

**HAMBURGER BEITRÄGE
ZUR MATHEMATIK**

Heft 305

**The fundamental group of locally finite
graphs with ends**

R. Diestel, Hamburg

P. Sprüssel, Hamburg

This is an extended version of our paper

The fundamental group of a locally finite graph with ends

It differs from that paper only in that it offers proofs for Lemmas 2, 4, 6, 7, 8, 9 and 20, and a longer example in Section 5.

February 2008
Revised January 2009

The fundamental group of locally finite graphs with ends

Reinhard Diestel and Philipp Sprüssel

January 16, 2009

Abstract

We characterize the fundamental group of a locally finite graph G with ends, by embedding it canonically as a subgroup in the inverse limit of the free groups $\pi_1(G')$ with $G' \subseteq G$ finite. As an intermediate step, we characterize $\pi_1(|G|)$ combinatorially as a group of infinite words.

1 Introduction

The purpose of this paper is to give a combinatorial characterization of the fundamental group of the compact space $|G|$ formed by a locally finite graph G —such as a Cayley graph of a finitely generated group—together with its ends. The space $|G|$, known as the *Freudenthal compactification* of G , is the standard setting in which locally finite graphs are studied from a topological point of view [6]. However, no combinatorial characterization of its fundamental group has so far been known.

When G is finite, $\pi_1(|G|) = \pi_1(G)$ is the free group on the set of (arbitrarily oriented) *chords* of a spanning tree of G , those edges of G that are not edges of the tree. 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 , one whose closure in $|G|$ contains no non-trivial loop.

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

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

where F_∞ is a group formed by ‘reduced’ infinite words of chords of T . 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.) 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 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 $F_\infty \rightarrow F^*$.

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 [18] and Cannon & Conner [3]. Our characterization of $\pi_1(|G|)$ coincides with their combinatorial description of this group when G has only one non-trivial end.

Our motivation for this paper is primarily that we feel that the fundamental group of such a standard space as $|G|$ ought to be understood. Besides, we apply our results in [11] to show that, in contrast to finite graphs, the first singular homology of $|G|$ differs essentially from the (topological) *cycle space* of G , the object commonly used in graph theory to describe the homology of locally finite graphs.

This paper is organized as follows. We begin with a section collecting together the definitions and known background that we need; some elementary general lemmas are also included here. In Section 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 4 we embed $\pi_1(|G|)$ in F_∞ , leaving the proof of the main lemma to Section 5.

2 Terminology and basic facts

In this section we briefly run through any non-standard terminology we use. We also list a few easy lemmas that we shall need, and use freely, later on. Some of these are given with references, the others are proved for the sake of completeness. The reader is encouraged to skim this section for definitions, but to turn to the proofs of the lemmas only as needed.

For graphs we use the terminology of [6], for topology that of Hatcher [17]. Our graphs 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 $\sigma: [0, 1] \rightarrow X$, where X is any topological space. If σ is a loop, it is *based at* the point $\sigma(0) = \sigma(1)$. We write σ^- for the path $s \mapsto \sigma(1 - s)$. The image of an injective path is an *arc* in X , the image of an ‘injective loop’ (a subspace of X homeomorphic to S^1) is a *circle* in X .

Lemma 1 ([16]). *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 .*

All homotopies between paths that we consider are relative to the first and last point of their domain, usually $\{0, 1\}$. We shall often construct homotopies between paths segment by segment. The following lemma enables us to combine certain homotopies defined separately on infinitely many segments.

Lemma 2. Let α, β be paths in a topological space X . Assume that there is a sequence $(a_0, b_0), (a_1, b_1), \dots$ of disjoint subintervals of $[0, 1]$ such that α and β coincide on $[0, 1] \setminus \bigcup_n (a_n, b_n)$, while each segment $\alpha \upharpoonright [a_n, b_n]$ is homotopic in $\alpha([a_n, b_n]) \cup \beta([a_n, b_n])$ to $\beta \upharpoonright [a_n, b_n]$. Then α and β are homotopic.

