that $f_{s(q)}$ agrees with σ on $D(q) := \big(a(q), a^-(q)\big) \cup \big(b^-(q), b(q)\big)$ and maps $\big[a^-(q), b^-(q)\big]$ to T;
- if $q \in \mathcal{R}'$: that $f_{s(q)}$ maps $\big[a(q), b(q)\big]$ to T.

Let us prove the first statement, as well as the second statement for $a(q)$ and $b(q)$. At times $s(q)$ and $t(q)$, both $a(q)$ and $b(q)$ were unchanged, because they can lie in an interval $\big(a(q'), b(q')\big)$ only when $q < q'$ and hence $t(q) < t(q')$. Therefore $f_{s(q)}$ and $f_{t(q)}$ both map $a(q)$ and $b(q)$ to $\sigma\big(a(q)\big) = \sigma\big(b(q)\big) = z(q)$. (Recall (3.9) and (3.20).)

Next, we prove the second statment for $x \in \big(a(q), b(q)\big)$. At time $t(q)$, none of these x was unchanged, since $q \in Q^{t(q)}_x$ for all these x. In fact, q is the greatest element of each of the $\le$-chains $Q^{t(q)}_x$. Therefore $(a^{t(q)}_x, b^{t(q)}_x) = \big(a(q), b(q)\big)$, and hence $f_{t(q)}(x) = \sigma\big(a(q)\big) = z(q)$ by (3.25).

To prove the first part of the third statement note that, if $q \in \mathcal{P}$, all the points in $D(q)$ are still unchanged at time $s(q)$. Hence $f_{s(q)}$ agrees with σ on $D(q)$.

To prove the rest of the third and the fourth statement, consider any point x in $\big[a(q), b(q)\big]$, but not in $D(q)$ if $q \in \mathcal{P}$. We have to show that $f_{s(q)}(x) \in T$. This follows from (3.26) if x is not unchanged at time $s(q)$, so assume that it is. Then $f_{s(q)}(x) = \sigma(x)$. If this point is not in T, then x is an inner point of the domain of a pass contained in some $p \le q$. If $p = q$ this means that $x \in D(q)$, contradicting our choice of x. Hence $p < q$, and $t(p) < s(q)$ by (3.24). Thus $p \in Q_x^{s(q)} \neq \emptyset$, contradicting our assumption that x was unchanged at time $s(q)$.

Having checked the prerequisites, we may now apply (3.6) as outlined earlier to choose f_t for all $t \in \big(s(q), t(q)\big)$ as follows:

For every $r \in \mathcal{R}'$, the paths $(f_t)_{t \in [s(r), t(r)]}$ form a homotopy retracting $f_{s(r)} \restriction \big[a(r), b(r)\big]$ to $z(r)$. (3.28)

For every $p \in \mathcal{P}$, the paths $(f_t)_{t \in [s(p), t(p)]}$ form a homotopy that first retracts $f_{s(p)} \restriction \big[a^-(p), b^-(p)\big]$ to $z^-(p)$ in T and then retracts the resulting path on $\big[a(p), b(p)\big]$ to $z(p)$ in $e(p)$. (3.29)

For our later proof that F is continuous, let us note an important property of the homotopies in (3.28) and (3.29). Let U be a neighbourhood in $|G|$ of an end ω; then $U \cap H$ is a neighbourhood of ω in H. By Lemma 3.3 there is a basic open neighbourhood $\hat{U} \subseteq U$ of ω in $|G|$ (i.e., $\hat{U} = \hat{C}(S, \omega)$ for some finite set S of vertices) such that for any $x, y \in \hat{U} \cap T$, the arc xTy is contained in U (and thus in $U \cap H$). Let S' be the set of neighbours of S in $C(S, \omega)$, note that these are finitely many. Call $U' := \hat{C}(S', \omega)$ a *core of U around ω*. If U' contains a vertex $z(p)$, then $\hat{U}$ contains its neighbour $z^-(p)$ and the edge $e(p)$. The next statement therefore follows from the fact that the homotopies used in (3.28) and (3.29) either run inside $e(p)$ or else are time-injective (by our definition of retracting).

Let $q \in \mathcal{P} \cap \mathcal{R}'$ and $x \in \big(a(q), b(q)\big)$, and let $U' \subseteq |G|$ be a core of a neighbourhood U around an end. If both $f_{s(q)}(x)$ and $f_{t(q)}(x)$ lie in U', then $f_t(x) \in U$ for all $t \in \big[s(q), t(q)\big]$. (3.30)

The definition of F is now complete. Note finally that

$$\operatorname{Im} F = \operatorname{Im} \sigma. \tag{3.31}$$

This completes part three of our proof.

3.5.4 The continuity proof

It remains to show that the family $(f_t)_{t\in[0,1]}$ of paths in H is a homotopy between σ and τ. Since $Q^0 = \emptyset$, we have $f_0 = \sigma$. Rather than proving that $f_1 = \tau$, let us show that $\operatorname{Im} f_1 \subseteq T$; then (3.6) and our assumption that $\operatorname{Im}\tau \subseteq T$ imply $f_1 \sim \tau$, which is good enough. For a proof that $\operatorname{Im} f_1 \subseteq T$ consider any $x \in [0,1]$. If x is not unchanged at time $t = 1$, then $f_1(x) \in T$ by (3.26). If it is, then $f_1(x) = \sigma(x)$. Since σ is straight, we have $f_1(x) \in T$ unless x lies in the interior of the domain of a pass of σ. But this pass is contained in some pair $p \in Q^1_x$, so x was not unchanged at time 1, contradiction.

We now prove that F is continuous. Let $x, t \in [0,1]$ be given, and let U be any neighbourhood of $F(x,t)$ in $|G|$; then $U \cap H$ is a neighbourhood of $F(x,t)$ in H. If $F(x,t)$ is an end, let U' be a core of U around that end. We shall find an $\varepsilon > 0$ such that $F\big((x-\varepsilon, x+\varepsilon), (t-\varepsilon, t+\varepsilon)\big) \subset U$, proceeding in two steps.

1. We find an ε for which $F\big((x-\varepsilon, x+\varepsilon), (t-\varepsilon, t]\big) \subseteq U$.

   For every $\varepsilon > 0$, let

   $$Q_\varepsilon := \big\{q \in \mathcal{P} \cup \mathcal{R}' \mid (t-\varepsilon, t] \cap \big(s(q), t(q)\big) \neq \emptyset\big\}.$$

   If $Q_\varepsilon = \emptyset$ for some ε, then for all $t' \in (t-\varepsilon, t]$ we have $Q^{t'} = Q^t$ and hence $f_{t'} = f_t$. As f_t is continuous, there is an $\varepsilon_0 < \varepsilon$ such that $F\big((x-\varepsilon_0, x+\varepsilon_0), (t-\varepsilon_0, t]\big) \subseteq U$.

   If Q_ε is never empty but finite for some ε, there exists $q \in \mathcal{P}\cup\mathcal{R}'$ such that $t \in \big(s(q), t(q)\big]$. Since $(f_t)_{t\in[s(q),t(q)]}$ was defined as a homotopy, in (3.28) or (3.29), we have $F\big((x-\varepsilon_0, x+\varepsilon_0), (t-\varepsilon_0, t]\big) \subseteq U$ for small enough $\varepsilon_0 < \varepsilon$.

   We may thus assume that Q_ε is infinite for every $\varepsilon > 0$, so it has no maximal element with respect to $\preceq$. By the definition of $\preceq$, we can choose ε_0 small enough that all pairs in Q_{ε_0} lie in the same chain M_i. Then $i = i(r)$ for all partitions r in Q_{ε_0}, e.g. by (3.19). Hence $Q_{\varepsilon_0} \subseteq \hat{M}_i$, so the intervals $\big(a(q), b(q)\big)$ with $q \in Q_{\varepsilon_0}$ are nested; put

   $$(a,b) := \bigcup_{q\in Q_{\varepsilon_0}} \big(a(q), b(q)\big).$$

   By (3.9) and (3.20) every point $a(q)$ with $q \in Q_{\varepsilon_0}$ is mapped by σ to a vertex or an end, and by Lemma 2.9 and the fact that G is locally finite every vertex can appear only finitely often as the image of such an $a(q)$. Thus, $\sigma(a)$ is the limit of an infinite sequence of ends or distinct vertices $\sigma\big(a(q)\big) = z(q)$, so $\sigma(a)$ must be an end.

Our assumption that Q_{ε_0} has no maximal element also implies that $t \notin C$, so $F(x,t) = f_t(x) = \sigma(a_x^t)$ by (3.25). Moreover, there are points $t' \notin C$ arbitrarily close below t. If $F(x,t)$ is an end, it will suffice to find an $\varepsilon < \varepsilon_0$ such that $t - \varepsilon \notin C$ and $F(x',t') \in U'$ for all $x' \in (x-\varepsilon, x+\varepsilon)$ and all $t' \in [t-\varepsilon, t] \setminus C$: then for all such x' and all $t'' \in (t-\varepsilon, t) \cap C$ we shall have $F(x', t'') \in U$ by (3.30).

We distinguish two cases: that x lies in (a,b) or not.

(i) Suppose that $x \in (a,b)$. Then $a_x^t = a$, so (3.25) yields $F(x,t) = \sigma(a)$, which we know is an end. Since σ is continuous, there is a $\delta > 0$ such that σ maps $[a, a+\delta)$ to U'. Choose $\varepsilon_1 < \varepsilon_0$ so that for all $q \in Q_{\varepsilon_1}$ we have $a(q) \in [a, a+\delta)$ and $(x - \varepsilon_1, x + \varepsilon_1) \subseteq \big(a(q), b(q)\big)$.

Pick $q_0 \in Q_{\varepsilon_1}$. Choose $\varepsilon_2 < \varepsilon_1$ small enough that $t - \varepsilon_2 > t(q_0)$, and so that $t - \varepsilon_2 \notin C$. Then for all $x' \in (x - \varepsilon_2, x + \varepsilon_2)$ and $t' \in [t - \varepsilon_2, t] \setminus C$ we have $q_0 \in Q_{x'}^{t'} \neq \emptyset$, and (3.25) yields $F(x',t') = \sigma(a_{x'}^{t'})$. These points lie in U', since $a \le a_{x'}^{t'} \le a(q_0) < a + \delta$.

(ii) Suppose now that $x \notin (a,b)$, say $x \le a$. Note that for all $x' \notin (a,b)$ and $t' \in (t - \varepsilon_0, t] \setminus C$ we have $Q_{x'}^{t'} = Q_{x'}^{t}$, and hence $f_{t'}(x') = f_t(x')$.

If $x = a$, pick $q_0 \in Q_{\varepsilon_0}$. Then $f_{t(q_0)}(x) = f_t(x) \in U'$. (Note that as $a = a_y^t$ for all $y \in (a,b)$, we have $f_t(x) = f_t(a) = \sigma(a)$ by (3.27), hence $f_t(x)$ is an end, so U' is defined.) As $f_{t(q_0)}$ is continuous, there is an $\varepsilon_1 \le t - t(q_0)$ $(\le \varepsilon_0)$ such that $t - \varepsilon_1 \notin C$ and $f_{t(q_0)}\big((x - \varepsilon_1, x + \varepsilon_1)\big) \subseteq U'$. We show that $F(x',t') \in U'$ for all $x' \in (x - \varepsilon_1, x + \varepsilon_1)$ and $t' \in [t - \varepsilon_1, t] \setminus C$. If $Q_{x'}^{t'} = Q_{x'}^{t(q_0)}$, then $f_{t'}(x') = f_{t(q_0)}(x') \in U'$. Otherwise $Q_{x'}^{t'} \supsetneq Q_{x'}^{t(q_0)}$ (since $t(q_0) < t'$), and hence $x' \in (a,b)$. Then

$$x - \varepsilon_1 < a \le a_{x'}^{t'} < x' < x + \varepsilon_1.$$

As $a_{x'}^{t'}$ is unchanged at time t' by (3.27) and hence also at time $t(q_0) < t'$, we have $f_{t'}(x') = \sigma(a_{x'}^{t'}) = f_{t(q_0)}(a_{x'}^{t'})$ by (3.25). his last point lies in U', by the above inequality and the choice of ε_1. Thus, $F(x',t') \in U'$.

If $x < a$, then, as f_t is continuous, there is an $\varepsilon_1 < \varepsilon_0$ such that $x + \varepsilon_1 < a$ and $F(x',t) \in U$ for all $x' \in (x - \varepsilon_1, x + \varepsilon_1)$. Choose ε_1 so that $t - \varepsilon_1 \notin C$. And as noted earlier, $F(x',t') = F(x',t)$ for every $t' \in [t - \varepsilon_1, t] \setminus C$ and all these x', so $F(x',t') \in U$. On the other hand for $t' \in (t - \varepsilon_1, t) \cap C$, say $t' \in \big(s(q), t(q)\big)$

with $q \in Q_{\varepsilon_1}$, we have $f_{t'}(x') = f_{t(q)}(x')$ for these same x', because $x' \notin (a,b) \supseteq \big(a(q), b(q)\big)$ and hence $f_{t'}(x')$ remained constant throughout the homotopy defined in (3.28) or (3.29). Since $t(q) \notin C$, we are thus home by the case of $t' \in [t-\varepsilon_1, t] \setminus C$.

2. We find an ε for which $F\big((x-\varepsilon, x+\varepsilon), [t, t+\varepsilon)\big) \subseteq U$.

   For every $\varepsilon > 0$, let

   $$Q_\varepsilon := \big\{ q \in \mathcal{P} \cup \mathcal{R}' \mid [t, t+\varepsilon) \cap \big(s(q), t(q)\big) \neq \emptyset \big\}.$$

   As in the first step, we may assume that Q_ε is infinite for every ε. Thus $Q_\varepsilon \neq \emptyset$, but Q_ε has no least element in $\preceq$. Then $t \notin C$, so the sets

   $$P^+ := \{ p \in \mathcal{P} \mid s(p) \geq t \} \quad \text{and} \quad P^- := \{ p \in \mathcal{P} \mid t(p) \leq t \}$$

   partition $\mathcal{P}$. From (3.22) we know that there are pairs p with $s(p)$ arbitrarily close after t, thus P^+ meets every Q_ε. So P^+ has no least element in $\preceq$, and

   $$t = \inf\{ t(p) \mid p \in P^+ \}. \tag{3.32}$$

   Let us write $r := (P^-, P^+)$. But note that P^- may be empty, in which case $r \notin \mathcal{R}$ and every Q_ε might meet infinitely many chains M_i.

   If $F(x,t)$ is an end, then, as in Step 1, it will suffice to find an $\varepsilon < \varepsilon_0$ such that $t + \varepsilon \notin C$ and $F(x', t') \in U'$ for all $x' \in (x-\varepsilon, x+\varepsilon)$ and all $t' \in [t, t+\varepsilon] \setminus C$.

   We distinguish two cases.

   (i) Our first case is that for every ε there is a $q \in Q_\varepsilon$ with $x \in \big(a(q), b(q)\big)$. Depending on whether P^- is empty or not, we shall in two different ways define an end $z(r)$ and an interval $\big[a(r), b(r)\big]$ containing x, and in each case prove that

   $$f_t \textit{ maps } \big(a(r), b(r)\big) \textit{ to } z(r). \tag{3.33}$$

   We shall then use (3.33) to find the ε desired in Step 2.

   We first assume that $P^- \neq \emptyset$. Then r is a critical partition, so $z(r)$ is defined and is an end. Every $Q_\varepsilon \cap \mathcal{P}$ meets only finitely many chains M_i, and for the largest of these i we have $Q_{\varepsilon_0} \subseteq \hat{M}_i$ for some small enough ε_0. Then the intervals $\big[a(q), b(q)\big]$ with $q \in Q_{\varepsilon_0}$ are nested, and

   $$\big[a(r), b(r)\big] = \bigcap_{p \in P_i^+} \big[a(p), b(p)\big] = \bigcap_{q \in Q_{\varepsilon_0}} \big[a(q), b(q)\big]$$

by definition of $a(r)$ and $b(r)$, and (3.22). By our assumption for Case 2(i), we have $x \in \big[a(r), b(r)\big]$.

If r is bad, then r lies in $\mathcal{R}'$ and precedes all pairs in P^+. By (3.24) and (3.32) we have $t(r) \leq t$, so $r \in Q^t$. In fact, r is the greatest element of Q^t. For r succeeds every $p \in Q^t \cap \mathcal{P}$, since these p lie in P^-. But then r also succeeds any $r' \in Q^t \cap \mathcal{R}'$: otherwise $r \prec p \prec r'$ for some $p \in P^+$ by (3.22), while $t(p) \leq t(r') \leq t$ (by (3.24) and $r' \in Q^t$) implies that $p \in P^-$. Now as r is the greatest element of Q^t, it is also the greatest element of Q^t_y for every $y \in \big(a(r), b(r)\big)$, giving $a^t_y = a(r)$. As $t \notin C$, (3.25) yields

$$f_t(y) = \sigma(a^t_y) = \sigma\big(a(r)\big) \stackrel{(3.17)}{=} z(r),$$

completing the proof of (3.33) for the case that $P^- \neq \emptyset$ and r is bad.

Let us suppose now that r is not bad (so $r \in \mathcal{R} \setminus \mathcal{R}'$), and once more show that $f_t(y) = z(r)$ for every $y \in \big(a(r), b(r)\big)$. As before, we have $f_t(y) = \sigma(a^t_y)$ by (3.25). If y is critical, then $Q^t_y \neq \emptyset$, and hence $a^t_y = a^-_y$ by (3.22) and the definitions of r, a^-_y and a^t_y. Thus

$$f_t(y) = \sigma(a^t_y) = \sigma(a^-_y) = z_y = z(r)$$

since y is not bad. Now assume that y is not critical. By definition of r, this means that

$$Q^t_y \cap \mathcal{P} = \emptyset. \tag{3.34}$$

We first prove that y is unchanged at time t. Indeed, otherwise $Q^t_y \neq \emptyset$, and by (3.34) there exists an $\tilde{r} = (\tilde{P}^-, \tilde{P}^+) \in \mathcal{R}'$ such that $y \in \big(a(\tilde{r}), b(\tilde{r})\big)$ and $t(\tilde{r}) \leq t$. By (3.24), $\tilde{r}$ precedes all pairs in P^+, which by (3.23) implies that $\tilde{P}^+ \supseteq P^+$. As $\tilde{r} \neq r$ (since $\tilde{r} \in \mathcal{R}'$ but $r \notin \mathcal{R}'$), there exists a pair $p \in \tilde{P}^+ \cap P^-$. As all sufficiently early pairs of $\tilde{P}^+$ lie in $M_{i(\tilde{r})}$, we can find this p in $M_{i(\tilde{r})}$, giving $\tilde{r} \leq p$. But then $y \in \big(a(\tilde{r}), b(\tilde{r})\big) \subseteq \big(a(p), b(p)\big)$ and $t(p) \leq t$, contradicting (3.34).

Thus y is unchanged at time t. Then (3.25) yields $f_t(y) = \sigma(y)$, so let us show that $\sigma(y) = z(r)$. Our aim is to prove $\sigma(y) = z(r)$ using (3.8). Let $[a, b] \ni y$ be a maximal interval with the property that (a, b) (which is allowed to be empty) avoids every domain of a pass of σ. As every neighbourhood of $a(r)$ or $b(r)$ meets the domain of a pass of σ (namely, in $a(p)$ or $b(p)$ for every sufficiently small $p \in P^+_{i(r)}$) we have $[a, b] \subseteq \big[a(r), b(r)\big]$.

We first prove that every point in $[a,b]$ is unchanged at time t. Indeed, otherwise there is a $q \in Q^t$ for which $\big(a(q), b(q)\big)$ meets $[a,b]$. As y is unchanged at time t we have $\big(a(q), b(q)\big) \not\supset [a,b]$ and thus $a(q) \in (a,b)$ or $b(q) \in (a,b)$, say $a(q) \in (a,b)$. If $q \in \mathcal{P}$, then $\big[a(q), a^-(q)\big]$ is the domain of a pass of σ that meets (a,b), a contradiction. If $q = (\tilde{P}^-, \tilde{P}^+) \in \mathcal{R}'$, then since $\big[a(q), b(q)\big] = \bigcap_{p \in \tilde{P}^+_{i(q)}} \big[a(p), b(p)\big]$ there is a pair $p \in \tilde{P}^+_{i(q)}$ with $a(p) \in \big(a, a(q)\big) \subseteq (a,b)$, with a similar contradiction. Hence every point in $[a,b]$ is unchanged at time t and thus $f_t \restriction [a,b] = \sigma \restriction [a,b]$.

Let us show that $\sigma(a) = \sigma(b) = z(r)$: then either $[a,b] = \{y\}$ and thus $\sigma(y) = z(r)$, or σ maps all of $[a,b]$, including y, to $z(r)$ by (3.8). We prove $\sigma(a) = z(r)$; the proof that $\sigma(b) = z(r)$ is analogous. If $a = a(r)$, then $\sigma(a) = \sigma\big(a(r)\big) = z(r)$, by definition of $z(r)$. If $a \neq a(r)$ then, by the choice of $[a,b]$, there is a sequence $y_0 < y_1 < \cdots$ of points in $\big[a(r), b(r)\big]$ such that $\lim_n y_n = a$ and such that every y_n lies in the interior of the domain of a pass of σ (possibly the same for all n), and hence is critical (for r). As a is unchanged at time t, and hence $a \notin \big(a(q), b(q)\big)$ for every $q \in Q^t$, we have $y_n \leq b^t_{y_n} \leq a$ and hence also $\lim_n b^t_{y_n} = a$. As σ is continuous, the points $\sigma(b^t_{y_n})$ converge in H to $\sigma(a)$. Since each y_n is critical but not bad, we have $\sigma(b^t_{y_n}) = z(r)$ for every n, and thus $\sigma(a) = z(r)$. This completes the proof of (3.33) for the case of $P^- \neq \emptyset$.

We now assume that $P^- = \emptyset$. Then every $x \in [0,1]$ is unchanged at time t, since $Q^t_x \neq \emptyset$ would imply $P^- \neq \emptyset$ by (3.22) and (3.24). Thus, $f_t = \sigma$. Consider the $\leq$-chain Q^1_x of all $q \in \mathcal{P} \cup \mathcal{R}'$ with $x \in \big(a(q), b(q)\big)$. By our assumption for Case 2(i),

$$\forall \varepsilon > 0 : \quad Q^1_x \cap Q_\varepsilon \neq \emptyset. \tag{3.35}$$

Since $Q^t_x = \emptyset$ this means that Q^1_x, like Q_ε, has no least element, and by (3.22) neither does $Q^1_x \cap \mathcal{P} = Q^1_x \cap P^+$. Therefore

$$\big[a(r), b(r)\big] := \bigcap_{p \in Q^1_x \cap P^+} \big[a(p), b(p)\big] = \bigcap_{q \in Q^1_x} \big[a(q), b(q)\big].$$

Pick $p_0, p_1, \ldots \in Q^1_x \cap P^+$ with $\lim_n a(p_n) = a(r)$. As σ is continuous, $\lim_n z(p_n) = \sigma\big(a(r)\big) =: z(r)$. (Recall that $z(p_n) = \sigma\big(a(p_n)\big)$.) By (3.9), Lemma 2.9, and the fact that G is locally finite, $z(r)$ is an end.

To complete the proof of (3.33), we show that $\big(a(r), b(r)\big)$ avoids every domain of a pass of σ: then it is a maximal interval with this property, and (3.33) follows from (3.8) and the fact that $f_t = \sigma$. (For the application of (3.8) note that $\sigma\big(b(r)\big) = \sigma\big(a(r)\big)$, by taking limits in (3.9).) Suppose that $\big(a(r), b(r)\big)$ contains a point y from the domain of a pass of σ, say of the pair p. If p was an element of Q_x^1, there would be another pair $p' \leq p$ in Q_x^1, with $y \notin \big(a(p'), b(p')\big)$. As this contradicts the choice of y as a point in $\big(a(r), b(r)\big)$, we have $p \notin Q_x^1$. Thus, $x \notin \big(a(p), b(p)\big)$, and hence $p \not\geq q$ for all $q \in Q_x^1$. But p is nested with every such q, since $y \in \big(a(p), b(p)\big) \cap \big(a(q), b(q)\big)$. Therefore $p < q$, and hence

$$\forall q \in Q_x^1 : \quad t < t(p) < s(q),$$

by (3.24) and since $p \notin P^- = \emptyset$. But then for $\varepsilon < t(p) - t$ we have $Q_\varepsilon \cap Q_x^1 = \emptyset$, contradicting (3.35). This completes the proof of (3.33).

We have thus shown for both sets of definitions that f_t maps $\big(a(r), b(r)\big)$, and hence also $\big[a(r), b(r)\big]$, to the end $z(r)$. As $x \in \big[a(r), b(r)\big]$, we in particular have $z(r) = f_t(x) \in U'$. As f_t is continuous, there is a $\delta > 0$ such that f_t maps $\big(a(r) - \delta, b(r) + \delta\big)$ to U'. By the definition of $\big[a(r), b(r)\big]$ (in either case) we can find a pair $p_0 \in P^+$ such that $a(p_0) \in \big(a(r) - \delta, a(r)\big)$ and $b(p_0) \in \big(b(r), b(r) + \delta\big)$. Choose ε so that $t + \varepsilon \notin C$, and small enough that $t + \varepsilon < t(p_0)$ as well as $(x - \varepsilon, x + \varepsilon) \subseteq \big(a(p_0), b(p_0)\big)$. Then for all $t' \in [t, t + \varepsilon] \setminus C$ and $x' \in (x - \varepsilon, x + \varepsilon)$ we have $p_0 \in Q_{x'}^{t(p_0)} \supseteq Q_{x'}^{t'}$ and hence $a(p_0) = a_{x'}^{t(p_0)} \leq a_{x'}^{t'}$. Thus,

$$a(r) - \delta < a(p_0) \leq a_{x'}^{t'} \leq x' < x + \varepsilon < b(r) + \delta,$$

giving $f_t(a_{x'}^{t'}) \in U'$ by the choice of δ. But $a_{x'}^{t'}$ is unchanged at time t' by (3.27), and hence also at time $t \leq t'$. So this latter point is just $\sigma(a_{x'}^{t'})$, giving $F(x', t') = \sigma(a_{x'}^{t'}) = f_t(a_{x'}^{t'}) \in U'$ by (3.25).

(ii) Our second case is that there is an ε_0 such that $x \notin \big(a(q), b(q)\big)$ for all $q \in Q_{\varepsilon_0}$. Suppose first that there is even an $\varepsilon_1 < \varepsilon_0$ such that $(x - \varepsilon_1, x + \varepsilon_1)$ avoids $\big(a(q), b(q)\big)$ for all $q \in Q_{\varepsilon_1}$. Then consider any $x' \in (x - \varepsilon_1, x + \varepsilon_1)$. For every $t' \in [t, t + \varepsilon_1) \setminus C$ we have $F(x', t') = F(x', t)$ by (3.25), since $Q_{x'}^{t'} = Q_{x'}^t$. For $t' \in [t, t + \varepsilon_1) \cap C$, say $t' \in \big(s(q), t(q)\big)$ with $q \in Q_{\varepsilon_1}$, this implies $F(x', t') = F\big(x', s(q)\big) = F(x', t)$ by (3.28) or (3.29), since $x' \notin \big(a(q), b(q)\big)$ and $s(q) \in [t, t + \varepsilon_1) \setminus C$. As f_t is continuous, there is an $\varepsilon_2 < \varepsilon_1$ such that f_t maps $(x - \varepsilon_2, x + \varepsilon_2)$ to U. Then $F(x', t') = F(x', t) \in U$ for all $x' \in (x - \varepsilon_2, x + \varepsilon_2)$ and $t' \in [t, t + \varepsilon_2)$.

We may thus assume that there is no such ε_1. Then:

$$\text{\textit{For every } } \varepsilon\text{\textit{, the interval} } (x-\varepsilon, x+\varepsilon) \text{ \textit{meets} } \big(a(q), b(q)\big) \text{ \textit{for infinitely many} } q \in Q_\varepsilon. \tag{3.36}$$

By (3.36) there is a sequence $q_0, q_1, \dots$ of pairs and partitions in Q_{ε_0} with $\lim_n a(q_n) = x$ or $\lim_n b(q_n) = x$; we assume that $\lim_n a(q_n) = x$. As σ is continuous, we have $\lim_n \sigma(a_n) = \sigma(x)$. By (3.9) and (3.20), the sequence $\big(\sigma(a_n)\big)_{n\in\mathbb{N}}$ is a sequence of vertices and ends, and every vertex appears only finitely often (Lemma 2.9). Therefore $\sigma(x)$ is an end.

Let us show that x is unchanged at time t. By (3.36), any interval $\big(a(q), b(q)\big)$ that contains x meets some $\big(a(q_n), b(q_n)\big)$, but is not contained in it since $q_n \in Q_{\varepsilon_0}$ and hence $x \in \big(a(q), b(q)\big) \setminus \big(a(q_n), b(q_n)\big)$. Thus $q > q_n$ by (3.21), and $t(q) > t(q_n) > t$ by (3.24), giving $q \notin Q^t_x$ as desired. As $x \notin \big(a(q), b(q)\big)$ for every $q \in Q_{\varepsilon_0}$, x remains unchanged at all times $t' \in [t, t+\varepsilon_0] \setminus C$.

As x is unchanged at time t, we have $\sigma(x) = F(x,t) \in U'$. As σ is continuous, there is an $\varepsilon < \varepsilon_0$ such that $\sigma\big((x-\varepsilon, x+\varepsilon)\big) \subset U'$ and $t+\varepsilon \notin C$. Then for every $x' \in [x, x+\varepsilon)$ and $t' \in [t, t+\varepsilon] \setminus C$ we have $x \leq a^{t'}_{x'}$ $(\leq x')$, because x is unchanged at time t' and hence $x \notin \big(a(q), b(q)\big)$ for every $q \in Q^{t'}_{x'}$. Likewise, for $x' \in (x-\varepsilon, x]$ we have $x' \leq b^{t'}_{x'} \leq x$. Thus for every $x' \in (x-\varepsilon, x+\varepsilon)$ and $t' \in [t, t+\varepsilon] \setminus C$ we have $a^{t'}_{x'} \in (x-\varepsilon, x+\varepsilon)$ or $b^{t'}_{x'} \in (x-\varepsilon, x+\varepsilon)$, say $b^{t'}_{x'} \in (x-\varepsilon, x+\varepsilon)$. By (3.27), $b^{t'}_{x'}$ is unchanged at time t' and hence also at time $t \leq t'$. We thus have $F(x', t') = \sigma(a^{t'}_{x'}) = \sigma(b^{t'}_{x'}) = f_t(b^{t'}_{x'}) \in U'$ by (3.25).

This completes the fourth and final part of our proof that σ and τ are homtopic if τ does not traverse chords.

Before we consider the general case where τ too travserses chords, let us remind ourselves in which subsets of H the homotopies considered so far run. In part one of the above proof, we first used (3.5) to straighten σ to a path σ', which we then trimmed further to obtain a path σ'' satisfying (3.8). This path σ'' served as f_0 for our homotopy F, which ended with a path f_1 in T. This path was homotopic to the straightened version τ' of the original path τ (not traversing any chords). We thus found homotopies

$$\sigma \sim \sigma' \sim f_0 \sim f_1 \sim \tau' \sim \tau.$$

The first of these homotopies runs in $\operatorname{Im}\sigma$; see (3.5). In the second homotopy we retracted segments $\sigma' \restriction [a,b] \subseteq T$ to paths $\sigma'' \restriction [a,b]$ with image $\sigma'(a)T\sigma'(b)$. This is the unique $\sigma'(a)$–$\sigma'(b)$ arc in T, so Lemma 2.1 implies that the image of $\sigma' \restriction [a,b]$ contains it. The homotopy between

$\sigma' \restriction [a,b]$ and $\sigma'' \restriction [a,b]$, which (3.6) says runs in the union of the images of those two paths, thus in fact runs in $\operatorname{Im}\sigma' \subseteq \operatorname{Im}\sigma$, and hence so does the entire second homotopy $\sigma' \sim \sigma'' = f_0$. The third homotopy, $f_0 \sim f_1$, runs in $\operatorname{Im} f_0 \subseteq \operatorname{Im}\sigma' \subseteq \operatorname{Im}\sigma$ by (3.31). Similarly, the last homotopy $\tau \sim \tau'$ runs in $\operatorname{Im}\tau$, and the penultimate one, $f_1 \sim \tau'$, runs in $\operatorname{Im} f_1 \cup \operatorname{Im}\tau' \subseteq \operatorname{Im}\sigma \cup \operatorname{Im}\tau$. All in all, we have shown the following:

If τ traverses no chords and w_σ reduces to the empty word, then there is a homotopy in $\operatorname{Im}\sigma \cup \operatorname{Im}\tau$ between σ and τ. (3.37)

3.5.5 The general case

To complete our proof of Lemma 3.13, we now consider the case in which both σ and τ traverse chords. By Lemma 3.20, there are paths σ' and τ' such that $w_{\sigma'} = r(w_\sigma)$ and $w_{\tau'} = r(w_\tau)$, and we may further assume that every pass of σ' is also a pass of σ while every pass of τ' is also a pass of τ. Statement (3.37), applied to every non-trivial interval $[a,b]$ that is maximal with the property that it avoids the interior of every domain of a pass of σ', yields $\sigma \sim \sigma'$ by Lemma 3.1: note that $\sigma \restriction [a,b]$ and $\sigma' \restriction [a,b]$ have the same first and last point (because a and b are either boundary points of domains of common passes of σ and σ' or limits of such points, and σ and σ' are continuous), and the reduction of w_σ to $w_{\sigma'}$ defines a reduction of $w_{\sigma\restriction[a,b]}$ to $w_{\sigma'\restriction[a,b]} = \emptyset$. Likewise, we obtain $\tau \sim \tau'$. It thus suffices to prove $\sigma' \sim \tau'$, from our assumptions that

$$w_{\sigma'} = r(w_\sigma) = r(w_\tau) = w_{\tau'} \tag{3.38}$$

(cf. Lemma 3.12).

Our aim is to use (3.7) to obtain the desired homotopy $\sigma' \sim \tau'$. But (3.7) requires that the two paths considered have the same passes, not just the same trace. In order to make (3.7) applicable, we therefore have to 'synchronize' corresponding passes of σ' and τ': make their domains coincide by reparametrizing σ' and τ', and make σ' and τ' agree on those domains by applying local homotopies inside the corrsponding chords.

Every path $\alpha\colon [0,1] \to H$ defines a partition of $[0,1]$ into intervals: the interiors of domains of passes, and the (closed) components of the rest of $[0,1]$. The set $\mathcal{I}_\alpha$ of all those intervals, including trivial 'intervals' $[x,x] = \{x\}$, inherits a linear ordering from $[0,1]$. The bijection between the passes of σ' and τ' provided by (3.38) defines an order-preserving bijection $\pi\colon \mathcal{I}_{\sigma'} \to \mathcal{I}_{\tau'}$.

Although π maps open to open and closed to closed intervals, it might map non-trivial closed intervals to trivial ones or vice versa. In order to synchronize σ' with τ' as planned, we therefore have to expand trivial closed intervals to non-trivial ones in our reparametrizations of σ' and τ'. This will be possible, since clearly $\mathcal{I}_{\sigma'}$ and $\mathcal{I}_{\tau'}$ contain only countably many trivial intervals whose corresponding interval in the other set is non-trivial.

We may thus partition $[0,1]$ into a set $\mathcal{I}$ of intervals so that there exist order-preserving bijections $\pi_{\sigma'}\colon \mathcal{I} \to \mathcal{I}_{\sigma'}$ and $\pi_{\tau'}\colon \mathcal{I} \to \mathcal{I}_{\tau'}$ that map open to open intervals bijectively, and trivial to trivial intervals, and which commute with $\pi\colon \mathcal{I}_{\sigma'} \to \mathcal{I}_{\tau'}$. We can now define surjective maps $\varphi, \psi\colon [0,1] \to [0,1]$ such that φ maps every $I \in \mathcal{I}$ onto $\pi_{\sigma'}(I) \in \mathcal{I}_{\sigma'}$ and ψ maps every $I \in \mathcal{I}$ onto $\pi_{\tau'}(I) \in \mathcal{I}_{\tau'}$.

Clearly, $\sigma'' := \sigma' \circ \varphi$ is homotopic to σ', and $\tau'' := \tau' \circ \psi$ is homotopic to τ'. So it suffices to show that $\sigma'' \sim \tau''$. But these maps now have not only the same trace but also the same domains of corresponding passes. Combining homotopies between corresponding passes inside their respective chords with a homotopy between the rests of σ'' and τ'' as in (3.7) yields the desired homotopy $\sigma'' \sim \tau''$, by Lemma 3.1.

This completes the proof of Lemma 3.13 and with it our proof that $\pi_1(|G|)$ embeds canonically into F_∞.

Chapter 4

$\mathcal{C}(G)$ and singular homology

4.1 Introduction

The topological cycle space $\mathcal{C}$ had not been considered in graph theory before [16] appeared, and it has been surprisingly successful at extending the classical cycle space theory of finite graphs to locally finite graphs; see [2, 3, 4, 5, 7, 8, 9, 13, 18, 30, 31, 32, 40], or [15] for a survey. However, a question raised in [13] but still unanswered is how new, from a topological viewpoint, is the homology described implicitly by $\mathcal{C}$. In this chapter we shall take the first step to clarify this relationship.

Since topological circles are (images of simplices representing) singular 1-cycles in $|G|$, it is also natural to ask how closely $\mathcal{C}(G)$, respectively its oriented version $\vec{\mathcal{C}}(G)$, is related to the first singular homology group of $|G|$. Indeed it is not clear whether the two coincide by some natural canonical isomorphism, so that $\vec{\mathcal{C}}(G)$ would be just another way of looking at $H_1(|G|)$.

Our aim in this chapter is to answer this question. We begin by studying the homomorphism $f\colon H_1(|G|) \to \vec{\mathcal{C}}(G)$ that should serve as the desired canonical isomorphism if indeed there is one. Surprisingly, this homomorphism is easily seen to be surjective. However, it turns out that it usually has a non-trivial kernel. Thus $\vec{\mathcal{C}}(G)$, despite looking 'larger' because we allow infinite sums in its generation from elementary cycles, turns out to be a (usually proper) quotient of $H_1(|G|)$.

For the proof that f has a non-trivial kernel we will need our characterization of the fundamental group of $|G|$ from Chapter 3. The embedding $\pi_1(|G|) \to F_\infty$ will enable us to show that some elements of the kernel of f are not boundaries of singular 2-chains, completing the proof that f need not be injective: By counting occurences of certain subwords of a reduced word we shall be able to define an invariant on 1-chains that vanishes on boundaries but is non-zero on some elements of the kernel of f.

4.2 Mapping $H_1(|G|)$ to $\mathcal{C}(G)$

Let $G = (V, E)$ be a connected locally finite graph. In this section we shall start our comparison of the first singular homology group H_1 of $|G|$ (with integer coefficients) with the oriented topological cycle space $\vec{\mathcal{C}}(G)$ of G. When G is finite, then $|G| = G$ and all circuits and their thin sums are finite. Hence in this case $\vec{\mathcal{C}}$ is just the first simplicial homology group of G, so the two groups are indeed the same.

When G is infinite, however, both circuits and thin sums can be infinite too. So they are not just the simplicial 1-cycles in G. But there is an obvious singular 1-cycle in $|G|$ associated with an oriented circuit φ_α: the circle path α, viewed as a singleton 1-chain. Our aim is to extend this correspondence to one between H_1 and $\vec{\mathcal{C}}$.

Our approach will be to define a homomorphism $f\colon H_1 \to \vec{\mathcal{E}}(G)$ that counts for a given homology class h how often the 1-simplices of a cycle representing h, when properly concatenated, traverse a given edge $\vec{e}$; we then let $f(h) \in \vec{\mathcal{E}}(G)$ map $\vec{e}$ to this number.[1] We shall prove that $f(h)$ always lies in $\vec{\mathcal{C}}$ and, perhaps surprisingly, that f maps H_1 onto $\vec{\mathcal{C}}$. However, we find that f is not normally injective. The main result of this chapter characterizes the graphs for which it is:

Theorem 4.1. *The map $f\colon H_1(|G|) \to \vec{\mathcal{E}}(G)$ is a group homomorphism onto $\vec{\mathcal{C}}(G)$, which has a non-trivial kernel if and only if G contains infinitely many (finite) circuits.*

Thus, $\vec{\mathcal{C}}$ turns out to be a canonical—but usually non-trivial—quotient of H_1. Taking this result mod 2 answers our original question: the topological cycle space $\mathcal{C}$ of G is a canonical—but usually non-trivial—quotient of the singular homology group of $|G|$ with $\mathbb{F}_2$ coefficients.

We remark that the last condition in Theorem 4.1 can be rephrased in various natural ways: that G has a non-trivial end; that G has a spanning tree with infinitely many chords; that every spanning tree of G has infinitely many chords; or that G contains infinitely many *disjoint* (finite) circuits [14, Ex. 37, Ch. 8]. The remainder of this chapter is devoted to the proof of Theorem 4.1.

Let us define f formally. Recall that we denote by S^1 the unit circle in the complex plane. The elements of $H_1(S^1)$ are represented by the loops $\eta_k\colon [0,1] \to S^1$, $t \mapsto e^{2\pi i k t}$, $k \in \mathbb{Z}$. Write $\pi\colon H_1(S^1) \to \mathbb{Z}$ for the group isomorphism $[\eta_k] \mapsto k$. For every edge e of G, let $f_e\colon |G| \to S^1$ wrap e round S^1 in its natural direction, defining $f_e \restriction \mathring{e}$ as $\eta_1 \circ \theta_e^{-1}$ (recall that θ_e is the isomorphism $[0,1] \to e$ fixed in the definition of G as a 1-complex) and putting $f_e(|G| \setminus \mathring{e}) := 1 \in \mathbb{C}$. Note that f_e is continuous.

[1]The precise definition of f will be given shortly.

The following lemma is easy to prove using homotopies in S^1, combined by Lemma 3.1:

Lemma 4.2. *Let* $\alpha\colon [0,1] \to |G|$ *be a loop based at a vertex. If* α *traverses* e *exactly* k *times in its natural direction and exactly* ℓ *times in the opposite direction, then* $\pi([f_e \circ \alpha]) = k - \ell$.

Proof. Composing a pass of α through e (in its natural direction) with f_e yields a map from a subinterval of $[0,1]$ to S^1 which, after reparametrization, is homotopic to η_1.

The domains of distinct passes of α through e are closed subintervals of $[0,1]$ meeting at most in their boundary points. The rest of $[0,1]$ is a finite disjoint union of open intervals (s,t) (or $(s,1]$ or $[0,t)$). Each of these is in turn a disjoint union, possibly infinite, of open intervals (s',t') which α maps to $\mathring{e}$ and closed intervals which $f_e \circ \alpha$ maps to $1 \in \mathbb{C}$. Since $[s,t]$, by definition, contains no pass through e, α always maps s' and t' to the same endvertex of e. Then $\alpha \restriction [s',t']$ is homotopic to the constant map to that vertex, and $(f_e \circ \alpha) \restriction [s',t']$ is homotopic to the constant map to 1. These homotopies combine to a homotopy of $(f_e \circ \alpha) \restriction [s,t]$ to the constant map with value 1.

We deduce that $f_e \circ \alpha$ is homotopic to a concatenation $\sigma_1 \cdot \ldots \cdot \sigma_n$ of loops in S^1 of which (after reparametrization) k are equal to η_1 and ℓ are equal to the inverse loop η_1^-, and the rest are constant with value 1. The result follows. □

Given $h \in H_1(|G|)$, we now let $f(h) \in \vec{\mathcal{E}}(G)$ assign $(\pi \circ (f_e)_*)(h) \in \mathbb{Z}$ to the natural orientation $\vec{e}$ of e:

$$f(h)\colon \vec{e} \mapsto (\pi \circ (f_e)_*)(h) \in \mathbb{Z}.$$

This completes the definition of $f\colon H_1(|G|) \to \vec{\mathcal{E}}$, which is clearly a group homomorphism.

We want f to be a homomorphism between $H_1(|G|)$ and $\vec{\mathcal{C}}(G)$, so it has to map homology classes to elements of the cycle space. This is indeed the case:

Lemma 4.3. $\operatorname{Im} f \subseteq \vec{\mathcal{C}}(G)$.