Proof. Write $D := \bigcup_n (a_n, b_n)$. For every $n \in \mathbb{N}$ let $F^n = (f_t^n)_{t \in [0,1]}$ be a homotopy in $\alpha([a_n, b_n]) \cup \beta([a_n, b_n])$ between $\alpha \upharpoonright [a_n, b_n]$ and $\beta \upharpoonright [a_n, b_n]$. We define the desired homotopy $F = (f_t)_{t \in [0,1]}$ 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 $F((x-\varepsilon, x), (t-\varepsilon, t+\varepsilon)) \subseteq U$; the case $F([x, x+\varepsilon], (t-\varepsilon, t+\varepsilon)) \subseteq U$ is analogous. Suppose first that there is an $\varepsilon_0 > 0$ such that $(x-\varepsilon_0, x) \subseteq D$. As the intervals (a_i, b_i) are disjoint, this means that $(x-\varepsilon, x) \subseteq (a_n, b_n)$ for some n . Then $(x-\varepsilon_0, x) \subseteq [a_n, b_n]$, and hence $F \upharpoonright (x-\varepsilon_0, x) \times [0, 1] = F^n \upharpoonright (x-\varepsilon_0, x) \times [0, 1]$. As F^n is continuous, there is an $\varepsilon < \varepsilon_0$ with $F((x-\varepsilon, x), (t-\varepsilon, t+\varepsilon)) \subseteq U$.

Now suppose that for every $\varepsilon > 0$ the interval $(x-\varepsilon, x)$ meets $[0, 1] \setminus D$. Then also $x \in [0, 1] \setminus D$, 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 . Then $F((x-\varepsilon, x), (t-\varepsilon, t+\varepsilon)) \subseteq U$. 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 $(a_n, b_n) \subseteq (x-\varepsilon, x)$ and hence $F(x', t') = F^n(x', t') \in \alpha([a_n, b_n]) \cup \beta([a_n, b_n]) \subseteq U$. \square

Locally finite CW-complexes can be compactified by adding their *ends*. This compactification can be defined, without reference to the complex, for any connected, locally connected, locally compact topological space X with a countable basis. Very briefly, an *end* of X is an equivalence class of sequences $U_1 \supseteq U_2 \supseteq \dots$ of connected non-empty open sets with compact frontiers and an empty overall intersection of closures, $\bigcap_n \overline{U_n} = \emptyset$, where two such sequences (U_n) and (V_m) are *equivalent* if every U_n contains all sufficiently late V_m and vice versa. This end is said to *live in* each of the sets U_n , and every U_n together with all the ends that live in it is *open* in the space whose point set is the union of X with the set $\Omega(X)$ of its ends and whose topology is generated by these open sets and those of X . This is a compact space, the *Freudenthal compactification* of X [13, 14]. More topological background on this can be found in [1, 2, 19]; for applications to groups see e.g. [1, 22, 23, 26, 28].

For graphs, ends and the Freudenthal compactification are more usually defined combinatorially [6, 15, 20], as follows. Let G be a connected locally finite graph. 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 . It is not hard to see [7] that this combinatorial definition of an end coincides with the topological one given earlier for locally finite complexes. We write $\Omega = \Omega(G)$ for the set of ends of G . The

Fractal compactification of G is now denoted by $|G|$; its topology is generated by the 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. $C(S, \omega) =: C$ is the unique component of $G - S$ in which ω *lives* (i.e., in which every ray of ω has a *tail*, or subray), and $\hat{C}(S, \omega)$ is the union of C with the set of all the ends of G that live in C and the (finitely many) open edges between S and C .¹ 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.

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 $|G|$. The subspaces we shall be considering will be *standard* subspaces of $|G|$: subspaces that contain every edge of which they contain an inner point.

We shall frequently use the following non-trivial lemma.