Proof. By Theorem 2.10 it suffices to show that for every finite oriented cut $\vec{F}$ of G and every $h \in H_1(|G|)$ we have $\sum_{\vec{e} \in \vec{F}} f(h)(\vec{e}) = 0$. Let $\vec{F} = \vec{E}(U,U')$ and h be given, let $F = \{e \mid \vec{e} \in \vec{F}\}$, and assume for simplicity that the orientations $\vec{e} \in \vec{F}$ of these edges are their natural orientations. Since f is a homomorphism, we may assume that h is represented by an elementary 1-cycle, which we may choose by Corollary 2.16 to consist of a loop α based at a vertex. We shall prove that α traverses the edges in F as often from U to U' as it does from U' to U. Then

$$\sum_{\vec{e} \in \vec{F}} f(h)(\vec{e}\,) = \sum_{e \in F} (\pi \circ (f_e)_*)([\alpha]) = \sum_{e \in F} \pi([f_e \circ \alpha]) = 0$$

by Lemma 4.2.

Let $[s_1, t_1], \ldots, [s_n, t_n]$ be the domains of the passes of α through edges of F, and let further $t_0 := 0$ and $s_{n+1} := 1$. In order to prove that as many of these passes are from a vertex in U to one in U' as vice versa, it suffices to show that each of the segments $\beta = \alpha \restriction [t_i, s_{i+1}],\ 0 \le i \le n$, has either all its vertices in U or all its vertices in U'. If $t_i = s_{i+1}$, then we are done, so assume that $t_i < s_{i+1}$, hence β is a path. If the starting vertex $\beta(t_i)$ of β lies in U, say, put

$$s := \sup\{\, r \in [t_i, s_{i+1}] : V \cap \beta([t_i, r]) \subseteq U \,\}.$$

We wish to show that $s = s_{i+1}$. If not, then $\beta(s)$ is an end by Lemma 2.5, and this end lies both in the closure of U and in the closure of U'. But these closures are disjoint: the set S of vertices incident with an edge in F is finite, and since S separates U from U', the neighbourhood $\hat{C}(S, \omega)$ of any end ω avoids either U or U'. $\square$

Next, we prove that f is surjective. At first glance, this may seem surprising: after all, we have to capture arbitrary thin sums of oriented circuits, which may well be disjoint, by finite 1-cycles.

Lemma 4.4. $\operatorname{Im} f \supseteq \vec{\mathcal{C}}(G)$.

Proof. Let $\varphi = \sum_{\alpha \in A} \varphi_\alpha \in \vec{\mathcal{C}}(G)$ be an arbitrary thin sum of oriented circuits, where each α is a circle path in $|G|$. We may assume that each α is based at a vertex $v(\alpha)$. (If the image of a non-trivial path contains no vertex it must lie inside an edge, because non-trivial sets of ends cannot be connected. Since edges do not contain circles, α has to meet a vertex.) We shall construct a loop τ in $|G|$ such that $f([\tau]) = \varphi$.

Let T be a spanning tree of G and pick a root $r \in V$. Write V_n for the set of vertices at distance n in T from r, and let T_n be the subtree of T induced by $V_0 \cup \cdots \cup V_n$. Our first aim will be to construct a loop σ in $|G|$ that traverses every edge of T once in each direction and avoids all other edges of G, similar to Example 3.8. We shall obtain σ as a limit of similar loops σ_n in $T_n \subseteq |G|$. We shall then incorporate our loops $\alpha \in A$ into σ, to obtain τ. When we describe these maps informally, we shall think of $[0, 1]$ as measuring time, and of a loop as a journey through $|G|$.

Let σ_0 be the unique (constant) map $[0, 1] \to T_0$. Assume inductively that $\sigma_n \colon [0, 1] \to T_n$ is a loop traversing every edge of T_n exactly once in each direction. Assume further that σ_n pauses every time it visits a vertex, remaining stationary at that vertex for some time. More precisely, we assume for every vertex $v \in T_n - r$ that $\sigma_n^{-1}(v)$ is a disjoint union of as many non-trivial closed intervals as v has incident edges in T_n, and of one more such interval in the case of $v = r$. Let us call the restriction of σ_n to such an interval a *pass* of σ_n through v. We are thus assuming that σ_n is the union of its passes through the vertices and edges of T_n.

Let σ_{n+1} be obtained from σ_n by replacing, for each leaf v of T_n, the unique pass of σ_n through v by a path that starts out remaining stationary at v for some time, then visits all the neighbours of v in V_{n+1} in turn, pausing at each and shuttling back and forth between v and those neighbours, and finally returns to v to pause there. Outside the passes of σ_n through leaves of T_n, let σ_{n+1} agree with σ_n. Note that σ_{n+1} satisfies our inductive assumptions for $n+1$: it traverses every edge of T_{n+1} exactly once each way, pauses every time it visits a vertex, and is the union of its passes through the vertices and edges of T_{n+1}.

Let us now define σ. Let $s \in [0,1]$ be given. If the values $\sigma_n(s)$ coincide for all large enough n, let $\sigma(s) := \sigma_n(s)$ for these n. If not, then $s_n := \sigma_n(s) \in V_n$ for every n, and $s_0 s_1 s_2 \dots$ is a ray in T; let σ map s to the end of G containing that ray.

Clearly every σ_n is continuous, and σ is continuous at points not mapped to ends. To show that σ is continuous at every point s mapped to an end $\omega = \sigma(s)$, let a neighbourhood $\hat{C}(S,\omega)$ of ω in $|G|$ be given. Choose n large enough that the finite set S is contained in $V(T_{n-1})$; let T' be the component of $T \setminus T_{n-1}$ containing s_n. We claim that σ maps the interval $I := \sigma_n^{-1}(s_n)$ to $\hat{C}(S,\omega)$. Since σ_{n+1} agrees with σ_n on the boundary points of I but not on s, we know that I is a neighbourhood of s in $[0,1]$, so this will complete the proof that σ is continuous. Let $t \in I$ be given. Induction on m shows that $\sigma_m(I) \subseteq T'$ for every $m \geq n$. Hence if $\sigma(t)$ is not an end, then $\sigma(t) = \sigma_m(t) \in T' \subseteq \hat{C}(S,\omega)$ for some m. But if $\sigma(t)$ is an end, then this is the end ω' of a ray that starts at $\sigma_n(t) = s_n$ and lies in $T' \subseteq \hat{C}(S,\omega)$. Hence so does $\omega' = \sigma(t)$.

Let τ be obtained from σ by replacing, for every vertex v, one of the passes of σ through v with a concatenation of all the circle paths α with $\alpha \in A$ and $v(\alpha) = v$. Note that these are finitely many for each v, because G has only finitely many edges at v and $\sum_{\alpha \in A} \varphi_\alpha$ is a thin sum.

Let us prove that τ is continuous. As before, this is clear at points $x \in [0,1]$ which σ does not map to an end: for such x the map τ agrees on suitable intervals $[s,x]$ and $[x,t]$ with σ or some $\alpha \in A$, which we know to be continuous. The proof that τ is continuous at points x which σ maps to ends is similar to our earlier continuity proof for σ. The only difference now is that we have to choose n large enough also to ensure that none of the α with $v(\alpha) \in T'$ passes through a vertex of S. Such a choice of n is possible, because only finitly many edges are incident with vertices in S and the φ_α form a thin family of circuits. Then $\hat{C}(S,\omega)$ contains not only T' but also the images of all α with $v(\alpha) \in T'$, because $\operatorname{Im}\alpha$ is connected but does not meet the boundary of $\hat{C}(S,\omega)$, since the boundary is a subset of S.

Finally, recall that σ traverses every edge of T once in each direction, and that it does not traverse any other edges. Therefore $f([\sigma]) = 0 \in \vec{\mathcal{C}}(G)$, and hence $f([\tau]) = \sum_{\alpha \in A} \varphi_\alpha = \varphi$ as desired. □

In fact, we have just shown the following stronger statement:

Lemma 4.5. *For every $c \in \vec{\mathcal{C}}(G)$ there is a loop α in $|G|$ with $f([\alpha]) = c$.*

4.3 Distinguishing boundaries from other cycles

To complete the proof of Theorem 4.1 it remains to show that f has a non-trivial kernel if and only if G contains infinitely many circuits. The forward implication of this is easy. Indeed, suppose that G contains only finitely many circuits, and let T be a normal spanning tree of G. Then T has only finitely many chords, so $|G|$ is homotopy equivalent to a finite graph (Corollary 3.6). Hence, as is well known, $H_1(|G|)$ equals the first simplicial homology group of G viewed as a 1-complex, which in turn is clearly isomorphic to $\vec{\mathcal{C}}(G)$. Therefore f must be injective.

The converse implication, surprisingly, is quite a bit harder. Assuming that G contains infinitely many circuits, we shall define a loop ρ in $|G|$ that traverses every edge equally often in both directions (so that $f([\rho]) = 0$), and which is easily seen not to be null-homotopic. To prove that $[\rho] \neq 0$, however, i.e. that ρ is not a boundary, will be harder: it turned out that we need the characterization of the fundamental group of $|G|$ developed in Chapter 3. With this tool we shall be able to define an invariant of 1-chains that can distinguish ρ from boundaries.

Let T be a topological spanning tree of G. Each of the infinitely many circuits in G is a finite sum (mod 2) of distinct fundamental circuits of T by Theorem 2.11. Therefore T has infinitely many chords, $e_0, e_1, \dots$ say. Since $|G|$ is compact, there is a sequence $\vec{e}_{i_0}, \vec{e}_{i_1}, \dots$ of chords which converge to an end ω of G. There exists a loop ρ in $|G|$, based at a vertex, that traverses $\vec{e}_{i_0}, \vec{e}_{i_1}, \dots, \overleftarrow{e}_{i_0}, \overleftarrow{e}_{i_1}, \dots$ in this order and runs otherwise along $\overline{T}$.[2]

[2]Thus, ρ starts with passes through $\vec{e}_{i_0}, \vec{e}_{i_1}, \dots$, interspersed with finite segments of T between the endpoints of these passes, until it reaches ω, from where it returns along $\overline{T}$ to the starting vertex of $\overleftarrow{e}_{i_0}$; it then traverses $\overleftarrow{e}_{i_0}, \overleftarrow{e}_{i_1}, \dots$ interspersed with connecting segments of T to reach ω a second time, and finally returns from there along $\overline{T}$ to its starting vertex. Note that the convergence of $e_{i_0}, e_{i_1}, \dots$ is essential for ρ to be a path:

(See Figure 4.1 for an example of ρ.) Since ρ traverses the chords of T equally often in both directions, Theorem 2.10 (applied to the fundamental cuts of T, the cuts that contain precisely one edge from T) and Lemmas 4.2 and 4.3 imply that ρ also traverses the edges of T equally often in both directions.

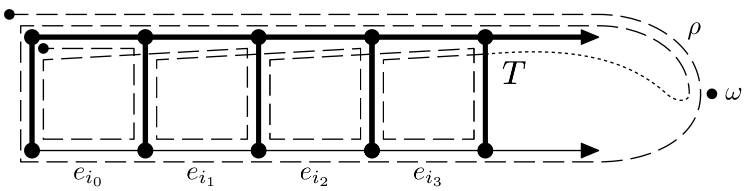

Figure 4.1: The loop ρ in the ladder with topological spanning tree T. The natural directions of the chords of T are pointing to the left side.

Hence $f([\rho]) = 0 \in \vec{\mathcal{C}}(G)$. To complete the proof that f is not injective, and thereby the proof of Theorem 4.1, we show that $[\rho] \neq 0$, i.e. that ρ is not the boundary of any 2-chain. In order to do so, we shall use our results and terminology from Chapter 3 to define an invariant of 1-chains that can distinguish ρ from boundaries. As in Chapter 3 we consider only paths whose boundary points are vertices or ends, so our invariant will be defined only for chains of 1-simplices with this property. However, it is easy to see that this entails no loss of generality: Indeed, if $\rho = \sum \lambda_n \partial \tau_n$ for 2-simplices τ_n, we can modify each τ_n into another 2-simplex τ'_n whose 0-faces are vertices or ends, and such that $\rho = \sum \lambda_n \partial \tau_n$, as follows. For every inner point x of an edge $e_x = u_x v_x$ in $|G|$ pick a fixed path π_x from x to v_x (say). Then append to every 1-simplex σ occurring in the boundary of a τ_n and ending in such a point x the path π_x after x, turning σ into a path σ' between two vertices by appending at most two such paths π_x. Then if $\partial \tau_n = \sigma_0 - \sigma_1 + \sigma_2$, say, it is easy to see that also $\sigma'_0 - \sigma'_1 + \sigma'_2$ is the boundary of a 2-simplex τ'_n. And clearly $\rho = \sum \lambda_n \partial \tau_n$ implies that also $\rho = \sum \lambda_n \partial \tau'_n$, since we modified only 1-simplices that cancelled out anyway in this sum.

We need some more notation. Given $k \in \mathbb{N}$ and a reduced word $w\colon S \to A$ (where $A = \{\vec{e}_0, \overleftarrow{e}_0, \vec{e}_1, \overleftarrow{e}_1, \dots\}$ as in Chapter 3), write $n^+(w,k)$ for the number of intervals in S (recall that an interval in S is a subset closed under betweenness) which can be written as $\{s_0, s_1, \dots\}$ with $s_0 < s_1 < \cdots$ and $w(s_j) = \vec{e}_{i_{k+j}}$ for every $j \in \mathbb{N}$. This number exists: there are at most $|w^{-1}(\vec{e}_{i_k})|$ such intervals, and this number is finite by our definition of 'word'. Put

$$n(w,k) := n^+(w,k) - n^+(w^-,k) \in \mathbb{Z}.$$

(Recall that w^- is w backwards with inverse letters, so $n^+(w^-,k)$ counts

there is no path in $|G|$ through an ω-sequence of chords that does not converge.

the intervals in S which can be written as $\{s_0, s_1, \dots\}$ with $s_0 > s_1 > \cdots$ and $w_\sigma(s_j) = \overleftarrow{e}_{i_{k+j}}$ for every $j \in \mathbb{N}$.)

Given $k \in \mathbb{N}$ and a path σ in $|G|$, let $N(\sigma, k) := n(r(w_\sigma), k)$. Given a 1-chain $\varphi = \sum_n \lambda_n \sigma_n$, let $N(\varphi, k) := \sum_n \lambda_n N(\sigma_n, k)$ for every fixed k, and put

$$N(\varphi) := \min_k |N(\varphi, k)|.$$

Unlike $N(\varphi, k)$ for fixed k, the function N is not a homomorphism. Nevertheless, it will help us distinguish our special path ρ from boundaries: we shall prove that N vanishes on boundaries, while clearly $N(\rho) = 1$.[3]

We begin by noting a property of the function $n(w, k)$:

If w is a reduced word and $w = w_1 w_2$, then there exists a $k \in \mathbb{N}$ such that $n(w, \ell) = n(w_1, \ell) + n(w_2, \ell)$ for all $\ell \geq k$. (4.1)

Indeed, denote the domains of w_1 and w_2 by S_1 and S_2 (chosen disjoint); then the domain of w is the disjoint union S of S_1 and S_2, with S_1 preceding S_2. If S_1 has a largest element, s_1 say, choose k large enough that $w(s_1) \notin \{\vec{e}_{i_k}, \overleftarrow{e}_{i_k}, \vec{e}_{i_{k+1}}, \overleftarrow{e}_{i_{k+1}}, \dots\}$. (Note that $w(s_1)$ does not have to be $\vec{e}_{i_j}$ or $\overleftarrow{e}_{i_j}$ for any j, since the e_{i_j} are not necessarily all edges of G; in this case any k would suffice.) Then for every $\ell \geq k$ none of the intervals in S counted by $n(w, \ell)$ meets both S_1 and S_2, since these intervals cannot contain s_1. Hence every such interval is either an interval of S_1 or one of S_2, so $n(w, \ell) = n(w_1, \ell) + n(w_2, \ell)$ as desired. On the other hand if S_1 has no largest element, then no interval in S that meets both S_1 and S_2 can be written as $\{s_0, s_1, \dots\}$ with $s_0 < s_1 < \dots$ or $s_0 > s_1 > \dots$, so none of the intervals counted by $n(w, \ell)$ for any ℓ meets both S_1 and S_2. Hence, in this case, $n(w, \ell) = n(w_1, \ell) + n(w_2, \ell)$ for all ℓ.

For our proof that N vanishes on boundaries φ, it suffices to show that every 2-simplex τ satisfies $N(\partial\tau, k) = 0$ for large enough k: then $N(\varphi, k) = 0$ for some (large) k, and hence $N(\varphi) = 0$ as claimed. So consider a 2-simplex τ, with boundary $\partial\tau = \sigma_2 - \sigma_1 + \sigma_0$ denoted so that σ_2 ends at the starting vertex of σ_0. Write $w_i := r(w_{\sigma_i})$ for the words to which the traces of the σ_i reduce ($i = 0, 1, 2$), and $w_{20} := r(w_{\sigma_{20}})$, where $\sigma_{20} := \sigma_2 \sigma_0$ is the path consisting of σ_2 followed by σ_0.

Note that $w_{20} = r(w_2 w_0)$. Indeed, we can reduce $w_{\sigma_{20}}$ by first applying to $w_{\sigma_2} \subseteq w_{\sigma_{20}}$ the reduction that turns w_{σ_2} into w_2, and then apply to $w_{\sigma_0} \subseteq w_{\sigma_{20}}$ the reduction that turns w_{σ_0} into w_0. Together this is a reduction of $w_{\sigma_{20}}$ to $w_2 w_0$. Let R be a reduction of $w_2 w_0$ to $r(w_2 w_0)$. Since we started with $w_{\sigma_{20}}$, the reduced word $r(w_2 w_0)$ we end up with is $r(w_{\sigma_{20}}) = w_{20}$ by Lemma 3.12.

Let us look at what R does. Since w_2 and w_0 are both reduced, every pair of positions in R has one position in w_2 and the other in w_0. Hence

[3] Indeed, the word w_ρ is easily seen to be reduced (cf. Lemma 3.10); hence $N(\rho, k) = n(w_\rho, k) = 1$ for all k, since $n^+(w_\rho, k) = 1$ and $n^+(w_\rho^-, k) = 0$.

if w denotes the subword of w_2 whose positions are deleted by R, we have found reduced words w, w_2', w_0' such that

$$w_2 = w_2' w \quad \text{and} \quad w_0 = w^- w_0' \quad \text{and} \quad w_{20} = w_2' w_0'.$$

By (4.1), therefore, we have for all large enough k

$$\begin{aligned} n(w_2, k) &= n(w_2', k) + n(w, k) \\ n(w_0, k) &= n(w^-, k) + n(w_0', k) \\ n(w_{20}, k) &= n(w_2', k) + n(w_0', k). \end{aligned}$$

As $n(w^-, k) = -n(w, k)$, we deduce that

$$n(w_2, k) + n(w_0, k) - n(w_{20}, k) = 0$$

for all these k.

Since σ_{20} is homotopic to σ_1 (across τ), Lemma 3.13 implies that $w_{20} = w_1$. We therefore deduce that

$$N(\partial\tau, k) = N(\sigma_2, k) + N(\sigma_0, k) - N(\sigma_1, k) = 0$$

for all large enough k, as desired. This completes the proof of Theorem 4.1.

Chapter 5

Interlude: Čech homology and homotopy

This chapter is devoted to the relationship between the topological cycle space of a graph with ends and its Čech homology. We shall see that their groups are canonically isomorphic, but also that this isomorphism is not enough to capture the relevance of $\vec{\mathcal{C}}(G)$ to the structure of G—the reason why cycle spaces are studied in the first place.

5.1 Čech homology

The Čech homology of a space is an alternative to singular homology for spaces that are not simplicial complexes. Consider a space X and an open cover $\mathcal{U}$ of X. Then $\mathcal{U}$ defines a simplicial complex $X_{\mathcal{U}}$, the *nerve* of $\mathcal{U}$: The 0-simplices of $X_{\mathcal{U}}$ are the elements of $\mathcal{U}$, and any $n+1$ elements of $\mathcal{U}$ form an n-simplex if and only if they have a nonempty overall intersection. For two open covers $\mathcal{U}$, $\mathcal{U}'$ of X, we write $\mathcal{U} \leq \mathcal{U}'$ if $\mathcal{U}'$ is a refinement of $\mathcal{U}$. In this case, it is easy to define a continuous map from $X_{\mathcal{U}'}$ to $X_{\mathcal{U}}$: For each 0-simplex U of $X_{\mathcal{U}'}$ (i.e. $U \in \mathcal{U}'$) there is a 0-simplex $\pi(U)$ of $X_{\mathcal{U}}$ (an element of $\mathcal{U}$) that contains it. Map each U to $\pi(U)$ and extend this map linearly to the higher-dimensional simplices in $X_{\mathcal{U}'}$ so as to obtain a map $\rho\colon X_{\mathcal{U}'} \to X_{\mathcal{U}}$.

Since $U \in \mathcal{U}'$ can be contained in more than one element of $\mathcal{U}$, the choice of $\pi\colon \mathcal{U}' \to \mathcal{U}$ is not unique and neither is ρ. But it is easy to see that all possible choices of π induce homotopic maps ρ: If π' is another possible choice of π, then their induced maps ρ and ρ' are homotopic by the homotopy $F\colon X_{\mathcal{U}'} \times [0,1] \to X_{\mathcal{U}}$ defined by letting $F(U,0) = \pi(U)$ and $F(U,1) = \pi'(U)$ for each $U \in \mathcal{U}'$ and extending linearly to the higher-dimensional simplices in $X_{\mathcal{U}'}$ and to all times $t \in (0,1)$. Thus, all choices of π induce a unique homomorphism $\rho^{\mathcal{U}}_{\mathcal{U}'}\colon H_n(X_{\mathcal{U}'}) \to H_n(X_{\mathcal{U}})$ on homology. Therefore, the homology groups $H_n(X_{\mathcal{U}})$ for all open covers $\mathcal{U}$, with the order $\leq$ defined above, together with the homomorphisms $\rho^{\mathcal{U}}_{\mathcal{U}'}$ form an inverse family. Define the nth *Čech homology group* $\check{H}_n(X)$ to be the inverse limit of this family. For more on Čech homology, see eg. [26] or [38].