Lemma 3 ([8]). *For a locally finite graph G , every closed, connected subspace of $|G|$ is arc-connected.*

A *topological tree* in $|G|$ is an arc-connected standard subspace of $|G|$ that contains no circle and is closed in $|G|$. Note that the subgraph that such a space induces in G need not be connected: its arc-connectedness may hinge on the ends it contains. A *chord* of a topological tree T is any edge of G that has both its endvertices in T but does not itself lie in T .

Lemma 4. *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 of them, 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 $\overline{C} \cap (C_x \setminus C) \subseteq \Omega$.

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 3. Let A be an arc in \overline{C}

¹The definition given in [6] is slightly different, but equivalent to the simpler definition given here when G is locally finite. Generalizations are studied in [21, 25].

that a point in C to one in $C_x \setminus C$. As S is finite and $\overline{C} \setminus (S \cup C) \subseteq \Omega$ contains no arc, we can choose A so that $A \cap S = \emptyset$. But then $A \subseteq C_x$, contradicting the definition of C as an arc-component of C_x . \square

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 5. *If a sequence z_0, z_1, \dots 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 4 implies that they lie in arbitrarily small neighbourhoods of z . \square

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.

Let us call a topological tree T in $|G|$ a *topological spanning tree* of G if T contains $V(G)$. Since T is closed in $|G|$, it then also contains $\Omega(G)$. Similarly, a topological tree T in $|G|$ is a *topological spanning tree* of a subspace H of $|G|$ if $T \subseteq H$ and T contains every vertex or end of G that lies in H .

Topological spanning trees are known to exist in all locally finite connected graphs (and in many more [6, 9, 10]). They also exist in all the relevant subspaces. We need a slight technical strengthening of this:

Lemma 6. *Let $T \subseteq H$ be closed, connected standard subspaces of $|G|$. If T is a topological tree, it can be extended to a topological spanning tree of H .*

Proof. As G is locally finite and connected, $|G|$ is a compact Hausdorff space [6]. Let \mathcal{S} be the set of connected 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 [27, p. 203], \mathcal{S} has a minimal element T' by Zorn’s Lemma. By Lemma 3, 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 . \square

Like graph-theoretical trees, topological trees in $|G|$ are contractible. 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 7. *For every point x in a topological tree T in $|G|$ there is a time-injective deformation retraction of T onto x .*

Proof. The space T is metrizable as follows. Choose an enumeration of the edges in T and give the n th 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 using Lemma 5 it is easy to check that it induces the given topology on T . Further, if $z \in yTy'$ for some $y, y' \in T$ we have $d(y, y') = d(y, z) + d(z, y')$. We construct a time-injective homotopy F in T from the identity on T to the map $T \rightarrow \{x\}$; then we have $F(y, t) \in xTy \subseteq X$ for every $y \in X$ and $t \in [0, 1]$, and hence $F \upharpoonright (X \times [0, 1])$ will be the desired time-injective homotopy for X . 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 $d(F(y, t), F(y', t)) \leq d(y, y')$ 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

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

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

$$d(F(y, t), F(y', t)) \leq d(F(y, t), z) + d(z, F(y', t)) \leq d(y, z) + d(z, y') = d(y, y').$$

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

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

□

Given a closed, connected 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 7. Trivial ends in larger subspaces can also be made visible by topological spanning trees:

Lemma 8. *Let T be a topological spanning tree of a closed, connected 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 4, there is an arc-connected neighbourhood $U' \subseteq U$ of ω in H , and we may clearly choose U' as standard subspace. Its closure T' in H lies in $\overline{U} \subseteq T$, is closed in $|G|$, and is therefore arc-connected (Lemma 3). So T' is a topological tree in $|G|$, and contractible by Lemma 7.

Conversely, suppose that ω has a contractible neighbourhood U in H ; this cannot contain a circle. By Lemma 7, the end ω has an open arc-connected 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. \square

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 \bar{e} . Thus, $\{\vec{e}, \bar{e}\} = \{(e, u, v), (e, v, u)\}$, but we cannot generally say which is which. However, from the definition of G as a CW-complex we have a fixed homeomorphism $\theta_e: [0, 1] \rightarrow e$. 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 $\sigma: [0, 1] \rightarrow |G|$ be a path in $|G|$. Given an edge $e = uv$ of G , if $[s, t]$ is a subinterval of $[0, 1]$ such that $\{\sigma(s), \sigma(t)\} = \{u, v\}$ and $\sigma((s, t)) = \vec{e}$, we say that σ *traverses* e on $[s, t]$. It does so *in the direction of* $(\sigma(s), \sigma(t))$, 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 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 $\sigma \upharpoonright [s_n, t_n]$ 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 $\epsilon > 0$, choose m large enough that for all $n > m$ both $|s_n - s| < \epsilon/2$ and $t_n - s_n < \epsilon/2$ (using that the lengths of the intervals $[s_n, t_n]$ converge to 0, which they clearly do); then $|t_n - s| < \epsilon$ 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 u, v has a neighbourhood not containing the other. \square

3 Infinite words, and limits of free groups

In this section and the next, we give a combinatorial description of $\pi_1(|G|)$ —indeed of $\pi_1(H)$ for any closed, connected 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 [12], Cannon & Conner [3], and by Chiswell and Müller [4].

When G is finite, $\pi_1(|G|)$ is the free group F on the set of *chords* (arbitrarily oriented) of any fixed spanning tree, the edges of G that are not edges of the 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 $\bar{e}_i\bar{e}_i$ or $\bar{e}_i\bar{e}_i$. 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 1). The path running along the bottom side of the ladder and back is a null-homotopic loop. Since it traces the chords $\bar{e}_0, \bar{e}_1, \dots$ all the way to ω and then returns the same way, the infinite word $\bar{e}_0\bar{e}_1\dots\bar{e}_1\bar{e}_0$ should reduce to the empty word. But it contains no cancelling pair of letters, such as $\bar{e}_i\bar{e}_i$ or $\bar{e}_i\bar{e}_i$.




Figure 1: The null-homotopic loop $\bar{e}_0\bar{e}_1\dots\omega\dots\bar{e}_1\bar{e}_0$

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 1. *There is a locally finite graph G with a null-homotopic loop σ in $|G|$ whose trace of chords contains no cancelling pair of subsequences, of any order type.*

Proof. 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 $V_0 \cup \dots \cup V_n$. 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 $T_n \subseteq |T|$.

Let σ_0 be the unique (constant) map $[0, 1] \rightarrow T_0$. Assume inductively that $\sigma_n: [0, 1] \rightarrow T_n$ 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 neighbourhoods 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 containing that ray. This map σ is easily seen to be continuous, and by Lemma 7 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 and subdivide the new edges once. 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 $|G|$. \square

Example 1 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 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 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 Sections 4 and 5 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

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

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

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

letters, and say that \bar{e}_i and \bar{e}_i are *inverse* to each other. A *word* in A is a map $w: S \rightarrow A$ from a totally ordered countable set S , the set of *positions* (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 $w: S \rightarrow A$ and $w': S' \rightarrow A$ will be considered the same if there is an order-preserving bijection $\varphi: S \rightarrow S'$ 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

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

and write $w \upharpoonright I$ as shorthand for the restriction $w \upharpoonright w^{-1}(A_I)$. Note that if I is finite then so is the word $w \upharpoonright I$, 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 $w: S \rightarrow A$ 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 \bar{e}_i, \bar{e}_i . We say that w *reduces to* the word $w \upharpoonright (S \setminus \bigcup R)$. If w has no nonempty reduction, we call it *reduced*.