The main result of this chapter is that for locally finite graphs the first Čech homology group and the topological cycle space coincide:

Theorem 5.1. *For a locally finite graph G we have a canonical isomorphism $\check{H}_1(|G|) \simeq \vec{\mathcal{C}}(G)$.*

Proof. To compute the inverse limit of the groups $H_1(X_{\mathcal{U}})$ it suffices to to consider a family $\mathfrak{U}$ of open covers of $|G|$ that contains a refinement for every open cover of $|G|$, and to compute the inverse limit of the inverse family $\big(H_1(|G|_{\mathcal{U}})\big)_{\mathcal{U}\in\mathfrak{U}}$. We will now construct a suitable $\mathfrak{U}$.

Let T be a normal spanning tree of G and denote the subtree induced by the first n levels by T_n. Now for each n let $\mathfrak{U}$ contain an open cover $\mathcal{U}_n$ consisting of the following sets: An open star of radius 2^{-n} around each vertex $v \in V(T_n)$, finitely many open subintervals of length 2^{-n} of each edge $e \in E(T_n)$, and the sets $\hat{C}(V(T_n), \omega)$ for each end ω of G. Note that $\mathcal{U}_n$ is a finite family as $G - V(T_n)$ has only finitely many components.

It is not hard to see that for each open cover $\mathcal{U}$ of $|G|$ some $\mathcal{U}_n$ is a refinement of $\mathcal{U}$: Since $|G|$ is compact, there is a finite subcover $\mathcal{U}'$ of $\mathcal{U}$. Choose n large enough so that each set in $\mathcal{U}_n$ has a diameter smaller than the Lebesgue number of $\mathcal{U}'$, where the metric of $|G|$ is that defined in Section 2.2. (Recall that in this metric, any set $\hat{C}(V(T_n), \omega)$ has diameter at most $2^{-(n-1)}$.) Then $\mathcal{U}_n$ is a refinement of $\mathcal{U}$.

Now every nerve $|G|_{\mathcal{U}_n}$ retracts to the graph G_n obtained from G by contracting all components of $G - T_n$, and hence the homology group $H_1(|G|_{\mathcal{U}_n}) = H_1(G_n)$ is a direct product of $\mathbb{Z}$'s, one for each chord of T with at least one endvertex in T_n. Thus $\check{H}_1(|G|)$ also is the direct product of copies of $\mathbb{Z}$, one for each chord of T. As the same is true for $\vec{\mathcal{C}}(G)$, we have that $\check{H}_1(|G|)$ and $\vec{\mathcal{C}}(G)$ are canonically isomorphic. $\square$

Theorem 5.1 shows that one can describe the topological cycle space in terms of the Čech homology. However, although $\check{H}_1(|G|)$ is isomorphic to $\vec{\mathcal{C}}(G)$ as a group, it does *not* sufficiently reflect the combinatorial properties of $\vec{\mathcal{C}}(G)$, i.e. its relation to the combinatorial structure of G. To make this precise, note that a number of classical results about the cycle space say which circuits generate it—as do the non-separating chordless circuits in a 3-connected graph, say. In the Čech homology, however, it is not possible to decide whether a given homology class in $\check{H}_1(|G|)$ corresponds to a circuit in G. Indeed, the obvious relation between $\check{H}_1(|G|)$ and the combinatorial structure of G is that every homology class $c \in \check{H}_1(|G|)$ corresponds to a family (c_n) of homology classes in the groups $H_1(|G|_{\mathcal{U}_n}) = H_1(G_n)$. One might think that the class c should correspond to a circuit if and only if every c_n with sufficiently large n corresponds to a circuit in G_n. But this is not the case: the limit of a sequence of cycle space elements in the G_n can be a circuit even if the elements of the sequence are not circuits in the G_n.

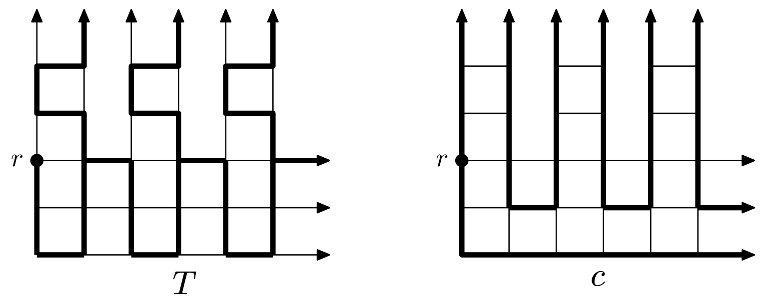

Figure 5.1: The graph G (drawn twice) with a normal spanning tree T and a circuit c.

Let G be the graph shown in Figure 5.1. G consists of a "wide ladder" with three stiles $x_1^1, x_2^1, \ldots,$ $x_1^2, x_2^2, \ldots,$ and $x_1^3, x_2^3, \ldots,$ and has attached infinitely many (oridinary) ladders by identifying the first rung of the nth ladder L_n with the edge $x_{2n-1}^1 x_{2n}^1$. It is not hard to prove that T from Figure 5.1 is a normal spanning tree of G with root $r = x_1^1$.

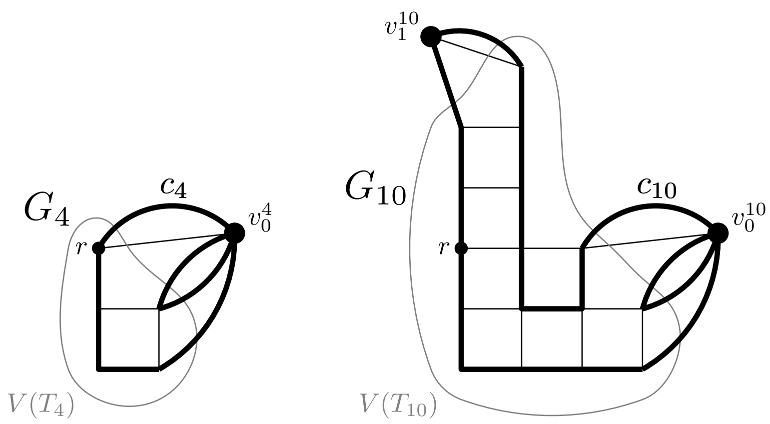

Figure 5.2: The edge sets c_4 in G_4 and c_{10} in G_{10}.

The edge set c from Figure 5.1 is a circuit, but each edge set c_n it induces on a contracted graph G_n with $n = 6k + 4$ is not a circuit. Indeed, each G_{6k+4} consists of $G[V(T_{6k+4})]$, for each $i \leq k$ a vertex v_i^{6k+4} corresponding to a contracted tail of the ladder L_i, and a vertex v_0^{6k+4} corresponding to the contracted tail of the wide ladder and all ladders L_j with $j > k$. The edge set c_{6k+4} is not a circuit since it has degree 4 at v_0^{6k+4}. Therefore, c is a circuit although it is the limit of the non-circuits c_{6k+4}.

One can easily manipulate the example so that no c_n with n large enough is a circuit. Indeed, for $i = 1, \ldots, 5$, we can attach a copy H_i of

G to L_i by connecting the root of the normal spanning tree of the wide ladder in H_i to a vertex v_i in L_i that has distance $6k+i-1$ from r in T for some k. We may choose k large enough so that v_i does not belong to one of the first two rungs of L_i; the other two vertices of the first rung of the wide ladder in H_i can then be connected to the two vertices u_i, w_i below v_i in L_i. For each $i < 6$, instead of letting c traversing $w_i u_i v_i$, we let it run along H_i just like the original c did in G. Then the new c is a circuit, but none of the edge sets c_n (with n large enough) it induces on the contracted graphs G_n is a circuit.

5.2 Čech homotopy

Čech homotopy is a way to define the fundamental group of a space (in fact homotopy groups of any dimension) similar to the way that the Čech homology is defined. One shows that the homotopy groups of the nerves of open coverings of a space X form an inverse system and defines the Čech homotopy groups $\check{\pi}_n(X)$ as the inverse limits of the homotopy groups of the nerves. Alternatively, one requires the coverings to be finite; then one obtains the Čech homotopy groups $\check{\pi}_n^F(X)$ based on finite covers. From this construction it is evident that the first Čech homotopy group of $|G|$, based on finite covers or not,[1] is the inverse limit of the fundamental groups of the finite graphs G_n and hence $\check{\pi}_1(|G|) = \check{\pi}_1^F(X)$ is the inverse limit F^* of the free groups F_I.

Eda and Kawamura [25] show that the fundamental group of every one-dimensional continuum X is isomorphic to a subgroup of its first Čech homotopy group based on finite covers. More precisely, they show that the 'canonical' homomorphism $\pi_1(X) \to \check{\pi}_1^F(X)$ (defined by mapping a homotopy class to the limit of its induced classes on the nerves of finite covers) is injective. Together with Lemma 2.9 and the embedding $F_\infty \hookrightarrow F^*$, this result implies our Lemma 3.13.

However, their proof does not explicitely construct a homotopy between paths with same image in $\check{\pi}_1^F(X)$, let alone a homotopy that procedes by contracting pairs of passes through chords. The main achievement of our proof is hence to have shown that homotopies between paths in $|G|$ basically work the same as homotopies in a finite graph, by contracting pairs of passes, one at a time.

[1]For compact spaces, $\check{\pi}_n$ and $\check{\pi}_n^F$ are clearly identical.

Chapter 6

A new homology for locally compact spaces

6.1 Introduction

We have seen in the last section that Čech homology—although its first group is isomorphic to the topological cycle space—fails to properly reflect its relation to the combinatorial structure of G. For this reason, we shall keep at our singular approach to define $\vec{\mathcal{C}}$ in terms of homology. Since by Theorem 4.1 standard singular homology is not the right theory to capture $\vec{\mathcal{C}}$, we shall define a singular-type homology that does so.

As advertised in Chapter 1, we shall define our homology for locally compact Hausdorff spaces with a (fixed) Hausdorff compactification. Recall that these properties are needed to reflect the properties of G and $|G|$ that are fundamental for the success of $\vec{\mathcal{C}}$. Therefore, this class of spaces is the broadest for which we can hope to obtain a homology theory with similar properties as $\vec{\mathcal{C}}$. Note that this class includes, for instance, all locally finite CW-complexes, of any dimension.

The loops constructed in Section 4.3 suggest that our homology should allow to subdivide a 1-simplex infinitely often: Then, every 1-chain in $|G|$ will be homologous to the sum of its passes through edges of G, and hence it will be null-homologous if and only if it lies in the kernel of f. The idea is thus to define the homology so that we obtain essentially the same 1-cycles as in standard singular homology but more boundaries.

The construction of $\vec{\mathcal{C}}$ is based on the idea to consider not only the graph itself but also its ends. Nevertheless, although ends do not play a different role in the definition of $\vec{\mathcal{C}}$ than points in G, elements of $\vec{\mathcal{C}}$ do behave differently at ends. Indeed, elements of $\vec{\mathcal{C}}$ are thin sums of circuits, and as G is locally finite, these circuits are also 'thin' at vertices, i.e. every vertex lies in only finitely many of the closures of the circuits in the family. This does not have to be the case for ends: An end can lie in the closures of infinitely many circuits, even when the circuits form a thin family.

This suggests to require a similar property from the chains in our homology: They will have to be locally finite in G but not at ends.[1] This will enable us to subdivide paths in $|G|$ infinitely often, but the required locally finiteness in G will keep us from obtaining undesired cycles, such as the edges of a double-ray (all directed the same way), which has zero boundary but does not correspond to an element of the cycle space. In the ad-hoc homology we shall define in Section 6.2 we will rule out such cycles by imposing an additional condition on cycles. This will lead to the desired result in dimension 1, i.e. our first homology group will be $\vec{\mathcal{C}}$, but generate problems elsewhere. More precisely, this homology will fail to satisfy the axioms for homology, which is caused precisely by this restriction on cycles.

In Section 6.3 we shall then change our approach slightly: Instead of restricting the group of cycles we shall define chains differently, so as to obtain 1-cycles that are essentially finite and 2-cycles that allow us to subdivide 1-simplices infinitely often. This homology theory will satisfy the axioms, which will be shown in Section 6.4. After briefly discussing the associated cohomology in Section 6.5, we shall show in Section 6.6 that our homology captures, for graphs and dimension 1, precisely the topological cycle space. In this last section, we will again need the preliminary work done in Section 6.2 to prove that the homomorphism f in Theorem 4.1 induces an isomorphism $H_1(G) \simeq \vec{\mathcal{C}}(G)$.

6.2 An ad-hoc homology for locally compact spaces

In this section we describe an ad-hoc way to define homology groups that extend the main properties of the cycle space of graphs to arbitrary dimensions.

While the homology defined in this section succeeds in capturing $\vec{\mathcal{C}}(G)$, it is not a homology theory: It fails to allow for long exact sequences as demanded by the axioms, see end of Section 6.2.1. Nevertheless, our proof that the homology theory we shall construct in Section 6.3 specializes in dimension 1 to yield $\mathcal{C}$ will rely on this section. Moreover, it serves to introduce some of the main ideas in a technically simpler setting.

6.2.1 Definition and examples

Let X be a locally compact Hausdorff space and let $\hat{X}$ be a Hausdorff compactification of X. (See e.g. [1] for more on such spaces.) Note that

[1] The formal definition of 'locally finite' will be given shortly.

every locally compact Hausdorff space is Tychonoff, and thus has a Hausdorff compactification. The kind of spaces we have in mind is that X is a locally finite CW-complex and $\hat{X}$ is its Freudenthal compactification, but formally we do not make any further assumptions. Nevertheless, we will call the points in $\hat{X} \setminus X$ *ends*, even if they are not ends in the usual, more restrictive, sense.

Although our chains, cycles etc. will live in $\hat{X}$, we shall denote their groups as $C_n(X)$, $Z_n(X)$ etc, with reference to X rather than $\hat{X}$: this is because ends will play a special role, so the information of which points of $\hat{X}$ are ends must be encoded in the notation for those groups.

Given points $v_0, \ldots, v_n \in \mathbb{R}^m$ (not necessarily in general position), we write $[v_0, \ldots, v_n]$ for their convex hull. The *natural map* $\Delta^n \to [v_0, \ldots, v_n]$ is the linear map $(t_0, \ldots, t_n) \mapsto \sum t_i v_i$.

If $v_0, \ldots, v_n$ are in general position, the natural map $\Delta^n \to [v_0, \ldots, v_n]$ is clearly a homeomorphism. Then $[v_0, \ldots, v_n]$ is an *n-simplex in* $\mathbb{R}^m$, the point v_i is its *ith vertex*. Every convex hull of $k+1 \leq n$ vertices is a *k-face* of $[v_0, \ldots, v_n]$. We use $[v_0, \ldots, \hat{v}_i, \ldots, v_n]$ to denote the $(n-1)$-face spanned by all the vertices but v_i.

Let us call a family $(\sigma_i \mid i \in I)$ of singular n-simplices in $\hat{X}$ *admissible* if

(i) $(\sigma_i \mid i \in I)$ is locally finite in X, that is, every $x \in X$ has a neighbourhood in X that meets the image of σ_i for only finitely many i;

(ii) every σ_i maps the 0-faces of Δ^n to X.

Note that as X is locally compact, (i) is equivalent to asking that every compact subspace of X meets the image of σ_i for only finitely many i. Condition (ii), like (i), underscores that ends are not treated on a par with the points in X: we allow them to occur on infinitely many σ_i (which (i) forbids for points of X), but not in the fundamental role of images of 0-faces: all simplices must be 'rooted' in X. If X is a countable union of compact spaces, (i) and (ii) together imply that admissible families are countable, i.e. that $|I| \leq \aleph_0$.

When $(\sigma_i \mid i \in I)$ is an admissible family of n-simplices, any formal linear combination $\sum_{i \in I} \lambda_i \sigma_i$ with all $\lambda_i \in \mathbb{Z}$ is an *n-sum in X*.[2] We regard n-sums $\sum_{i \in I} \lambda_i \sigma_i$ and $\sum_{j \in J} \mu_j \tau_j$ as *equivalent* if for every n-simplex ρ we have $\sum_{i \in I, \sigma_i = \rho} \lambda_i = \sum_{j \in J, \tau_j = \rho} \mu_j$. Note that these sums are well-defined since an n-simplex can occur only finitely many times in an admissible family. We write $C_n(X)$ for the group of *n-chains*, the equivalence classes of n-sums. The elements of an n-chain are its *representations*. Clearly

[2]In standard singular homology, one does not usually distinguish between formal sums and chains. It will become apparent soon why we have to make this distinction.

every n-chain c has a unique representation whose simplices are pairwise distinct—which we call the *reduced representation* of c—, but we shall consider other representations too. The subgroup of $C_n(X)$ consisting of those n-chains that have a finite representation is denoted by $C'_n(X)$.

The boundary operators $\partial_n \colon C_n \to C_{n-1}$ are defined by extending linearly from $\partial_n \sigma_i$, which are defined as usual in singular homology. Note that ∂_n is well defined (i.e., that it preserves the required local finiteness), and $\partial_{n-1}\partial_n = 0$. Chains in $\operatorname{Im} \partial$ will be called *boundaries*.

As n-cycles, we do *not* take the entire kernel of ∂_n. Rather, we define $Z'_n(X) := \operatorname{Ker} (\partial_n \restriction C'_n(X))$, and let $Z_n(X)$ be the set of those n-chains that are sums of such finite cycles:

$$Z_n(X) := \Big\{ \varphi \in C_n(X) \;\Big|\; \varphi = \sum_{j \in J} z_j \text{ with } z_j \in Z'_n(X) \ \forall j \in J \Big\}.$$

More precisely, an n-chain $\varphi \in C_n(X)$ shall lie in $Z_n(X)$ if it has a representation $\sum_{i \in I} \lambda_i \sigma_i$ for which I admits a partition into finite sets I_j $(j \in J)$ such that, for every $j \in J$, the n-chain $z_j \in C'_n(X)$ represented by $\sum_{i \in I_j} \lambda_i \sigma_i$ lies in $Z'_n(X)$. Any such representation of φ as a formal sum will be called a *standard representation* of φ *as a cycle*.[3] We call the elements of $Z_n(X)$ the *n-cycles* of X.

The chains in $B_n(X) := \operatorname{Im} \partial_{n+1}$ then form a subgroup of $Z_n(X)$: by definition, they can be written as $\sum_{j \in J} \lambda_j z_j$ where each z_j is the (finite) boundary of a singular $(n+1)$-simplex. We therefore have homology groups

$$H_n(X) := Z_n(X)/B_n(X)$$

as usual.

Note that if X is compact, then all admissible families and hence all chains are finite, so the homology defined above coincides with the usual singular homology. The characteristic feature of this homology is that while infinite cycles are allowed, they are always of 'finite character': in any standard representation of an infinite cycle, every finite subchain is contained in a larger finite subchain that is already a cycle.

For graphs and Freudenthal compactifications, the finite character of this homology is also shown in another aspect: We will show in the next section that every 1-cycle—finite or infinite—is homologous to a cycle whose reduced representation consists of a single loop.

Let us now define relative homology groups $H_n(X, A)$. Normally, these groups are defined for all subsets $A \subset X$. In our case, the subspace A

[3]Since the σ_i need not be distinct, φ has many representations by formal sums. Not all of these need admit a partition as indicated—an example will be given later in the section.

has to satisfy further conditions. Since we wish to consider chains in A, in our sense, A has to be locally compact and come with a compactification $\hat{A}$. Chains in A have to be chains also in X, hence we further need that $\hat{A} \subset \hat{X}$, and that ends of A lie in $\hat{X} \setminus X$, that is, they have to be ends of X.

Let A be a *closed* subset of X (but not necessarily closed in $\hat{X}$). Since X is locally compact, so is A. Let $\hat{A}$ denote the closure of A in $\hat{X}$. Then $\hat{A}$ is a compactification of A, and $\hat{A} \setminus A \subset \hat{X} \setminus X$. Clearly, admissible families of simplices in A are also admissible in X. We define $H_n(X, A)$ as follows. Let $C_n(X, A)$ be the quotient group $C_n(X)/C_n(A)$,[4] and let $C'_n(X, A)$ be the subgroup of all its elements $\varphi + C_n(A)$ with $\varphi \in C'_n(X)$. Define $Z'_n(X, A)$ as the kernel of the quotient map $C_n(X, A) \to C_{n-1}(X, A)$ of ∂_n restricted to $C'_n(X, A)$, and $B_n(X, A)$ as the image of the quotient map $C_{n+1}(X, A) \to C_n(X, A)$ of ∂_{n+1}. Then define $Z_n(X, A)$ from $Z'_n(X, A)$ as before, and put $H_n(X, A) = Z_n(X, A)/B_n(X, A)$. Clearly, $H_n(X, \emptyset) = H_n(X)$.

Before proving in the next section that this homology (for graphs and dimension 1) captures the cycle space, let us look at an example which might indicate whether we obtain the desired cycles. For simplicity, we will restrict our attention to absolute homology. Consider the *double ladder*. This is the 2-ended graph G with vertices v_n and v'_n for all integers n, and with edges e_n from v_n to v_{n+1}, edges e'_n from v'_n to v'_{n+1}, and edges f_n from v_n to v'_n. The 1-simplices corresponding to these edges, oriented in their natural directions, are θ_{e_n}, $\theta_{e'_n}$, and θ_{f_n}, see Figure 6.1.

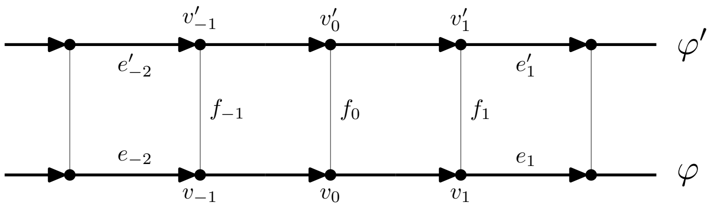

Figure 6.1: The 1-chains φ and φ' in the double ladder.

In order to let the elements of our homology be defined, let $\hat{G}$ be any Hausdorff compactification of G. (One could, for instance, choose the Freudenthal compactification $|G|$ of G.) For the infinite chains φ and φ' represented by $\sum \theta_{e_n}$ and $\sum \theta_{e'_n}$, respectively, and for $\psi := \varphi - \varphi'$ we have $\partial\varphi = \partial\varphi' = \partial\psi = 0$, and neither sum as written above contains a finite cycle. However, we can rewrite ψ as $\psi = \sum z_n$ with finite cycles

[4]Formally, $C_n(A)$ is not a subset of $C_n(X)$, because the equivalence classes of n-sums in X are larger than those in A. For instance, every formal sum $\sigma - \sigma$ with σ a singular n-simplex in X that does not live in A is part of the equivalence class of the empty n-sum in X, but not in A. But there is a natural embedding $C_n(A) \hookrightarrow C_n(X)$: map an n-chain in A to the n-chain in X with the same reduced representation.

$z_n = \theta_{e_n} + \theta_{f_{n+1}} - \theta_{e'_n} - \theta_{f_n}$. This shows that $\psi \in Z_1(G)$, although this was not visible from its original representation.

By contrast, one can show that $\varphi \notin Z_1(G)$ if $\hat{G}$ is the Freudenthal compactification of G. This follows from Theorem 6.3 below and Theorem 2.10, but is not obvious. For example, one might try to represent φ as $\varphi = \sum_{n=1}^{\infty} z'_n$ with $z'_n := \theta_{e_{-n}} + \theta_{n-1} + \theta_{e_n} - \theta_n$, where $\theta_n \colon [0,1] \to e_{-n} \cup \cdots \cup e_n$ maps 0 to v_{-n} and 1 to v_{n+1}, see Figure 6.2.

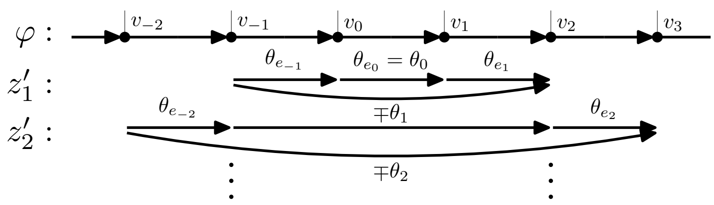

Figure 6.2: Finite cycles summing to φ—by an inadmissible sum.

This representation of φ, however, although well defined as a formal sum (since every simplex occurs at most twice), is not a legal 1-sum, because its family of simplices is not locally finite and hence not admissible. (The point v_0, for instance, lies in every simplex θ_i.)

We close this section with a proof that this homology is not a homology theory since it fails to satisfy the long exact sequence axiom. To see this, let $A \subset X$ consist of a single point a in X and assume there is a path π in $\hat{X}$ from a to an end. This assumption is satisfied in every path-connected, non-compact space X with a path-connected compactification $\hat{X}$, for instance in every connected locally finite graph and its Hausdorff compactification. The 0-chain $c = -\sigma$ in A, where $\sigma \colon \{0\} \to A$, is a 0-cycle whose homology class in $H_0(A)$ lies in the kernel of $\iota_* : H_0(A) \to H_0(X)$ (because $c = \partial\tau$ for $\tau = \sum_{i=1}^{\infty} \pi \restriction [1 - 2^{1-i}, 1 - 2^{-i}]$) but not in the image of $\partial_1 \colon H_1(X, A) \to H_0(A)$ (because clearly no finite 1-cycle in X can have boundary c, and no infinite 1-cycle in X that is a sum of finite cycles can have boundary c, since by Condition (i) only finitely many of those finite cycles meet a). Hence the long sequence for the pair (X, A) fails to be exact at $H_0(A)$.

6.2.2 $H_1(G)$ equals $\mathcal{C}(G)$

In this section we show that, for graphs G, the group $H_1(G)$ defined in Section 6.2.1 is canonically isomorphic to the topological cycle space $\vec{\mathcal{C}}(G)$ of G.

In analogy to our notation of Section 4.2, we shall denote this isomorphism by $f \colon H_1(G) \to \vec{\mathcal{C}}(G)$. In our definition of f we shall have to refer to the map which, in Section 4.2, was denoted as $f \colon H_1(|G|) \to \vec{\mathcal{E}}(G)$; this map will now be denoted as f'. (Recall that $|G|$ denotes the Freudenthal compactification of G, and that $H_1(|G|)$ is its usual first singular homology group.[5]) When G is finite, our new function f will coincide with f'.

In order to define f, let $\varphi \in Z_1(G)$ be given in any standard representation $\varphi = \sum_{i\in I} \lambda_i\sigma_i$ as a cycle, and let $\vec{e} \in \vec{E}$ be any oriented edge. We shall first define $f([\varphi])(\vec{e}\,) \in \mathbb{Z}$ with reference to φ and its given representation as a cycle, and then show that our definition does not depend on these choices.

To define $f([\varphi])(\vec{e}\,)$, we show that for all large enough finite subchains $\varphi' \in Z_1'(G)$ of φ the values of $f'([\varphi'])(\vec{e}\,)$ agree (the homology class $[\varphi']$ being taken in $H_1(|G|)$), and set $f([\varphi])(\vec{e}\,)$ to this common value. Write I_e for the set of those $i \in I$ whose σ_i meets e; since e is compact and $(\sigma_i \mid i \in I)$ is a good family, I_e is a finite set.

Let $\pi\colon H_1(S^1) \to \mathbb{Z}$ and $f_e\colon |G| \to S^1$ be defined as in Section 4.2, and write $(f_e)_\sharp\colon C_1(|G|) \to C_1(S^1)$ for the chain map induced by f_e.

Lemma 6.1. *For all finite sets I' such that $I_e \subseteq I' \subseteq I$ and $\varphi' := \sum_{i\in I'} \lambda_i\sigma_i \in Z_1(|G|)$, the values of $f'([\varphi'])(\vec{e}\,)$ agree.*

Proof. Let $\varphi_e := \sum_{i\in I_e} \lambda_i\sigma_i$. We show that even if φ_e is not a cycle in $|G|$, the chain $(f_e)_\sharp(\varphi_e)$ is a cycle in S^1 homologous to $(f_e)_\sharp(\varphi')$ for every φ' as stated. Then, by definition of f',

$$f'([\varphi'])(\vec{e}\,) = \pi((f_e)_*([\varphi'])) = \pi([(f_e)_\sharp(\varphi')]) = \pi([(f_e)_\sharp(\varphi_e)])$$

for all such φ', and the result follows.

For a proof of $[(f_e)_\sharp(\varphi_e)] = [(f_e)_\sharp(\varphi')]$, note that for all $i \in I \setminus I_e$ the map $f_e \circ \sigma_i$ is constant (with value $1 \in \mathbb{C}$). So for such i, $f_e \circ \sigma_i$ is a null-homologous cycle. But $(f_e)_\sharp(\varphi_e)$ differs from $(f_e)_\sharp(\varphi')$, which is a cycle, precisely by the terms $\lambda_i(f_e \circ \sigma_i)$ with $i \in I' \setminus I_e$. Hence $(f_e)_\sharp(\varphi_e)$ too is a cycle, and it is homologous to $(f_e)_\sharp(\varphi')$. □

We now define $f\colon H_1(G) \to \vec{\mathcal{E}}(G)$ by letting $f([\varphi])$ map an oriented edge $\vec{e}$ to the common value of $f'([\varphi'])(\vec{e}\,)$ for all φ' as in Lemma 6.1. In order to show that f is well defined, let $\varphi \in Z_1(G)$ and $\psi \in B_1(G)$ be given in any standard representations $\varphi = \sum_{i\in I} \lambda_i\sigma_i$ and $\psi = \sum_{i\in J} \lambda_i\sigma_i$ with $I \cap J = \emptyset$. We show that f assigns the same value to $[\varphi] = [\varphi + \psi]$ no matter whether we base its computation on φ or on $\varphi+\psi$: this proves that $f([\varphi])$ depends neither on the choice of φ as a representative of $[\varphi]$ nor on its representation as $\sum_{i\in I} \lambda_i\sigma_i$.

Given $\vec{e} \in \vec{E}$, let I_e be the set of all $i \in I$ such that σ_i meets e, and define J_e likewise. Let $I' \subseteq I$ and $J' \subseteq J$ be finite sets containing I_e and J_e,

[5]We shall use $C_1(G)$, $Z_1(G)$, $B_1(G)$ and $H_1(G)$ to refer to our new homology of $|G|$ that relies on the information of which points of $|G|$ are ends, while $C_1(|G|)$, $Z_1(|G|)$, $B_1(|G|)$ and $H_1(|G|)$ continue to refer to the usual singular homology of the space $|G|$.

respectively, such that $\varphi' := \sum_{i\in I'} \lambda_i\sigma_i \in Z_1(|G|)$ and $\psi' := \sum_{i\in J'} \lambda_i\sigma_i \in B_1(|G|)$; such sets exist since φ and ψ are given in standard representations. Then

$$f'([\varphi'])(\vec{e}\,) = f'([\varphi' + \psi'])(\vec{e}\,).$$

For our new function f, its value of $[\varphi] = [\varphi + \psi]$ computed with reference to φ equals the left-hand side of this equation, while its value computed with reference to $\varphi+\psi$ equals the right-hand side. This completes the proof that f is well defined. Note that if φ is finite, then trivially $f[\varphi] = f'[\varphi]$, where $[\varphi]$ is taken in $H_1(G)$ and in $H_1(|G|)$, respectively.

Since f' is a homomorphism with image $\vec{\mathcal{C}}(G)$ (see Section 4.3), Lemma 6.1 implies that so is f. Indeed, for a proof that $f([\varphi]) \in \vec{\mathcal{C}}(G)$ consider the finite oriented cuts $\vec{F}$ of G, and apply Theorem 2.10 to any finite subchain φ' of φ containing all the simplices that meet this cut. The proof that f is surjective is the same as in Section 4.2: every element of $\vec{\mathcal{C}}(G)$ has the form $f([\tau])$ with τ a single loop. Thus in fact,

$$\vec{\mathcal{C}}(G) \subseteq f(H_1(G)) \subseteq \vec{\mathcal{C}}(G)$$

with equality.

Our final goal is to show that f is injective. For finite G, the standard proof is to rewrite a given cycle $z \in Z_1(G)$ as a homologous sum of simplices each traversing exactly one edge. If $[z] \in \operatorname{Ker} f$, every edge is traversed equally often in both directions, and we can pair up the simplices traversing it accordingly. Each pair is a boundary, and hence so is z.

The reason why this proof does not work for f' on $H_1(|G|)$ is that the simplices even in a finite cycle can traverse infinitely many edges. The proof would therefore require us to break up the given finite cycle into a 'homologous' infinite chain, which is impossible in $H_1(|G|)$.

In our new setup, however, this can indeed be done. In fact, it turns out that our restriction that any boundary chains to be added must be locally finite exactly strikes the balance between being restrictive enough to rule out counterexamples like φ in Section 6.2.1 and being general enough to allow the subdivision into chains of single edges even of complicated cycles like our non-injectivity example from Section 4.3.

This is shown in the following lemma. Although its proof looks somewhat technical, the idea is very simple, so let us describe it informally first. Consider a 1-simplex τ traversing infinitely many edges. Our task is to 'subdivide it infinitely often', into 1-simplices $\sigma_1, \sigma_2, \ldots$ each traversing exactly one edge, by adding a locally finite sum of boundaries. We begin by targeting the first pass of τ through an edge, $e = uv$ say. Let σ_1 be this pass, and let τ' and τ'' be the segments of τ before and after σ_1. We now subdivide τ at u and v: we add to τ the boundary $\tau' + \sigma_1 + \tau'' - \tau$, to obtain the chain $\tau' + \sigma_1 + \tau''$. Next, we target the second pass of τ through an edge, σ_2. If this is a pass of τ', say, with segments α and β before and after σ_2, we add the boundary $\alpha + \sigma_2 + \beta - \tau'$ to insert σ_2 into our chain while eliminating τ'. Doing this for all passes of τ in turn should leave us at the limit with only the chain $\sigma_1 + \sigma_2 + \dots$, since all other simplices are eliminated again when the earliest pass they contain is targeted. The main task of the formal proof of this, except for the inevitable book-keeping, is to ensure that all the boundaries we add do indeed form a locally finite chain, i.e. an element of $B_1(G)$.

Lemma 6.2. *For every $z \in Z'_1(G)$ there exist a chain $\varphi = \sum_{i\in I} \sigma_i \in Z_1(G)$ and a chain $b \in B_1(G)$ such that $z + b = \varphi$, every σ_i maps $[0,1]$ homeomorphically to some edge e, and all these edges e as well as the images of the simplices in b are contained in the image of the 1-simplices in z.*

Proof. By the additivity of $Z'_1(G)$, we may clearly assume that z is an elementary cycle consisting of a single loop τ_0 that is based at a vertex and is not null-homotopic. In particular, τ_0 traverses an edge. Since τ_0 traverses every edge only finitely often (Lemma 2.9), τ_0 contains only countably many passes through edges, $\pi_1, \pi_2, \dots$ say, which we reparametrize as maps from $[0,1]$.

In each of at most ω steps we shall add to our then current finite cycle

$$z_n = \sum_{i=1}^{n} \sigma_i + \sum_{j\in J_n} \tau_j$$

(which initially is $z_0 = \tau_0$) finitely many simplices σ_i or τ_j with coefficients 1 or -1 so that the sum of simplices added lies in $B_1(G)$. We shall make sure that all these simplices added or deleted form a good family; in particular, their sum will not depend on the order of summation, although this order will help us with our book-keeping. The result will be a chain of the form $\sum_{i\in I} \sigma_i + \sum_{j\in J} \tau_j$ in which every τ_j is a null-homotopic loop (in particular $0 \notin J$) and the σ_i are those required in the statement of the lemma.

We shall choose the z_n inductively so as to satisfy the following conditions, which hold for $n = 0$ with $J_0 = \{0\}$:

(i) $\sigma_1, \dots, \sigma_n$ and all τ_j ($j \in J_n$) are paths in $|G|$ between (possibly identical) vertices;

(ii) if $n \geq 1$, every τ_j ($j \in J_n$) is a segment of some τ_i with $i \in J_{n-1}$;

(iii) if $n \geq 1$, there exists $j(n)$ such that $J_{n-1} \setminus J_n = \{j(n)\}$ and the finite chain $b_n := \sigma_n - \tau_{j(n)} + \sum_{j'\in J_n \setminus J_{n-1}} \tau_{j'}$ lies in $B_1(G)$;

(iv) σ_n is homotopic to π_n relative to $\{0,1\}$;

(v) suitably reparametrized, $(\pi_{n+1}, \pi_{n+2}, \dots)$ is the family of all edge-passes of the paths τ_j $(j \in J_n)$; specifically, the edge-passes in the paths $\tau_{j'}$ with $j' \in J_n \setminus J_{n-1}$ are precisely those in $\tau_{j(n)}$ other than π_n.

Assuming that z_{n-1} satisfies these conditions, let us define z_n. If π_n does not exist, we terminate the construction, putting $I := \{1, \dots, n-1\}$ and $J := J_{n-1}$. If it does, then by (v) for $n-1$ there is a unique $j \in J_{n-1}$ such that π_n is an edge-pass in τ_j. The path τ_j is a concatenation of three segments α, π_n, and β, where α and β may have trivial domain. Let J_n be obtained from J_{n-1} by removing $j =: j(n)$ and adding new indices j', j'' for $\alpha =: \tau_{j'}$ and $\beta =: \tau_{j''}$ whenever these maps are paths (i.e., have non-trivial domain), reparametrizing each to domain $[0, 1]$. Let σ_n be an injective path that is homotopic to π_n relative to $\{0, 1\}$. Clearly, z_n again satisfies the conditions. If the process continues for ω steps, we complete it by putting $I := \mathbb{N}$, and letting $J := \bigcap_{n\in\mathbb{N}} \bigcup_{k>n} J_k$ consist of those j that are eventually in J_n.

Let us take a look at the simplices τ_j with $j \in J$. By definition of J, we have $j \in J_n$ for all large enough n. By (i) and (ii), τ_j is a segment of τ_0 between two vertices, and by (v) it contains none of the passes $\pi_1, \pi_2, \dots$. So it does not traverse any edge. Hence,

$$\tau_j \textit{ is a null-homotopic loop based at a vertex.} \tag{6.1}$$

Notice that for only finitely many $j \in J$ can τ_j be based at the same vertex v. Indeed, given $j \in J$, let n be the unique integer such that $j \in J_n \setminus J_{n-1}$. Then, since $0 \notin J$, τ_j is a segment of τ_0 followed or preceded by π_n, and hence π_n is a pass through an edge at v. Since τ_0 contains only finitely many such passes, this can happen for only finitely many n, and indices j first appearing in J_n for different n are distinct.

Next, let us show the following:

The family of all simplices added or deleted in the construction, that is, of all σ_i, $i \in I$, all τ_j, $j \in J$, and all $\tau_{j(n)}$, is locally finite and hence good. (6.2)

To prove (6.2), let x be any point in G. If x is a vertex, let E_x be the set of edges at x; if $x \in \mathring{e}$ for an edge e, let $E_x := \{e\}$. Choose an open neighbourhood U of x contained in $\bigcup E_x$. Since τ_0 traverses each edge in E_x only finitely often, only finitely many of the paths σ_i, $i \in I$, meet U. Similarly, any path τ_j with $j \in J$ that meets U must be based at a vertex incident with an edge in E_x. Since there are only finitely many such vertices, and at each only finitely many τ_j are based, only finitely many τ_j with $j \in J$ meet U. Finally, consider a path $\tau_{j(n)}$. This path traverses an edge (in π_n), so if it meets U it must also traverse an edge in E_x or adjacent to an edge in E_x. Only finitely many of the passes π_k traverse such edges. By (v), any τ_j containing π_k satisfies $j \in J_1 \cup \cdots \cup J_{k-1}$, so $j(n) \in J_1 \cup \cdots \cup J_{k-1}$ for the largest such k. Since this is a finite set and the map $n \mapsto j(n)$ is injective, only finitely many n are such that $\tau_{j(n)}$ meets U. This completes the proof of (6.2).

To complete the proof of the lemma, we show that $z + b = \varphi$ for $b := \sum_{i \in I} b_i - \sum_{j \in J} \tau_j$, and in particular that $b \in B_1(G)$. By (6.2), the family of all simplices in b is good, so $b \in B_1(G)$ by (iii) and (6.1). Likewise, the family of all σ_i is good. Since

$$z + \sum_{i \in I} b_i = \sum_{i \in I} \sigma_i + \sum_{j \in J} \tau_j$$

by construction, we deduce that $z + b = \sum_{i \in I} \sigma_i = \varphi$ as desired. □

We can now easily complete the proof that our function $f\colon H_1(G) \to \vec{\mathcal{C}}(G)$ is injective. Consider any $[z] \in \operatorname{Ker} f$. As $z \in Z_1(G)$, it has a standard representation as $z = \sum_{j \in J} z_j$ with all $z_j \in Z_1'(G)$. By Lemma 6.2, there are $b_j \in B_1(G)$ $(j \in J)$ such that $z_j + b_j = \varphi_j$, where $\varphi_j = \sum_{i \in I_j} \sigma_i$ is a chain of simplices each traversing exactly one edge, and these edges as well as the images of the simplices in b_j lie in the image of z_j. The fact that z is a locally finite chain therefore implies that so are

$$b := \sum_{j \in J} b_j \textit{ and } \varphi := \sum_{j \in J} \varphi_j.$$

Indeed, every $x \in G$ has an open neighbourhood U that meets the images of simplices in z_j for only finitely many j; let J_x be the set of those j. Hence U does not meet the images of any simplices in b_j or φ_j for $j \notin J_x$. For each $j \in J_x$, we can find an open neighbourhood $U_j \subseteq U$ of x that meets only finitely many simplices in b_j or φ_j, because b_j and φ_j are well-defined chains. The intersection of these finitely many U_j thus is an open neighbourhood of x that meets only finitely many simplices in b or in φ, showing that b and φ are well-defined chains.

For $I := \bigcup_{j \in J} I_j$, we thus have $Z_1 \ni z + b = \varphi = \sum_{i \in I} \sigma_i$, with $b \in B_1(G)$. Since $[z] \in \operatorname{Ker} f$, we thus have $[\varphi] \in \operatorname{Ker} f$. Therefore the loops formed by the elementary cycles in $\varphi = \sum_{i \in I} \sigma_i$ traverse, in total, each edge of G equally often in both directions (see Lemma 4.2). Since each of the σ_i traverses precisely one edge, we can thus pair them up into cancelling pairs $\sigma_i + \sigma_{i'} \in B_1(G)$, where σ_i and $\sigma_{i'}$ traverse the same edge but in opposite directions. Hence $\varphi = \sum_{i \in I} \sigma_i \in B_1(G)$, giving $z = \varphi - b \in B_1(G)$ as desired.

We have thus shown that f is a group isomorphism between $H_1(G)$ and $\vec{\mathcal{C}}(G)$. Moreover, if we restrict f to those homology classes that are represented by finite cycles, then by Lemma 4.4 and the fact that f and f'

coincide on finite cycles we obtain that this this restriction is still onto and hence an isomorphism. This implies that all homology classes have finite representatives. We thus have our first main result of this chapter:

Theorem 6.3. *The function f is a group isomorphism between $H_1(G)$ and $\vec{\mathcal{C}}(G)$. Moreover, for every class $c \in H_1(G)$ there is a finite cycle $z \in Z_1'(G)$ with $c = [z]$.*

6.3 A new homology for locally compact spaces

In this section we define a homology theory that implements the same ideas as our ad-hoc homology of Section 6.2, but which will satisfy all the usual axioms. To achieve this, we shall encode all the properties we need into the definition of chains—rather than restricting both chains and cycles, as in Section 6.2. Our homology will also be defined for disjoint unions of compactifications, i.e. for $X = \bigsqcup X_k$ and $\hat{X} = \bigsqcup \hat{X}_k$ where each $\hat{X}_k$ is a compactification of X_k. Nevertheless, we will start with the definition for compact $\hat{X}$ and then extend it to unions of compactifications.

Let X be a locally compact Hausdorff space, and let $\hat{X}$ be a Hausdorff compactification of X. We define admissible families and n-sums as in Section 6.2. All other notation will now be defined differently.

In order to capture $\mathcal{C}(G)$ in dimension 1 for locally finite graphs, we have to consider chains consisting of infinitely many simplices, by Theorem 4.1. On the other hand, if one allows infinite chains without further restrictions, one obtains cycles like φ in Figure 6.1, which does not correspond to an element of $\mathcal{C}(G)$. The solution to this dilemma is to allow infinitely many simplices only if they are of a certain type.