Informally, we 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 $w = \bar{e}_0 \bar{e}_0 \bar{e}_0$ then $r(w) = \bar{e}_0$, but this letter \bar{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 \emptyset . This is the free group with free generators $\bar{e}_0, \bar{e}_1, \dots$ and inverses $\bar{e}_0, \bar{e}_1, \dots$. For finite $I \subseteq \mathbb{N}$, the subgroup

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

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

Consider a word w , finite or infinite, and $I \subseteq \mathbb{N}$. It is easy to check the following:

$$\text{If } R \text{ is a reduction of } w \text{ then } \{\{s, s'\} \in R \mid w(s) \in A_I\}, \quad (1)$$

with the ordering induced from R , is a reduction of $w \upharpoonright I$.

In particular:

$$\text{Any result of first reducing and then restricting a word can} \quad (2)$$

also be obtained by first restricting and then reducing it.

By (2), mapping $w \in F_J$ to $r(w \upharpoonright I) \in F_I$ for $I \subseteq J$ defines an inverse system of homomorphisms $F_J \rightarrow F_I$. Let us write

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

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 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 into a group, our desired group F_∞ .

Existence is immediate:

Lemma 10. *Every word reduces to some reduced word.*

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

To prove uniqueness, we begin with a characterization of the reduced words in terms of reductions of their finite subwords. Let $w: S \rightarrow A$ 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 $w \upharpoonright I$. By (2), a permanent position of $w \upharpoonright I$ is also permanent in $w \upharpoonright J$ 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 $w \upharpoonright I$ but $w \upharpoonright J$ has a letter from $A_J \setminus A_I$ whose position lies between s and s' , so that s and s' are permanent in $w \upharpoonright J$.

Lemma 11. *A word is reduced if and only if all its positions are permanent.*

Proof. The assertion is clear for finite words, so let $w: S \rightarrow A$ 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 $w \upharpoonright I$. By (1), the pairs in R whose elements map to A_I form a reduction of $w \upharpoonright I$, 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 $w \upharpoonright I$ with $I = \{0, \dots, n\}$, the restriction of w to S_n . For any reduction R of w_{n+1} , the set $R^- := \{\{t, t'\} \in R \mid t, t' \in S_n\}$ with the induced ordering is a reduction of w_n , by (1).

Pick $N \in \mathbb{N}$ large enough 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, König's infinity lemma [6] 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 . Inductively, this implies:

$$\text{For all } m \leq n, \text{ we have } R_m = \{\{t, t'\} \in R_n \mid t, t' \in S_m\}, \text{ and} \quad (3)$$

$$\text{the ordering of } R_m \text{ on this set agrees with that induced by } R_n.$$

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).

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). 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), 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 $p' \not\prec p$. 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 . \square

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

Lemma 11 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. By Lemma 11 it suffices to check whether every position becomes permanent in some large enough but finite $w \upharpoonright I$.

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:

Lemma 12. *Two words w, w' can be reduced to the same (abstract) word if and only if $r(w \upharpoonright I) = r(w' \upharpoonright I)$ for every finite $I \subseteq \mathbb{N}$.*

Proof. The forward implication follows easily from (2). Conversely, suppose that $r(w \upharpoonright I) = r(w' \upharpoonright I)$ for every finite $I \subseteq \mathbb{N}$. By Lemma 10, 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 $\varphi: S \rightarrow S'$ 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

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

Hence for every possible domain $S_I \subseteq S$ of $r(v \upharpoonright I)$ and every possible domain $S'_I \subseteq S'$ of $r(v' \upharpoonright I)$ there exists an order isomorphism $S_I \rightarrow S'_I$ that commutes with v and v' . For every I , there are only finitely many such maps $S_I \rightarrow S'_I$, since there are only finitely many such sets S_I and S'_I . And for $I \subseteq J$, every such map $S_J \rightarrow S'_J$ induces such a map $S_I \rightarrow S'_I$ with $S_I \subseteq S_J$ and $S'_I \subseteq S'_J$, by (2). Hence by the infinity lemma [6] there exists a sequence $\varphi_0 \subseteq \varphi_1 \subseteq \dots$ of such maps $\varphi_n: S_{\{0, \dots, n\}} \rightarrow S'_{\{0, \dots, n\}}$, whose union φ maps all of S onto S' , since by Lemma 11 every position of v and of v' is permanent. \square