Call a singular n-simplex σ in $\hat{X}$ *degenerate* if it is lower dimensional in the following sense: There is a compact Hausdorff space X_σ of dimension at most $n-1$ such that σ can be written as the composition of continuous maps $\Delta^n \to X_\sigma \to \hat{X}$. (Note that X_σ is normal as it is compact and Hausdorff. Hence the topological dimension of X_σ is defined as in Section 2.1) The idea behind this definition is that the 2-simplices whose boundaries are added to a 1-simplex to subdivide it, like in the proof of Lemma 6.2, are easily seen to be degenerate, and we need to be able to add infinitely many of those simplices. (In fact, all 2-simplices in $|G|$ are trivially degenerate by Lemma 2.8)

As the empty space is the only space of dimension -1, and every 0-dimensional space is totally disconnected, we have that no singular 0-simplex is degenerate and a singular 1-simplex is degenerate if and only if it is constant.

Denote by $C'_n(X)$ the group of equivalence classes of n-sums. (Recall that two n-sums are called equivalent if every n-simplex appears equally often—taking account of the multiplicities λ_i—in both sums.) As before, the elements of a class $c \in C'_n(X)$ are its *representations*, its unique representation $\sum \lambda_i \sigma_i$ with pairwise distinct σ_i is the *reduced representation* of c. We call c *good* if the simplices σ_i in its reduced representation are degenerate for all but finitely many $i \in I$. An *n-chain in X* is an equivalence class $c \in C'_n(X)$ that can be written as $c = c_1 + \partial c_2$, where both $c_1 \in C'_n(X)$ and $c_2 \in C'_{n+1}(X)$ are good. In other words, c is an n-chain if and only if it has a representation $\sum_{i \in I} \lambda_i \sigma_i$ for which I is the disjoint union of a finite set I_0, a (possibly infinite) set I_1, and finite sets I_j, $j \in J$, such that each σ_i, $i \in I_1$, is degenerate, and each sum $\sum_{i \in I_j} \lambda_i \sigma_i$ is the boundary of a degenerate singular $(n+1)$-simplex.[6] We call such a representation a *standard representation* of c. Note that a standard representation will not, in general, be a reduced representation, and vice versa, a reduced representation does not have to be standard.

We write $C_n(X)$ for the group of all n-chains in X. As usual, we write $Z_n(X) := \mathrm{Ker}\, \partial_n$ and $B_n(X) := \mathrm{Im}\, \partial_{n+1}$. The elements of Z_n are *n-cycles*, those of B_n are *boundaries*. Clearly, $B_n \subseteq Z_n$, so we can define the *homology groups* $H_n(X) := Z_n/B_n$ as usual.

Since a cycle $c_1 + \partial c_2$ as above represents the same homology class as c_1 does, we have at once:

Proposition 6.4. *Every homology class is represented by a good n-cycle.*

As no singular 0-simplex is degenerate, this means that every homology class in $H_0(X)$ is represented by a finite 0-cycle. Moreover, as every degenerate 1-simplex is constant and hence equivalent to the boundary of a constant (and thus degenerate) 2-simplex, we have the same in dimension 1:

Proposition 6.5. *Every homology class in $H_0(X)$ or in $H_1(X)$ is represented by a finite cycle.*

Let us now define relative homology groups. Consider a closed subset A of X and write $\hat{A}$ for the closure of A in $\hat{X}$. In order to make all the axioms work, we additionally require the boundary of $\hat{A}$ in $\hat{X}$ to be a (compact) subset of X. In the case of graphs and their Freudenthal compactification, this is the case for instance if A is a component of $G - S$ for S a finite vertex set. In infinite graph theory, it is an often used procedure to contract such components (see, for instance, the construction of the graphs

[6]Hence I_0 contains all indices of the non-degenerate n-simplices in the representation of c_1 certifying it as a good chain, as well as all indices of the boundary simplices of the non-degenerate $(n+1)$-simplices in the according representation of c_2.

G_i in Chapter 5), so for our purposes it does not seem too restrictive to only consider such subsets. We call (X, A) an *admissible pair*. Like in Section 6.2, we have immediately that admissible families of simplices in A are admissible also in X. Now let $C_n(X, A) = C_n(X)/C_n(A)$, let $Z_n(X, A)$ be the kernel of the quotient map $C_n(X, A) \to C_{n-1}(X, A)$ of ∂_n, and $B_n(X, A)$ the image of the quotient map $C_{n+1}(X, A) \to C_n(X, A)$ of ∂_{n+1}, and define $H_n(X, A) := Z_n(X, A)/B_n(X, A)$.

Having defined the homology groups for compactifications, we now extend it to disjoint unions of compactifications as follows: If $X = \bigsqcup X_k$, $\hat{X} = \bigsqcup \hat{X}_k$, and A is a closed subspace of X such that for each k the pair (X_k, A_k) is admissible, where $A_k := A \cap X_k$, we call (X, A) an *admissible pair*. For an admissible pair (X, A), define $C_n(X, A)$ as the direct sum $\bigoplus C_n(X_k, A_k)$. The homology groups $H_n(X)$ and $H_n(X, A)$ are then defined in the obvious way.

Our earlier definitions of admissible families, n-sums, and n-chains for compact $\hat{X}$ also extend naturally to disjoint unions $X = \bigsqcup X_k$ as follows: A family of singular n-simplices in X is *admissible* if its subfamily of simplices in X_k is admissible for finitely many k and empty for all other k. (Note that every simplex lives in a unique $\hat{X}_k$ as we assumed the $\hat{X}_k$ to be disjoint.) An *n-sum in X* is a formal sum $\sum_{i \in I} \lambda_i \sigma_i$ where $(\sigma_i)_{i \in I}$ is an admissible family. The equivalence classes of n-sums form a group $C'_n(X)$, an element c of $C'_n(X)$ is *good* if it has a representation in which all but finitely many simplices are degenerate, and an *n-chain in X* is a class $c \in C'_n(X)$ that can be written as $c = c_1 + \partial c_2$ with good $c_1 \in C'_n(X)$ and good $c_2 \in C'_{n+1}(X)$. It is easy to see that $C_n(X)$, defined earlier as $\bigoplus_k C_n(X_k)$, is indeed the group of n-chains in X.

In standard homology, it is trivial that a chain in X all of whose simplices live in $\hat{A}$ is also a chain in A. In our case, this is not immediate: If all simplices in the reduced representation of a chain in X live in $\hat{A}$, this does not imply directly that there is a *standard* representation all whose simplices live in $\hat{A}$. Indeed, if there is an infinite admissible family of degenerate $(n+1)$-simplices that do not live in $\hat{A}$ but whose boundaries do, then the sum of their boundaries is the representation of an n-chain in X, and all simplices in the reduced representation of this chain live in $\hat{A}$. But as soon as this reduced representation consists of infinitely many non-degenerate n-simplices, we do not know whether it does also represent a chain in A. Here we can use that (X, A) is an admissible pair: As each $\hat{A}_k$ has a compact boundary that is contained in X_k, there is no admissible family as above. An one can indeed show that a chain in X is a chain in A as soon as their simplices live in $\hat{A}$. More generally, we have the following:

Lemma 6.6. *Let $\sum_{i \in I} \lambda_i \sigma_i$ be a reduced representation of a chain c in X and let $I' \subseteq I$ be the set of those indices with* $\operatorname{Im} \sigma_i \subseteq \hat{A}$. *Then $\sum_{i \in I'} \lambda_i \sigma_i$*

is the reduced representation of a chain c_A in A.

Proof. Choose a standard representation of c, i.e.

$$c = \sum_{i \in I_0} \lambda_i \sigma_i + \sum_{i \in I_1} \lambda_i \sigma_i + \sum_{j \in J} \lambda_j \partial \tau_j, \tag{6.3}$$

where I_0 is finite and each simplex $\sigma_i,\ i \in I_1$, and $\tau_j,\ j \in J$, is degenerate. Note that not all simplices occurring in this representation have to live in $\hat{A}$, this only has to hold for the n-simplices that are part of the reduced representation of c.

Let $I_0' \subseteq I_0$, $I_1' \subseteq I_1$, and $J' \subseteq J$ be the index sets of those simplices that live in $\hat{A}$. Let further $(\sigma_k)_{k \in K}$ be the family of those n-simplices living in $\hat{A}$ that are a face of some τ_j with $j \in J \setminus J'$, and let λ_k be the multiplicity in which σ_k occurs in the sum $\sum_{j \in J \setminus J'} \lambda_j \partial \tau_j$. Note that K is finite since every $\tau_j,\ j \in J \setminus J'$, with a face $\sigma_k,\ k \in K$, meets the compact boundary of $\hat{A}$ and $(\tau_j)_{j \in J \setminus J'}$ is admissible. Now the n-sum

$$\sum_{i \in I_0'} \lambda_i \sigma_i + \sum_{i \in I_1'} \lambda_i \sigma_i + \sum_{j \in J'} \lambda_j \partial \tau_j + \sum_{k \in K} \lambda_k \sigma_k \tag{6.4}$$

is a standard representation of a chain c_A in A, and by construction the reduced representation of c_A is precisely $\sum_{i \in I'} \lambda_i \sigma_i$: Each simplex in (6.4) also occurs in the representation from (6.3), and it does so in the same multiplicity. Hence all simplices in (6.4) occur with the same multiplicity in the reduced representation of c and c_A, showing that c_A is represented by $\sum_{i \in I'} \lambda_i \sigma_i$. □

The homology defined in this section captures $\vec{\mathcal{C}}$ in dimension 1. We will prove this in Section 6.6 using Theorem 6.3, but first we show in Section 6.4 that it is indeed a homology theory.

6.4 Verifying the axioms

In this section we show that the homology theory we defined satisfies the axioms for homology (see Section 2.3). As a preliminary step we have to show that continuous functions between spaces induce homomorphisms between their homology groups. This will not work for arbitrary continuous functions: As we distinguish between ends and other points, our functions will have to respect this distinction.

Let locally compact Hausdorff spaces $X, \hat{X}$ and $Y, \hat{Y}$ be given, where $X = \bigsqcup X_k$ and $\hat{X} = \bigsqcup \hat{X}_k$ with $\hat{X}_k$ a compactification of X_k, and similarly for $Y = \bigsqcup Y_l$. Let $A \subseteq X$ and $B \subseteq Y$ be closed subspaces such that (X, A) and (Y, B) are admissible pairs. As before, we write $\hat{A}$ and $\hat{B}$ for the closures of A in $\hat{X}$ and B in $\hat{Y}$, and note that $\hat{A} \setminus A \subseteq \hat{X} \setminus X$ and $\hat{B} \setminus B \subseteq \hat{Y} \setminus Y$. Let us call a continuous function $f\colon \hat{X} \to \hat{Y}$ a *standard map* if $f(X) \subset Y$ and $f(\hat{X} \setminus X) \subset \hat{Y} \setminus Y$. If, in addition, $f(A) \subset B$, we write $f\colon (X, A) \to (Y, B)$. (As before, we refer to X even though the functions live on $\hat{X}$.)

Let us show that every standard map $f : (X, A) \to (Y, B)$ induces a homomorphism $f_* : H_n(X, A) \to H_n(Y, B)$, defined as follows. For a homology class $[c] \in H_n(X, A)$, choose a standard representation $\sum_{i\in I} \lambda_i \sigma_i$ of c and map $[c]$ to the homology class in $H_n(Y, B)$ that contains the n-cycle represented by $\sum_{i\in I} \lambda_i f\sigma_i$. In ordinary singular homology this map is always well defined. To see that it is well defined in our case, note first that f preserves the equivalence of sums, maps boundaries to boundaries, and maps degenerate simplices to degenerate simplices. Hence all that remains to check is that f maps chains to chains. The following lemma implies that it does:

Lemma 6.7. *For every standard map $f : (X, A) \to (Y, B)$, if $(\sigma_i)_{i\in I}$ is an admissible family of n-simplices in $\hat{X}$ (resp. $\hat{A}$), then $(f\sigma_i)_{i\in I}$ is an admissible family of n-simplices in $\hat{Y}$ (resp. $\hat{B}$).*

Proof. As f is standard and $(\sigma_i)_{i\in I}$ is admissible, every $f\sigma_i$ maps the 0-faces of Δ^n to Y. It therefore remains to show that every $y \in Y$ has a neighbourhood that meets the image of $f\sigma_i$ for only finitely many i. Let U be a compact neighbourhood of y in Y, its preimage $f^{-1}(U)$ is a subset of $X = \bigsqcup X_k$ that is closed in $\hat{X}$ as f is continuous. Hence $f^{-1}(U) \cap \hat{X}_k$ is compact for each k. Since $(\sigma_i)_{i\in I}$ is admissible, it contains simplices in only finitely many $\hat{X}_k$, and as the subfamilies of simplices in those $\hat{X}_k$ are admissible, only finitely many σ_i meet $f^{-1}(U)$. Hence U meets the image of $f\sigma_i$ for only finitely many i and hence $(f\sigma_i)_{i\in I}$ is admissible. The analogous claim for A and B follows as $f(A) \subset B$. □

By Lemma 6.7 the map $\sum_{i\in I} \lambda_i \sigma_i \mapsto \sum_{i\in I} \lambda_i f\sigma_i$ defines a homomorphism

$$f_\sharp\colon C_n(X, A) \to C_n(Y, B)$$

with $\partial f_\sharp = f_\sharp \partial$, i.e. $f_\sharp$ is a *chain map.* Thus every standard map $f : (X, A) \to (Y, B)$ induces a homomorphism $f_* : H_n(X, A) \to H_n(Y, B)$. It is easy to see that if $g : (Y, B) \to (Z, C)$ is another standard map we have $(gf)_* = g_* f_*$, and that $\mathbb{1}_* = \mathbb{1}$.

We thus have shown that our homology admits induced homomorphisms if the continuous functions satisfy the natural condition that they map ends to ends and points in X to points in Y. We now show that, subject only to similarly natural constraints, our homology satisfies the axioms for a homology theory.

We will verify the axioms in the order of Section 2.3.

Theorem 6.8 (Homotopy equivalence)**.** *If standard maps $f, g : (X, A) \to (Y, B)$ are homotopic via standard maps $(X, A) \to (Y, B)$ then $f_* = g_*$.*

Proof. Denote by $F = (f_t)_{t\in[0,1]}$ the homotopy between f and g consisting of standard maps $f_t : (X, A) \to (Y, B)$ and satisfying $f_0 = f$ and $f_1 = g$. We first consider the absolute groups $H_n(X)$, $H_n(Y)$.

The main ingredient in the proof of homotopy equivalence for standard singular homology is a decomposition of $\Delta^n \times [0, 1]$ into $(n + 1)$-simplices $D_0, \dots, D_n$ (see eg. [34]). This decomposition works as follows: In $\Delta^n \times [0, 1]$ let $\Delta^n \times \{0\} =: [v_0, \dots, v_n]$ and $\Delta^n \times \{1\} =: [w_0, \dots, w_n]$, and put $D_j := [v_0, \dots, v_j, w_j, \dots, w_n]$. Each D_j is an $(n+1)$-simplex, and hence the natural map between Δ^{n+1} and D_j is a homeomorphism which we denote by τ_j.

In standard singular homology, for an n-chain $z = \sum_{i\in I} \lambda_i\sigma_i$ in X one considers the $(n + 1)$-chain

$$P(z) = \sum_{i\in I} \sum_{j=0}^{n} (-1)^j \lambda_i F \circ (\sigma_i \times \mathbb{1}) \circ \tau_j \tag{6.5}$$

in Y, where $\sigma \times \mathbb{1} \colon \Delta^n \times [0, 1] \to X \times [0, 1]$ is the map $(a, b) \mapsto (\sigma(a), b)$, and then shows that $\partial P(z) + P(\partial z) = g_\sharp(z) - f_\sharp(z)$. If z is an n-cycle, then $g_\sharp(z) - f_\sharp(z) = \partial P(z) + P(\partial z) = \partial P(z)$, thus $g_\sharp(z) - f_\sharp(z)$ is a boundary, which means that $g_\sharp$ and $f_\sharp$ take z to the same homology class and hence $f_*([z]) = g_*([z])$.

In our case, we first have to show that, given an n-chain z in X with representation $\sum_{i\in I} \lambda_i\sigma_i$, the expression $P(z)$ in (6.5) is indeed an $(n+1)$-sum, i.e. that $(F \circ (\sigma_i \times \mathbb{1}) \circ \tau_j)_{i\in I, j\in\{0,\dots,n\}}$ is an admissible family of $(n+1)$-simplices in $\hat{Y}$. Then we have to show that the $c \in C'_{n+1}$ represented by $P(z)$ has a standard representation. If these two claims are true, we will also have $\partial P(z) + P(\partial z) = g_\sharp(z) - f_\sharp(z)$ and hence $f_*([z]) = g_*([z])$.

To show that the family $(F \circ (\sigma_i \times \mathbb{1}) \circ \tau_j)_{i\in I, j\in\{0,\dots,n\}}$ is admissible, note first that, since $(\sigma_i)_{i\in I}$ is an admissible family of simplices in $\hat{X}$, their images meet only finitely many $\hat{X}_k$; let $\hat{X}^-$ be their (compact) union. Now let $y \in Y$ be given, and choose a compact neighbourhood U of y. As $\hat{Y}$ is Hausdorff, U is closed in $\hat{Y}$. Consider the preimage of U under F. As U is closed and F is continuous, this is a closed subset of $\hat{X}^- \times [0, 1]$, and hence compact. Its projection

$$\tilde{U} := \{x \in \hat{X} \mid \exists t \in [0, 1] : F(x, t) \in U\}$$

to $\hat{X}^-$, then, is also compact. Since $U \subset Y$ and each f_t is standard, we have $\tilde{U} \subset X$, so $\tilde{U}$ meets $\operatorname{Im} \sigma_i$ for only finitely many i. And for only those i does U meet the image of any $F \circ (\sigma_i \times \mathbb{1}) \circ \tau_j$, $j \in \{0, \dots, n\}$. Hence $P(z)$ is an $(n + 1)$-sum.

To verify our second claim, let $[z] \in H_n(X)$ be given, and assume without loss of generality that z is good (cf. Proposition 6.4), i.e. it has a representation $\sum_{i\in I}\lambda_i\sigma_i$ such that only finitely many of the σ_i are not degenerate. We show that if σ_i is degenerate then $F\circ(\sigma_i\times\mathbb{1})\circ\tau_j$ is degenerate for each j; from this it follows directly that $P(z)$ as stated in (6.5) is a standard representation of an $(n+1)$-chain in Y.

Suppose that σ_i is degenerate; then there is a compact Hausdorff space X_{σ_i} of dimension at most $n-1$, and continuous maps $\alpha\colon\Delta^n\to X_{\sigma_i}$ and $\beta\colon X_{\sigma_i}\to\hat{X}$ with $\sigma_i=\beta\circ\alpha$. Now let $\gamma:\Delta^{n+1}\to X_{\sigma_i}\times[0,1]$ be the composition of the natural map τ_j from Δ^{n+1} to $D_j\subseteq\Delta^n\times[0,1]$ and the map $\alpha\times\mathbb{1}$ from $\Delta^n\times[0,1]$ to $X_{\sigma_i}\times[0,1]$. Then $F\circ(\sigma_i\times\mathbb{1})\circ\tau_j=(F\circ\beta)\circ\gamma$, so all that remains to show is that $X_{\sigma_i}\times[0,1]$ has dimension at most n. But this is immediate by Lemma 2.3 and the fact that X_{σ_i} has dimension at most $n-1$ while $[0,1]$ has dimension 1.

We thus have $f_*=g_*:H_n(X)\to H_n(Y)$. As P takes sums in A to sums in B, the formula $\partial P+P\partial=g_\sharp-f_\sharp$ remains valid also for relative chains, and thus we also have $f_*=g_*\colon H_n(X,A)\to H_n(Y,B)$. □

Theorem 6.9 (The Long Exact Sequence of a Pair)**.** *There are boundary homomorphisms $\partial\colon H_n(X,A)\to H_{n-1}(A)$ such that*

$$\cdots\xrightarrow{\partial}H_n(A)\xrightarrow{\iota_*}H_n(X)\xrightarrow{\pi_*}H_n(X,A)\xrightarrow{\partial}H_{n-1}(A)\xrightarrow{\iota_*}H_{n-1}(X)\xrightarrow{\pi_*}\cdots$$

is an exact sequence, where ι denotes the inclusion $(A,\emptyset)\to(X,\emptyset)$ and π denotes the inclusion $(X,\emptyset)\to(X,A)$. These boundary homomorphisms are natural*, i.e. given a continuous map $f\colon(X,A)\to(Y,B)$ the diagrams*

$$\begin{array}{ccc} H_n(X,A) & \xrightarrow{\partial} & H_{n-1}(A) \\ \downarrow f_* & & \downarrow f_* \\ H_n(Y,B) & \xrightarrow{\partial} & H_{n-1}(B) \end{array}$$

commute.

Proof. As clearly $\operatorname{Im}\iota_\sharp=\operatorname{Ker}\pi_\sharp$ we have a short exact sequence of chain complexes

$$
\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
\cdots & \xrightarrow{\partial} & C_{n+1}(A) & \xrightarrow{\partial} & C_n(A) & \xrightarrow{\partial} & C_{n-1}(A) & \xrightarrow{\partial} & \cdots \\
 & & \downarrow{\scriptstyle \iota_\sharp} & & \downarrow{\scriptstyle \iota_\sharp} & & \downarrow{\scriptstyle \iota_\sharp} & & \\
\cdots & \xrightarrow{\partial} & C_{n+1}(X) & \xrightarrow{\partial} & C_n(X) & \xrightarrow{\partial} & C_{n-1}(X) & \xrightarrow{\partial} & \cdots \\
 & & \downarrow{\scriptstyle \pi_\sharp} & & \downarrow{\scriptstyle \pi_\sharp} & & \downarrow{\scriptstyle \pi_\sharp} & & \\
\cdots & \xrightarrow{\partial} & C_{n+1}(X,A) & \xrightarrow{\partial} & C_n(X,A) & \xrightarrow{\partial} & C_{n-1}(X,A) & \xrightarrow{\partial} & \cdots \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}
$$

It is a general algebraic fact (see eg. [34]) that for every short exact sequence of chain complexes there exists a natural boundary homomorphism ∂ of the corresponding homology groups giving the desired long exact sequence. □

Theorem 6.10 (Excision)**.** *Let (X, A) be an admissible pair and let B be a closed subset of X such that the interiors* $\operatorname{int}\hat{A}$ *of $\hat{A}$ and* $\operatorname{int}\hat{B}$ *of $\hat{B}$ cover $\hat{X}$. Then the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms $H_n(B, A \cap B) \to H_n(X, A)$.*

To prove Theorem 6.10, we first sketch the proof of excision for ordinary singular homology, and then point out the differences to our case. We start with barycentric subdivision of simplices. The aim is to find a sufficiently fine barycentric subdivision so as to construct a homomorphism from $C_n(X)$ to $C_n(A+B) := C_n(A) + C_n(B) \subseteq C_n(X)$.