With Lemma 12 we are now able to prove:

Lemma 13. *Every word reduces to a unique reduced word.*

Proof. By Lemma 10, 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 12, we have

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

for every finite $I \subseteq \mathbb{N}$. By the nontrivial direction of Lemma 12, 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 $w' = w''$. \square

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 \emptyset , and inverses w^- of $w: S \rightarrow A$ defined as the map on the same S , but with the inverse ordering, satisfying $\{w(s), w^-(s)\} = \{\bar{e}_i, \bar{e}_i\}$ 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 13.

In the notation of Cannon & Conner [3] we have $F_\infty = BF(\aleph_0)$. Eda defined infinite words in a more general setting [12], in his notation $\times_{n \in \mathbb{N}} \mathbb{Z} = \times_{n \in \mathbb{N}}^\sigma \mathbb{Z}$ equals F_∞ .

As indicated earlier, we claim that F_∞ embeds canonically in the inverse limit F^* of the groups F_I . By (2), the maps $h_I: w \mapsto r(w \upharpoonright I)$ are homomorphisms $F_\infty \rightarrow F_I$ that commute with the homomorphisms $F_J \rightarrow F_I$ from the inverse system, so they define a homomorphism

$$h: F_\infty \rightarrow F^*$$

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

$$r(w \upharpoonright I) = h_I(w) = \pi_I(h(w)) = \pi_I(\text{id}) = \emptyset,$$

where id 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 11, this means that $w = 0$. Thus, h is a group embedding of F_∞ in F^* , as claimed.

We remark that h is never surjective. Indeed, while every letter occurs only finitely often in a given word, there are elements of F^* whose projections to the F_I contain some fixed letter unboundedly often; such an element will not lie in the image of h . (For example, the words $\bar{e}_1\bar{e}_0\bar{e}_1\bar{e}_0\bar{e}_2\bar{e}_0\bar{e}_2\bar{e}_0\dots\bar{e}_i\bar{e}_0\bar{e}_i\bar{e}_0\in F_I$ for $I = \{1, \dots, i\}$ define such an element of F^* .) 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 $\bar{e} \in A$ there exists a $k \in \mathbb{N}$ such that $|w_I^{-1}(\bar{e})| \leq k$ for all I .

Theorem 15 (ii) below summarizes what we have shown so far.

4 Embedding $\pi_1(H)$ in F_∞

Let G be a locally finite connected graph. Let H be a closed, connected 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—apply Lemma 7 to the maximal topological subtrees not meeting the interior of an arc between two chords—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

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

Let

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

be the group of infinite reduced words with letters in A , as defined in Section 3.

Unless otherwise mentioned, the endpoints of all paths considered from now on will be vertices or ends, and any homotopies between paths will be relative to $\{0, 1\}$. 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 $\pi_1(H) \rightarrow F_\infty$.

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 12 it suffices to show that $r(w_\alpha \upharpoonright I) = r(w_\beta \upharpoonright I)$ for every finite $I \subseteq \mathbb{N}$. Consider the space obtained from H by attaching a 2-cell to H for every $j \notin I$, by an injective attachment map from the boundary of the 2-cell 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 Lemma 7 to the wedge sum W_I of $|I|$ circles, one for every e_i . Composing α and β with the map $H \rightarrow W_I$ from this homotopy equivalence yields homotopic loops α' and β' in W_I , whose traces in F_I are $w_{\alpha'} = w_\alpha \upharpoonright I$ and $w_{\beta'} = w_\beta \upharpoonright I$. Since $\langle \gamma \rangle \mapsto r(w_\gamma)$ is known to be well defined for wedge sums of finitely many circles, we deduce that

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

This completes the proof that $\langle \alpha \rangle \mapsto r(w_\alpha)$ is well defined. By (2), it is a homomorphism. For injectivity, we shall prove in Section 5 the following extension to paths that need not be loops:

Lemma 14. *Paths σ, τ in H with the same endpoints are homotopic in H if (and only if) their traces reduce to the same word.*

We remark that the map $\langle \alpha \rangle \mapsto r(w_\alpha)$ will not normally be surjective. For example, a sequence $\bar{e}_0, \bar{e}_1, \dots$ of distinct chords will always be a reduced word, but no loop in $|G|$ can pass through these chords in order if they do not converge to an end. Hence if two ends are non-trivial in H , then by Lemma 8 there is a non-convergent sequence $\bar{e}_0, \bar{e}_1, \dots$ of chords of T in H (picked alternately from smaller and smaller neighbourhoods of the two ends), which forms a reduced word in $F_\infty(\bar{A})$ 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 $w' := w \upharpoonright S'$ of a word $w: S \rightarrow A$ *monotonic* if S' is infinite and can be written as $S' = \{s_0, s_1, \dots\}$ so that either $s_0 < s_1 < \dots$ or $s_0 > s_1 > \dots$. Let us say that w' *converges* (in $|G|$) if there exists an end to which every sequence x_0, x_1, \dots 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 15. *Let G be a locally finite connected graph, and let H be a closed, connected standard subspace of $|G|$. Let T be a topological spanning tree of H , and let e_0, e_1, \dots 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 $\{\bar{e}_0, \bar{e}_1, \dots\} \cup \{\bar{e}_0, \bar{e}_1, \dots\}$. Its image consists of those reduced words whose monotonic subwords all converge in $|G|$.*
- (ii) *The homomorphisms $w \mapsto r(w \upharpoonright I)$ 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 $r(w \upharpoonright I)$ 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 14 (which will be proved in Section 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 $w: S \rightarrow A$ be such a word. Enumerate S as s_0, s_1, \dots . 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 $\alpha: [0, 1] \rightarrow H$.

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 $s_0 < s_2 < \dots \mid \dots < s_3 < s_1$ (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 P . If not, there are uncountably many 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, 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 $S^-(P_0) \subsetneq S^-(P_1) \subsetneq \dots$ or $S^+(P_0) \subsetneq S^+(P_1) \subsetneq \dots$. We assume that $S^-(P_0) \subsetneq S^-(P_1) \subsetneq \dots$, 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 $w \upharpoonright \{s_0^-, s_0^+, s_1^-, s_1^+, \dots\}$ 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 $a \leq b$ of each I , and then extend it continuously to map \bar{I} onto the arc $\alpha(a)T\alpha(b)$. To help with our subsequent proof that α is continuous, let us make sure when we choose $\alpha(a)$ and $\alpha(b)$ that they satisfy the following 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 . (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 $0 < a \leq b < 1$. 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. In both these cases, (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 coinital sequence $\dots < s_{n_1} < s_{n_0}$, 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 coinital 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 (4). As $\alpha(a)$, if $a \neq b$, we choose any vertex or end in H , satisfying (4) trivially. In the case of $b = 1$, we proceed analogously.

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

It remains to check that α is continuous at boundary points x of intervals I_n or of connecting intervals. Suppose first that x is the boundary point of a connecting interval. Let x_0, x_1, \dots 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 (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 $\alpha \upharpoonright I$. Disregarding these x_j , we may thus assume that each I^i contains only finitely many x_j .