Lemma 6.11. *For every n-simplex $[v_0, \dots, v_n]$ there is a finite family of degenerate simplices $\Delta^{n+1} \to [v_0, \dots, v_n]$ such that adding the boundaries of those $(n+1)$-simplices, as well as the n-simplices in the corresponding families of the $(n-1)$-faces of $[v_0, \dots, v_n]$, to the natural map $[v_0, \dots, v_n]$ yields the sum of simplices in its barycentric subdivision (with suitable signs).*

Proof. Induction on n. The lemma is clearly true for $n = 0$. For $n > 0$, let b be the barycentre of $[v_0, \dots, v_n]$. Then $[v_0, \dots, v_n]$ is homologous to $\sum_{k=0}^{n} (-1)^k \Delta_k$ with $\Delta_k := [b, v_0, \dots, \hat{v}_k, \dots, v_n]$, since it differs from this sum by the boundary of the degenerate $(n+1)$-simplex $[b, v_0, \dots, v_n]$. By induction, every $(n-1)$-face of $[v_0, \dots, v_n]$ is homologous via boundaries of degenerate simplices to a sum of the simplices in its barycentric subdivision plus a sum of (degenerate) simplices for each of its $(n-2)$-faces. Hence Δ_k, being the cone over the $(n-1)$-face $[v_0, \dots, \hat{v}_k, \dots, v_n]$ is a corresponding sum of boundaries of degenerate simplices one dimension higher. As each

$(n-2)$-face appears equally often as a face of an $(n-1)$-face of $[v_0, \ldots, v_n]$ with positive and negative sign, so does the sum of (degenerate) simplices belonging to this face. Hence those sums cancel in the sum of all boundaries, which implies that $[v_0, \ldots, v_n]$ is homologous to a sum of the desired type. $\square$

For every singular n-simplex σ, let $T(\sigma)$ be the sum consisting of the compositions of σ and each of the degenerate $(n+1)$-simplices provided by Lemma 6.11 applied to Δ^n, and let $S(\sigma)$ be the sum of restrictions of σ to the simplices in the barycentric subdivision of Δ^n. Then Lemma 6.11 says that (with appropriate choice of the signs in T and S)

$$\partial T(\sigma) = \sigma - T(\partial\sigma) - S(\sigma).$$

Now S and T extend to a chain map $S : C_n(X) \to C_n(X)$, that is, a map with $\partial S = S\partial$ (which follows immediately from the definition of the barycentric subdivision), and a map $T : C_n(X) \to C_{n+1}(X)$ with

$$\partial T + T\partial = \mathbb{1} - S. \tag{6.6}$$

Next, let us define, for every positive integer m, the map $D_m \colon C_n(X) \to C_{n+1}(X)$ like in standard homology, i.e. $D_m := \sum_{0 \le j < m} TS^j$. Note that

$$\partial D_m + D_m\partial = \mathbb{1} - S^m \tag{6.7}$$

by (6.6) and the fact that S is a chain map.

Finally, define maps $D \colon C_n(X) \to C_{n+1}(X)$ and $\rho \colon C_n(X) \to C_n(A+B)$ as follows: For every singular simplex σ, let $m(\sigma)$ be the smallest number m for which every simplex in $S^m(\sigma)$ lives in the interior of $\hat{A}$ or of $\hat{B}$. Now define $D(\sigma) := D_{m(\sigma)}$ and extend linearly to $C_n(X)$. The map ρ is defined by $\rho(\sigma) := S^{m(\sigma)}(\sigma) + D_{m(\sigma)}(\partial\sigma) - D(\partial\sigma)$ and extending linearly. Note that $\rho(\sigma)$ is indeed in $C_n(A+B)$, see [34]. With this notation, we have

$$\partial D + D\partial = \mathbb{1} - \iota\rho, \tag{6.8}$$

where ι is the inclusion $C_n(A+B) \to C_n(X)$. Moreover, we clearly have

$$\rho\iota = \mathbb{1}. \tag{6.9}$$

The relations (6.8) and (6.9) are the main ingredients for the proof of excision in the case of standard singular homology. In the case of our homology, we have to confront three major problems in order to define D and ρ so as to satisfy (6.8) and (6.9):[7] Firstly, these maps will map a singular simplex to a sum of simplices, but the underlying family of this

[7] In order to avoid confusion with the notation of the case of standard homology, we will from now label the maps from standard homology by adding the index $_{\mathrm{fin}}$.

sum need not be admissible as its simplices may map 0-faces to ends. Hence we have to change the maps so that the simplices in their image map 0-faces to X. The second problem is that, while we change the image simplices, we have to ensure that each of them still lives in the interior of $\hat{A}$ or of $\hat{B}$. Hence we are not allowed to change them too much. The third problem will be to guarantee that the image of a chain is a chain, i.e. that it has a standard representation. We shall overcome the first two problems by subdividing the simplices at points that are mapped to X contrary to the barycentres of Δ^n and its faces.

To make this precise, we define the notion of a *σ-pseudo-linear m-simplex*, where σ is a given singular n-simplex. See Figure 6.3 for a low-dimensional example. Let points $w_0, \ldots, w_m, w'_0, \ldots, w'_m \in \Delta^n,\ m \geq 1$, be given such that σ maps each w'_i to X and each w_i with $w_i \neq w'_i$ to $\hat{X} \setminus X$. The σ-pseudo-linear m-simplex with *centre* $[w_0, \ldots, w_m]$ and *antennae* $w_i w'_i,\ 0 \leq i \leq m$, is a singular simplex $\tau : \Delta^m \to [w_0, \ldots, w_m] \cup \bigcup_{i=0}^m w_i w'_i$ defined as follows. Let $v'_0, \ldots, v'_m$ be the vertices of Δ^m and consider the following simplex $[v_0, \ldots, v_m] \subseteq \Delta^m$: Put $v_i := v'_i$ if $w_i = w'_i$ and $v_i := \frac{1}{m+2}\left(2v'_i + \sum_{j \neq i} v'_j\right)$ otherwise. Then map $[v_0, \ldots, v_m]$ to $[w_0, \ldots, w_m]$ by sending v_i to w_i and extending linearly, and map each line $v_i v'_i$ to the line $w_i w'_i$. Call the union of $[v_0, \ldots, v_m]$ and the lines $v_i v'_i$ the *kernel* of Δ^m with respect to the points w_i and w'_i.

For $m = 1$, this already defines the simplex τ. For $m > 1$ and each l-face of $[v'_0, \ldots, v'_m]$ $(1 \leq l < m)$ define τ on the kernel of this face (with respect to the corresponding w_i and w'_i) the same way it is defined on the kernel of Δ^m. Now consider a point x on the boundary of $[v_0, \ldots, v_m]$ and write $x = \sum_{i=0}^m \mu_i v'_i$. For every face D of $[v_0, \ldots, v_m]$ that contains x, spanned by the vertices $v_i,\ i \in I$ say, we say that the projection

$$x_D := \frac{1}{\sum_{i \in I} \mu_i} \sum_{i \in I} \mu_i v'_i$$

of x to the corresponding face of Δ^m is *associated with* x (see Figure 6.3). Note that x_D lies in the kernel of this face of Δ^m and that τ maps x and x_D to the same point in the corresponding face of $[w_0, \ldots, w_m]$. Together with x these points x_D span internally disjoint simplices as follows: For every maximal descending sequence $D_m, D_{m-1}, \ldots, D_k$ of faces that contain x,[8] the points x_{D_j} span an $(m-k)$-simplex, where the number k depends on x but not on the choice of the sequence. For a point on a line $v_i v'_i$ we obtain a set of $(m-1)$-simplices defined in the same way. It is easy to see that these simplices are disjoint for distinct points x, y and that they cover all of Δ^m apart from the interior of its kernel. We can thus define τ on each such simplex as the constant function with image the image of x.

[8]Note that $D_m = [v_0, \ldots, v_m]$ and hence $x_{D_m} = x$. Furthermore, D_k is the (unique)

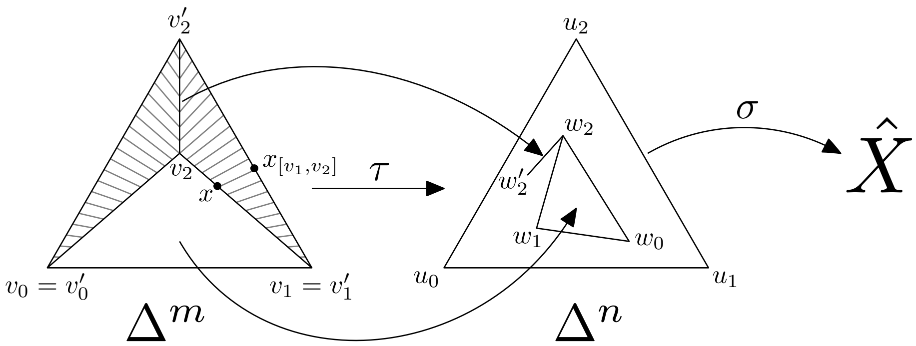

Figure 6.3: A singular 2-simplex σ in $\hat{X}$ and a σ-pseudo-linear 2-simplex τ with centre $[w_0, w_1, w_2]$ and antennae $w_i w_i'$, where $w_0' = w_0$ and $w_1' = w_1$. The points w_0, w_1, and w_2' are mapped to X, the point w_2 is mapped to an end. The simplex τ is constant on each of the grey lines on the left.

The definition of σ-pseudo-linear simplices immediately yields that the boundary of a σ-pseudo-linear $(m+1)$-simplex τ is the sum (with appropriate signs) of the σ-pseudo-linear m-simplices with centres the m-faces of the centre of τ (and the corresponding antennae). This implies

Lemma 6.12. *If an m-simplex $[w_0, \ldots, w_m] \subseteq \Delta^n$ is homologous to a sum of m-simplices, then this remains true if we choose a point w' for every vertex w of those simplices and replace each simplex S by a σ-pseudo-linear simplex with centre S and antennae all lines from a vertex w of S to its w'.* □

The maps D and ρ will map a singular simplex σ to a sum consisting of compositions of σ and σ-pseudo-linear simplices, and correspondingly a chain c to a sum of compositions with σ-pseudo-linear simplices for all simplices σ in a representation of c still to be chosen. In order to choose the antennae of the σ-pseudo-linear simplices, we shall use a subset B' of B defined as follows: For every point in the boundary of $\hat{A}$, choose a compact neighbourhood that is contained in B. This is possible because the boundary of $\hat{A}$ is contained in X—since (X, A) is an admissible pair—and because $\hat{X} \setminus \operatorname{int} \hat{A} \subseteq \operatorname{int} \hat{B}$. Since the boundary of each $\hat{A}_k = \hat{A} \cap \hat{X}_k$ is compact, finitely many such neighbourhoods suffice to cover it. Let B' be the union of $\hat{B} \setminus \hat{A}$ and the neighbourhoods for all k. Write $\hat{B}'$ for the closure of B' in $\hat{X}$. Note that the interiors of $\hat{A}$ and $\hat{B}'$ cover $\hat{X}$ and that the boundary of each $\hat{B}_k' = \hat{B}' \cap \hat{X}_k$ is a compact subset of X_k.

Now consider a singular n-simplex σ. Let b be the barycentre of Δ^n. If $\sigma(b) \in X$, then we set $b' := b$. Otherwise consider the line bu_0, where $\Delta^n =: [u_0, \ldots, u_n]$, see Figure 6.3. As $\hat{X} \setminus X$ is closed and σ is continuous,

face of $[v_0, \ldots, v_m]$ of smallest dimension that contains x.

there is a last point $\tilde{b}$ on this line for which $\sigma(b\tilde{b}) \subseteq \hat{X} \setminus X$. Since the boundaries of $\hat{A}$ and $\hat{B}$ are contained in X, we can find a point b' on the line bu_0 so that

$$\textit{if } \sigma(b) \textit{ lies in the interior of } \hat{A} \textit{ then so does } \sigma(bb') \tag{6.10}$$

and

$$\textit{if } \sigma(b) \textit{ lies in the interior of } \hat{B}' \textit{ then so does } \sigma(bb'). \tag{6.11}$$

Proceed analogously if b is a barycentre of a face of Δ^n. The only difference is that we consider the line bu_j, where j is the smallest index with u_j belonging to that face. It is not hard to see that the points b' can be chosen so that, for singular simplices with a common face, the choices of the points on this face coincide.

We are now ready to define the maps D and ρ. For every singular simplex σ, let $m(\sigma)$ be the smallest number m for which every simplex in $S^m_{\text{fin}}(\sigma)$ lives in the interior of $\hat{A}$ or of $\hat{B}'$. Now for a chain $c \in C_n(X)$ with reduced representation $c = \sum_{i\in I} \lambda_i \sigma_i$, consider the sum

$$\sum_{i\in I} \lambda_i (D_{m(\sigma_i)})_{\text{fin}}(\sigma_i)$$

and define $D(c)$ to be the sum obtained from the above sum by replacing each simplex in each $(D_{m(\sigma_i)})_{\text{fin}}(\sigma_i)$ by the composition of σ_i and a σ_i-pseudo-linear simplex defined as above. (Note that each simplex in $(D_{m(\sigma_i)})_{\text{fin}}(\sigma_i)$ is the concatenation of σ_i and a standard map of a simplex in Δ^n.) For ρ, consider the sum

$$\sum_{i\in I} \Big(S^{m(\sigma_i)}_{\text{fin}}(\sigma_i) + (D_{m(\sigma_i)})_{\text{fin}}(\partial\sigma_i) - D_{\text{fin}}(\partial\sigma_i) \Big)$$

and again replace each simplex in it by the composition of σ_i and a σ_i-pseudo-linear simplex so as to obtain $\rho(c)$.

We need to show that $D(c)$ and $\rho(c)$ are indeed chains, i.e. that they have a standard representation. For both sums, the underlying families of simplices are admissible as the family $(\sigma_i)_{i\in I}$ is and both $D(c)$ and $\rho(c)$ consist of finitely many restrictions of each σ_i (with their 0-faces mapped to X). Now $D(c)$ clearly has a standard representation since each of its simplices can be written as $\sigma_i \circ \tau$ with $\tau\colon \Delta^{n+1} \to \Delta^n$ and thus is degenerate. A standard representation of $\rho(c)$ can be found by combining standard representations of $\partial D(c)$, $D(\partial c)$, and c, according to (6.8). Hence $D(c)$ and $\rho(c)$ are chains.

Proof of Theorem 6.10. Since the inclusion $\iota : C_n(A+B) \hookrightarrow C_n(X)$ maps chains in A to chains in A, it induces a homomorphism $C_n(A+B, A) \to$

$C_n(X,A)$. By (6.8) and (6.9) we obtain that for an n-cycle z in $C_n(A+B,A)$ or in $C_n(X,A)$ the sum $(\rho\circ\iota)(z)-z$, respectively $(\iota\circ\rho)(z)-z$, is a boundary. Hence we have $\rho_*\circ\iota_*=\mathbb{1}$ and $\iota_*\circ\rho_*=\mathbb{1}$ and thus $\iota_*: H_n(A+B,A)\to H_n(X,A)$ is an isomorphism.

We claim that the map $C_n(B)/C_n(A\cap B)\to C_n(A+B)/C_n(A)$ induced by inclusion is an isomorphism and thus induces an isomorphism $H_n(B,A\cap B)\to H_n(A+B,A)$. Then we will have $H_n(B,A\cap B)\simeq H_n(X,A)$ as desired. Indeed, $C_n(A+B)/C_n(A)$ can be obtained by starting with $C_n(B)$ and factoring out those chains whose reduced representation consists of simplices living in $\hat{A}$ (and hence in $\hat{A}\cap\hat{B}$). By Lemma 6.6 and the fact that the boundary of each $\hat{A}_k$ and each $\hat{B}'_k$ is a compact subset of X_k, the latter are precisely the chains in $C_n(A\cap B)$, hence the map $C_n(B)/C_n(A\cap B)\to C_n(A+B)/C_n(A)$ is an isomorphism. □

The last axiom follows directly from the definition.

Theorem 6.13 (Disjoint unions)**.** *For a disjoint union $X=\bigsqcup_\alpha X_\alpha$ (with $\hat{X}$ the disjoint union of all $\hat{X}_\alpha$) with inclusions $\iota_\alpha: X_\alpha\hookrightarrow X$, the direct sum map $\bigoplus_\alpha(\iota_\alpha)_*: \bigoplus_\alpha H_n(X_\alpha,A_\alpha)\to H_n(X,A)$, where $A=\bigsqcup_\alpha A_\alpha$, is an isomorphism.* □

6.5 Cohomology

The cohomology belonging to the homology constructed in Section 6.3 is defined as usual by dualization: Given an admissible pair (X,A) and an abelian group G, the nth group $C^n(X,A;G)$ of cochains is the group of homomorphisms from the nth chain group $C_n(X,A)$ (with integer coefficients) to G. The boundary operators $\partial_n\colon C_n\to C_{n-1}$ dualize to the coboundary operators $\delta^n\colon C^{n-1}\to C^n$ by letting $\delta^n f:=f\circ\partial_n$. In analogy to the notation for homology, the group $\operatorname{Ker}\delta^{n+1}$ of n-cocycles is denoted by Z^n, while B^n denotes the group $\operatorname{Im}\delta^n$ of n-coboundaries. The cohomology groups are then defined by $H^n:=Z^n/B^n$.

We show that this cohomology satisfies the axioms for cohomology. The proofs will be easier than the proofs of the axioms for homology as we can use the results of the last section.

As in the case of homology, we first have to show that we have induced homomorphisms. This follows by dualization of the corresponding fact for homology: Every standard map $f\colon(X,A)\to(Y,B)$ defines a chain map $f_\sharp\colon C_n(X,A)\to C_n(Y,B)$ whose dual is a cochain map $f^\sharp\colon C^n(Y,B;G)\to C^n(X,A;G)$, i.e. which satisfies $\delta f^\sharp=f^\sharp\delta$. Therefore, $f^\sharp$ induces a homomorphism $f^*\colon H^n(Y,B;G)\to H^n(X,A;G)$. Like for homology, if $g\colon(Y,B)\to(Z,C)$ is another standard map, then $(gf)^*=f^*g^*$ and $\mathbb{1}^*=\mathbb{1}$.

Theorem 6.14 (Homotopy equivalence)**.** *If standard maps $f, g : (X, A) \to (Y, B)$ are homotopic via standard maps, then $f^* = g^* \colon H^n(Y, B; G) \to H^n(X, A; G)$.*

Proof. In the proof of Theorem 6.8 we constructed a chain homotopy P between $f_\sharp$ and $g_\sharp$, i.e. a map $P \colon C_n \to C_{n+1}$ with $\partial P + P\partial = g_\sharp - f_\sharp$. The dual of P hence is a map $P^* \colon C^{n+1} \to C^n$ with $P^*\delta + \delta P^* = g^\sharp - f^\sharp$, i.e. a cochain homotopy between $f^\sharp$ and $g^\sharp$. This implies that $f^* = g^*$. □

Theorem 6.15 (The Long Exact Sequence of a Pair)**.** *There are coboundary homomorphisms $\delta \colon H^n(A; G) \to H^{n+1}(X, A; G)$ such that*

$$\cdots \xrightarrow{\delta} H^n(X, A; G) \xrightarrow{\pi^*} H^n(X; G) \xrightarrow{\iota^*} H^n(A; G) \xrightarrow{\delta} H^{n+1}(X, A; G) \xrightarrow{\pi^*} H^{n+1}(X; G) \xrightarrow{\iota^*} \cdots$$

is an exact sequence, where ι denotes the inclusion $(A, \emptyset) \to (X, \emptyset)$ and π denotes the inclusion $(X, \emptyset) \to (X, A)$. These coboundary homomorphisms are natural, *i.e. given a standard map $f : (X, A) \to (Y, B)$ the diagrams*

$$\begin{array}{ccc} H^n(B; G) & \xrightarrow{\delta} & H^{n+1}(Y, B; G) \\ \big\downarrow{\scriptstyle f^*} & & \big\downarrow{\scriptstyle f^*} \\ H^n(A; G) & \xrightarrow{\delta} & H^{n+1}(X, A; G) \end{array}$$

commute.

Proof. The short exact sequence

$$0 \longrightarrow C_n(A) \xrightarrow{\iota_\sharp} C_n(X) \xrightarrow{\pi_\sharp} C_n(X, A) \longrightarrow 0$$

dualizes to

$$0 \longleftarrow C^n(A; G) \xleftarrow{\iota^\sharp} C^n(X; G) \xleftarrow{\pi^\sharp} C^n(X, A; G) \longleftarrow 0\,,$$

where $\iota^\sharp$ and $\pi^\sharp$ denote the cochain maps induced by the inclusions $\iota : (A, \emptyset) \to (X, \emptyset)$ and $\pi : (X, \emptyset) \to (X, A)$. (Note that the cochain maps are the duals of the corresponding chain maps $\iota_\sharp$ and $\pi_\sharp$.) This short sequence is exact: Injectivity of $\pi^\sharp$ is immediate and so is $\operatorname{Ker} \iota^\sharp = \operatorname{Im} \pi^\sharp$. The surjectivity of $\iota^\sharp$ follows from Lemma 6.6, since every given homomorphism $\varphi \colon C_n(A) \to G$ can be extended to a homomorphism $\psi \colon C_n(X) \to G$ by defining $\psi(c)$ as $\varphi(c_A)$, where c_A is the chain in A which c defines by Lemma 6.6. We thus have a short exact sequence of cochain complexes. Like in the homology case, this gives us the desired long exact sequence. □

Theorem 6.16 (Excision)**.** *If (X, A) is an admissible pair and B is a subspace of X such that the interiors of $\hat{A}$ and $\hat{B}$ cover $\hat{X}$, the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms $H^n(X, A; G) \to H^n(B, A \cap B; G)$ for all n.*

Proof. The chain homotopy D and the chain maps ρ and ι from (6.8) and (6.9) induce dual maps D^*, ρ^*, and ι^* that satisfy the dual equations $\iota^*\rho^* = \mathbb{1}$ and $\mathbb{1} - \rho^*\iota^* = D^*\delta + \delta D^*$. Therefore, ι^* and ρ^* induce isomorphisms between the cohomology groups $H^n(X; G)$ and $H^n(A + B; G)$. The inclusion $\pi\colon C_n(A + B) \hookrightarrow C_n(X)$ is the identity on $C_n(A)$ and hence induces an inclusion $C_n(A + B, A) \hookrightarrow C_n(X, A)$ which we also denote by π. Therefore, it induces a homomorphism $\pi^*\colon H^n(X, A; G) \to H^n(A + B, A; G)$. Now by the long exact sequence axiom we have a commutative diagram

$$
\begin{array}{ccc}
H^{n-1}(X; G) & \xrightarrow{\iota^*} & H^{n-1}(A + B; G) \\
\downarrow & & \downarrow \\
H^{n-1}(A; G) & \xrightarrow{\mathbb{1}^*} & H^{n-1}(A; G) \\
\downarrow & & \downarrow \\
H^n(X, A; G) & \xrightarrow{\pi^*} & H^n(A + B, A; G) \\
\downarrow & & \downarrow \\
H^n(X; G) & \xrightarrow{\iota^*} & H^n(A + B; G) \\
\downarrow & & \downarrow \\
H^n(A; G) & \xrightarrow{\mathbb{1}^*} & H^n(A; G)
\end{array}
$$

and since the two maps ι^* as well as the two maps $\mathbb{1}^*$ are isomorphisms, the Five Lemma (Lemma 2.17) shows that π^* is an isomorphism $H^n(X, A; G) \to H^n(A + B, A; G)$. Since the map $C_n(B)/C_n(A \cap B) \hookrightarrow C_n(A + B)/C_n(A)$ induced by inclusion is an isomorphism, this also induces an isomorphism $C^n(A + B, A; G) \to C^n(B, A \cap B; G)$ and hence an isomorphism on cohomology. We thus have an isomorphism $H^n(X, A; G) \to H^n(B, A \cap B; G)$ as desired. □

Theorem 6.17 (Disjoint unions)**.** *If $X = \bigsqcup_\alpha X_\alpha$ with inclusions $\iota_\alpha\colon X_\alpha \hookrightarrow X$, the direct product map $\prod_\alpha (\iota_\alpha)_*\colon H^n(X, A; G) \to \prod_\alpha H^n(X_\alpha, A_\alpha; G)$, where $A = \bigsqcup_\alpha A_\alpha$, is an isomorphism.*

Proof. This follows immediately from the definition of the chains for disjoint unions and the (easy) algebraic fact that the dual of a direct sum is the direct product of the duals. □

6.6 Our homology on graphs

In this section we wind up the analysis of our new homology theory in the case of graphs. We start by computing its homology groups for the case that the space X is a locally finite graph and $\hat{X}$ its Freudenthal compactification. In particular, our homology captures (in dimension 1) precisely the topological cycle space of X:

Theorem 6.18. *Let X be a locally finite connected graph and $\hat{X}$ its Freudenthal compactification. Then*

(i) *mapping a 0-chain in X with finite reduced representation $\sum \lambda_i \sigma_i$ to $\sum \lambda_i$ defines a homomorphism $C_0(X) \to \mathbb{Z}$ which in turn induces an isomorphism $H_0(X) \to \mathbb{Z}$,*

(ii) *the map f defined in Section 6.2.2 defines an isomorphism $H_1(X) \to \mathcal{C}(X)$, and*

(iii) *$H_n(X) = 0$ for every $n > 1$.*

Proof. (i) As no 0-simplex is degenerate and every degenerate 1-simplex is constant and hence a cycle, we obtain that every 0-chain has a finite reduced representation, hence the map defined above is indeed a homomorphism $C_0(X) \to \mathbb{Z}$. Moreover, the boundaries of 1-chains are precisely the boundaries of the finite 1-chains and hence the group $H_0(X)$ is the same as in standard singular homology. In particular, the above map induces an isomorphism to $\mathbb{Z}$.

(ii) By Theorem 6.3 it suffices to prove that our homology coincides with that of Section 6.2 for X and dimension 1. Let us denote the chain groups and homology groups defined in Section 6.2 by $\tilde{C}_n(X)$ and $\tilde{H}_n(X)$, similar for the group of cycles and for that of boundaries.

By Lemma 2.8, every n-simplex in $\hat{X}$, $n \geq 2$, is degenerate. Therefore, every $c \in C'_n(X)$, $n \geq 2$, is good, and so we have $C_n(X) = C'_n(X) = \tilde{C}_n(X)$ for every $n \geq 2$. By Proposition 6.5 every homology class in $H_1(X)$ is represented by a finite cycle, and by Theorem 6.3 the same is true also for $\tilde{H}_1(X)$. Since $C_2(X) = \tilde{C}_2(X)$, we have $B_1(X) = \tilde{B}_1(X)$ and hence $H_1(X) = \tilde{H}_1(X)$.

(iii) Let $n > 1$ and an n-cycle z be given; we show that z is a boundary. To this end, choose an enumeration $e_0, e_1, \ldots$ of the edges of X. Let B_1 be the union of $X - e_0$ and two disjoint closed half-edges of e_0, one at each endvertex. Then the interiors of e_0 and $\hat{B}_1$ cover $\hat{X}$. We may thus apply excision.

Let ρ be the map $C_n(X) \to C_n(e_0 + B_1)$ from (6.8) and (6.9). Then $\rho(z)$ is the sum of a chain in e_0 and a chain in B_1. The boundary of both of those chains is an $(n-1)$-cycle in $e_0 \cap B_1$. As $e_0 \cap B_1$ is the disjoint union of two closed intervals, all its homology groups of dimension at least 1 vanish, hence the boundary of the two chains is also a boundary in $e_0 \cap B_1$. Choose an n-chain in $e_0 \cap B_1$ with the right boundary and subtract it from our two chains in e_0 and B_1 so as to obtain cycles z_0 in e_0 and z'_1 in B_1. Note that $z_0 + z'_1$ is homologous to $z =: z'_0$.

Now repeat the construction with z'_1, e_1, and B_2 the union of $B_1 - e_1$ and two disjoint half-edges of e_1 so as to obtain cycles z_1 in e_1 and z'_2 in B_2. Working through the edges e_i in turn this way, we obtain cycles z_i in e_i and z'_{i+1} in B_{i+1}. Since X is locally finite, for every vertex v there exists an i such that the component C_v of B_i containing v is a closed star around v. In all later B_j, this component remains unchanged, and hence the simplices of z'_i living in C_v are not touched by ρ, i.e. all later z'_j agree on C_v; let z_v be the cycle in C_v formed by those simplices.

Since each z'_i is homologous in B_i to $z_i + z'_{i+1}$ (with $B_0 := X$), the family of all simplices in the $(n+1)$-chains certifying these homologies is locally finite in X: For every $x \in X$ there is an i such that either $x \notin B_i$ or x is contained in a component C_v of B_i (if x is a vertex, then obviously $v = x$). In either case there is a neighbourhood around x that avoids all the $(n+1)$-chains of later steps. Since each $(n+1)$-simplex is degenerate, the family of those simlices is admissible.

Thus, z is homologous to the sum of all z_i and z_v. Since each e_i and each C_v has trivial homology in dimension n, each z_i is a boundary in e_i, of an $(n+1)$-chain c_i say, and so is each z_v in C_v, of an $(n+1)$-chain c_v say. As every point in X has a neighbourhood that meets only finitely many e_i and C_v, the infinite sum $\sum_i c_i + \sum_v c_v$ is an $(n+1)$-chain c in X. By construction, $\partial c = z$.

□

Note that Theorem 6.18 (i) holds for every connected locally compact Hausdorff space. Hence for any $X = \bigsqcup_{k \in K} X_k$ we have $H_0(X) = \bigoplus_{k \in K} \mathbb{Z}$.

Theroem 6.18 implies that $H_1(X)$ is the *strong abelianization* of the fundamental group $\pi_1(\hat{X})$, extending the fact that the first singular homology group is the abelianization of the fundamental group to locally finite graphs with our homology. Given a subgroup H of F_∞, the strong abelianization is defined as follows (see also [10]).

A word w in H is an element of the *big commutator subgroup* of H if its domain S can be partitioned into (possibly infinitely many) intervals S_i, $i \in I$, so that every word $w_i := w \restriction S_i$ lies in H and there is a bijection $\varphi \colon I \to I$ with $w_{\varphi(i)} = w_i^-$ for every $i \in I$. In other words, the big commutator subgroup consists of words that can be split into subwords which in turn can be partitioned into inverse pairs. This definition is a direct translation of the definition of the commutator subgroup of a finite free group, which can be defined the same way but is restricted to finite partitions of S. The strong abelianization of H is now defined as the quotient of H over its big commutator subgroup.

In the case of $H = \pi_1(\hat{X})$ it is easy to see that a word lies in the big commutator subgroup of $\pi_1(\hat{X})$ if and only if every letter apears as often as its inverse. Thus, the strong abelianization H' of $\pi_1(\hat{X})$ is canonically isomorphic to a subgroup of the topological cycle space: For an element h of H' count, for each chord $\vec{e}$ of the topological spanning tree T on which our description of $\pi_1(\hat{X})$ is based, how often the words in h contain the letter $\vec{e}$ and $\overleftarrow{e}$, respectively, and let $\lambda_h(\vec{e})$ be the difference of those two numbers. (This difference is well defined by the definition of the big commutator subgroup.) Then map h to the sum of $\lambda_h(\vec{e})\vec{C}_e$ for all chords of T.

This map clearly is an embedding, and by Lemma 4.5 it is even an isomorphism. By Theorem 6.18 (ii) this implies

Theorem 6.19. *If X is a locally finite connected graph and $\hat{X}$ its Freudenthal compactification, then $H_1(X)$ is the strong abelianization of $\pi_1(\hat{X})$.*

Another aspect of graph (co-)homology mentioned in Section 2.4 is that homology and cohomology coincide in dimension 1 if the coefficients of the homology groups lie in the same abelian group G used in the definition of the cohomology groups. We close this chapter by showing that this does not remain true for our homology.

To see this, let us study the cohomology group $H^1(X;\mathbb{Z})$. The 0-cochains are the functions $X \to \mathbb{Z}$ (recall that every 0-chain is finite and hence a 0-cochain is determined uniquely by its images of the 0-simplices). A 1-cocycle c is a homomorphism from the group of 1-chains to $\mathbb{Z}$ that vanishes on boundaries. Since two 1-chains differ by a boundary if and only if they traverse the same edges (Theorem 6.18 (ii)), a 1-cocycle is determined by its images of (sums of) passes through edges and half-edges. Let T be a normal spanning tree of X. Then one can define a 0-cochain c_0 so that $c + \delta c_0$ maps each sum of passes through edges of T to 0. Therefore, we can characterize the cohomology class $[c]$ by the images of $c + \delta c_0$ on the chords of T. This characterization is easily seen to be an isomorphism ϕ between $H^1(X;\mathbb{Z})$ and the group of all homomorphisms $\mathbb{Z}^{E(G)\setminus E(T)} \to \mathbb{Z}$. If X has at least one non-trivial end, then T has infinitely many chords and $H^1(X;\mathbb{Z})$ is canonically isomorphic to the group of homomorphisms $\mathbb{Z}^\omega \to \mathbb{Z}$.

For each chord e of T, let χ_e be the element of $\mathbb{Z}^{E(G)\setminus E(T)}$ which is 1 at the component corresponding to e and 0 elsewhere. The canonical homomorphism $g\colon H^1(X,\mathbb{Z}) \to \vec{\mathcal{C}}$ is now defined by

$$g([c]) := \sum_{e \in E(G)\setminus E(T)} \phi([c])(\chi_e)\vec{C}_e.$$

For finite graphs, g is known to be an isomorphism. This is not the case for infinite graphs:

Theorem 6.20. *Let X be a locally finite graph with at least one non-trivial end. Then $g\colon H^1(X,\mathbb{Z}) \to \vec{\mathcal{C}}(X)$ is not surjective.*

Proof. Let $z \in \vec{\mathcal{C}}$ be given and write $z = \sum_{e\in E(G)\setminus E(T)} \phi_e \vec{C}_e$. A cohomology class $[c]$ with $g([c]) = z$ has to satisfy $\phi([c])(\chi_e) = \phi_e$. We claim that this is not possible if infinitely many ϕ_e are non-zero; this implies $\operatorname{Im} g = \vec{\mathcal{C}}_{\text{fin}}$ and hence g is not surjective.

To prove the claim, we need an algebraic notation. An abelian torsion-free group G is *slender* if every homomorphism $\mathbb{Z}^\omega \to G$ maps all but finitely many χ_i (with χ_i being 1 in the component i and 0 elsewhere) to 0. For an introduction to (and some examples of) slender groups, see [29, Section 94]. It is immediate that subgroups of slender groups are slender, hence $\mathbb{Z}$, being a subgroup of every torsion-free group, is slender. Thus, we can have $\phi([c])(\chi_e) = \phi_e$ for all chords e if[9] and only if all but finitely many ϕ_e are zero. □

[9]This direction is immediate: Every element φ of $\mathbb{Z}^{E(G)\setminus E(T)}$ has a unique representation as $\sum_{e\in E(G)\setminus E(T)} \lambda_e \chi_e$; let $\phi([c])$ map φ to $\sum_{e\in E(G)\setminus E(T)} \lambda_e \phi_e$. Note that this sum is well-defined since all but finitely many summands are zero.

Bibliography

[1] H. Abels and P. Strantzalos, *Proper transformation groups*, in preparation.

[2] E. Berger and H. Bruhn, *Eulerian edge sets in locally finite graphs*, preprint 2008.

[3] H. Bruhn, *The cycle space of a* 3*-connected locally finite graph is generated by its finite and infinite peripheral circuits*, J. Combin. Theory (Series B) **92** (2004), 235–256.

[4] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. Comput. **15** (2006), 75–90.

[5] H. Bruhn, R. Diestel, and M. Stein, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, J. Graph Theory **50** (2005), 150–161.

[6] H. Bruhn and A. Georgakopoulos, *Bases and closed spaces with infinite sums*, Preprint 2007.

[7] H. Bruhn, S. Kosuch, and M. Win Myint, *Bicycles and left-right tours in locally finite graphs*, Preprint 2007.

[8] H. Bruhn and M. Stein, *MacLane's planarity criterion for locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 225–239.

[9] ______, *On end degrees and infinite circuits in locally finite graphs*, Combinatorica **27** (2007), 269–291.

[10] J.W. Cannon and G.R. Conner, *The combinatorial structure of the Hawaiian Earring group*, Topology Appl. **106** (2000), 225–271.

[11] I.M. Chiswell and T.W. Müller, *A class of groups with canonical* $\mathbb{R}$*-tree action*, in preparation.

[12] G.R. Conner and K. Eda, *Fundamental groups having the whole information of spaces*, Topology Appl. **146–147** (2005), 317–328.

[13] R. Diestel, *The cycle space of an infinite graph*, Comb., Probab. Comput. **14** (2005), 59–79.

[14] ______, *Graph Theory* (3rd edition), Springer-Verlag, 2005, Electronic edition available at:
`http://diestel-graph-theory.com/GrTh.html`.

[15] ______, *Locally finite graphs with ends: a topological approach*, Hamburger Beitr. Math. **340** (2009), see
`http://www.math.uni-hamburg.de/math/research/preprints/hbm.html`.

[16] R. Diestel and D. Kühn, *On infinite cycles I*, Combinatorica **24** (2004), 68–89.

[17] ______, *On infinite cycles II*, Combinatorica **24** (2004), 91–116.

[18] ______, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.

[19] R. Diestel and P. Sprüssel, *The fundamental group of a locally finite graph with ends*, preprint 2008.

[20] ______, *The homology of locally finite graphs with ends*, Combinatorica (to appear).

[21] ______, *On the homology of locally compact spaces with ends*, preprint 2009.

[22] ______, *On the homology of locally finite graphs*, Hamburger Beitr. Math. **277** (2007), see
`http://www.math.uni-hamburg.de/research/hbm.html`.

[23] ______, *The fundamental group of locally finite graphs with ends*, Hamburger Beitr. Math. **305** (2008), see
`http://www.math.uni-hamburg.de/research/hbm.html`.

[24] K. Eda, *Free σ-products and noncommutatively slender groups*, J. Algebra **148** (1991), 243–263.

[25] K. Eda and K. Kawamura, *The fundamental groups of one-dimensional spaces*, Topology Appl. **87** (1998), 163–172.

[26] S. Eilenberg and N. Steenrod, *Foundations of algebraic topology*, Princeton Univ. Press, 1952.

[27] R. Engelking, *Theory of Dimensions, Finite and Infinite*, Sigma series in pure mathematics, vol. 10, Heldermann Verlag, 1995.

[28] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Math. Zeitschr. **33** (1931), 692–713.

[29] L. Fuchs, *Infinite abelian groups II*, New York: Academic Press, 1973.

[30] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, Advances in Mathematics **220** (2009), 670–705.

[31] ———, *Topological circles and Euler tours in locally finite graphs*, Electronic J. Comb. **16** (2009), #R40.

[32] A. Georgakopoulos and P. Sprüssel, *Geodetic topological cycles in locally finite graphs*, Electronic J. Comb. **16** (2009), #144.

[33] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.

[34] A. Hatcher, *Algebraic Topology*, Cambrigde Univ. Press, 2002.

[35] G. Higman, *Unrestricted free products and varieties of topological groups*, J. London Math. Soc. **27** (1952), 73–81.

[36] B. Hughes and A. Ranicki, *Ends of complexes*, Cambrigde Univ. Press, 1996.

[37] B. Krön, *End compactifications in non-locally-finite graphs*, Math. Proc. Cambridge Phil. Soc. **131** (2001), 427–443.

[38] J. Mayer, *Algebraic topology*, Prentice-Hall, 1972.

[39] O. Schreier, *Die Untergruppen der freien Gruppen*, Abh. Math. Sem. Univ. Hamburg **5** (1927), 161–183.

[40] M. Stein, *Arboriticity and tree-packing in locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 302–312.

[41] C. Thomassen and A. Vella, *Graph-like continua, augmenting arcs, and Menger's theorem*, Combinatorica **28** (2008), 595–623, DOI: 10.1007/s00493-008-2342-9.

[42] S. Willard, *General topology*, Dover Publications, 2004.

Name:	Philipp Sprüssel
Titel:	On the homology of infinite graphs with ends
Jahr der Drucklegung:	2010

Zusammenfassung

Der *topologische Zyklenraum* $\mathcal{C}(G)$ eines unendlichen Graphen G wurde 2004 von Diestel und Kühn eingeführt. Mit seiner Hilfe war es möglich, die Zyklenraum-Theoreme, welche den Zyklenraum eines endlichen Graphen mit seinen strukturellen Eigenschaften wie Zusammenhang oder Plättbarkeit in Verbindung setzen, auf lokal endliche Graphen zu erweitern. Die Definition von $\mathcal{C}(G)$ wurde dabei auf kombinatorischem Wege mit Hilfe der Freudenthal-Kompaktifizierung $|G|$ von G gegeben. Ziel der vorliegenden Arbeit ist die Untersuchung, ob sich $\mathcal{C}(G)$ auch über eine Homologietheorie definieren lässt.

Die simpliziale Homologie scheidet als Kandidat hierfür offensichtlich aus, näher untersucht werden die singuläre Homologie (Kapitel 4) und die Čech-Homologie (Kapitel 5) von $|G|$. Das Hauptresultat von Kapitel 4 ist, dass $\mathcal{C}(G)$ ein echter Quotient der ersten singulären Homologiegruppe $H_1(|G|)$ ist: Der natürliche Homomorphismus $H_1(|G|) \to \mathcal{C}(G)$ ist stets surjektiv aber im Allgemeinen nicht injektiv. Hingegen ist die erste Čech-Homologiegruppe kanonisch isomorph zu $\mathcal{C}(G)$, allerdings genügt dieser Gruppenisomorphismus nicht, um die Zyklenraum-Theoreme auf die Čech-Homologie zu erweitern: Aufgrund der Definition von $\check{H}_1(|G|)$ als inverser Limes geht etwa die Information verloren, welche Elemente von $\check{H}_1(|G|)$ Kreisen in $|G|$ entsprechen, was jedoch ein entscheidender Bestandteil vieler Zyklenraum-Theoreme ist.

Als Hilfsmittel für den Beweis der Nicht-Injektivität des Homomorphismus aus Kapitel 4 wird in Kapitel 3 eine kombinatorische Beschreibung der Fundamentalgruppe von $|G|$ entwickelt, ein bereits an sich wünschenswertes Resultat. Zur Beschreibung von $\pi_1(|G|)$ werden unendliche Wörter verwendet, deren Buchstaben die Sehnen eines fest gewählten topologischen Spannbaumes von G sind. Die erhaltene Beschreibung liefert als Korollar eine bekannte Beschreibung der Fundamentalgruppe des *Hawaiischen Ohrrings* durch Higman sowie Cannon und Conner.

Das letzte Ziel dieser Dissertation ist die Konstruktion einer Homologietheorie, basierend auf singulären Simplizes, welche das Konzept des topologischen Zyklenraumes verallgemeinert. In Kapitel 6 gelingt dieses Vorhaben: Es wird eine Homologietheorie für lokal kompakte Hausdorff-Räume definiert, welche für lokal endliche Graphen den topologischen Zyklenraum als erste Homologiegruppe besitzt und zusätzlich die Eilenberg-Steenrod-Axiome für Homologie erfüllt.

Lebenslauf

Philipp Sprüssel, geboren am 18.10.1981 in Hamburg.

Wissenschaftlicher Werdegang

02/2007 – 03/2007	**Israel Institute of Technology, Haifa** Forschungsaufenthalt bei Ron Aharoni
04/2006 – 01/2010	**Universität Hamburg** Promotionsstudium bei Reinhard Diestel
10/2001 – 03/2006	**Universität Hamburg** Diplom Mathematik (Note „sehr gut“)
08/1991 – 06/2000	**Gymnasium Osdorf, Hamburg** Abitur (Note $1,3$)
08/1987 – 06/1991	**Grundschule Mendelssohnstraße, Hamburg**

Stipendien & Anstellungen

04/2007 – 09/2009	**Studienstiftung des deutschen Volkes** Promotionsstipendium
10/2006 – 03/2007	**Universität Hamburg** Wissenschaftlicher Mitarbeiter

Sonstiges

09/2000 – 07/2001	**Ernst & Claere Jung Stiftung, Hamburg** Ersatzdienst im o.g. Seniorenheim