Let us show that the sequence $\alpha(y_0), \alpha(z_0), \alpha(y_1), \alpha(z_1), \dots$ converges to $\alpha(x)$. Let a neighbourhood $U := \hat{C}(S, \alpha(x))$ of $\alpha(x)$ be given, and let ε be as provided by (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 (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 , we obtain $\alpha(y_i) \in U$ as desired. We now apply Lemma 5 to the sequence $\alpha(y_0), \alpha(z_0), \alpha(y_1), \alpha(z_1), \dots$. By the lemma, the entire arcs $\alpha(I^i) = \alpha(y_i)T\alpha(z_i)$ converge to $\alpha(x)$. In particular, $\alpha(x_i) \rightarrow \alpha(x)$ 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. \square

Corollary 16. *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 6, T extends to a topological spanning tree T' of H' . Since every chord of T is also a chord of T' , Theorem 15 (i) implies the assertion. \square

Let us call a standard subspace H of $|G|$ *non-trivial* if it contains an end (of G) that is non-trivial in H . The fundamental groups of all such spaces contain and are contained in the abstract group F_∞ defined in Section 3; recall that this group is independent of G , as long as $|G|$ itself is non-trivial. Indeed:

Corollary 17. *For every non-trivial standard subspace H of any $|G|$ there are subgroup embeddings $F_\infty \rightarrow \pi_1(H) \rightarrow F_\infty$.*

Proof. Theorem 15 (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 8). 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 16. Since T is a topological spanning tree of H' all whose (infinite subwords of words of) chords converge in $|G|$, Theorem 15 (i) implies that $\pi_1(H')$ is isomorphic to F_∞ . \square

Corollary 18. *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 an end, which would be non-trivial in H by Lemma 8. By Lemma 7, 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 17, $\pi_1(H)$ contains F_∞ as a subgroup, so by the Nielsen-Schreier theorem [24] it suffices to prove that F_∞ is not free. As pointed out in [3], this was shown by Higman [18]. \square

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 15 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 . It does not answer the group-theoretical question of how interesting or varied are the groups occurring as $\pi_1(|G|)$ or $\pi_1(H)$.

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 occurring as $\pi_1(H)$ for non-trivial H . However, this is far from the truth:

Corollary 19. *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 Corollary 19 we need a lemma about F_∞ . By Theorem 4.1 of Conner and Eda [5], the image of any homomorphism $F_\infty \rightarrow A * B$ is contained in either $A' * B$ or $A * B'$ for finitely generated subgroups A' of A and B' of B . Hence,

Lemma 20. *If groups A, B are not finitely generated then there is no epimorphism $F_\infty \rightarrow A * B$. \square*

Proof of Corollary 19. 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 $\pi_1(H) \rightarrow F_\infty$ from Theorem 15 (i) is therefore surjective.

Now suppose that $\pi_1(H)$ is isomorphic to F_∞ . Then H is non-trivial (see the proof of Corollary 18), 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 $H_1 \cap H_2 = \{v_S\}$, the Seifert - van Kampen theorem yields that $\pi_1(H') \simeq \pi_1(H_1) * \pi_1(H_2)$.

We now define an epimorphism $f: \pi_1(H) \rightarrow \pi_1(H')$; as $\pi_1(H) \simeq F_\infty$ by assumption, this will induce an epimorphism $F_\infty \rightarrow \pi_1(H_1) * \pi_1(H_2)$ contradicting Lemma 20. By Theorem 15 (i), every element a of $\pi_1(H)$ corresponds to

a reduced word w_a in $F_\infty(\{\bar{e}_0, \bar{e}_1, \dots\})$ all of whose monotonic subwords converge in $|G|$. Consider the word $r(w_a \upharpoonright \{n, n+1, \dots\})$. 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 (2) and Theorem 15 (i) the map $w_a \mapsto r(w_a \upharpoonright \{n, n+1, \dots\})$ induces a homomorphism $f: \pi_1(H) \rightarrow \pi_1(H')$. For the same reason f is surjective: by Theorem 15 (i), every element of $\pi_1(H')$ corresponds to a reduced word w in $\{\bar{e}_n, \bar{e}_{n+1}, \dots\} \cup \{\bar{e}_n, \bar{e}_{n+1}, \dots\}$ 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 15 (i). \square

5 Proof of Lemma 14

We conclude our proof of Theorem 15 with the proof of Lemma 14. In this proof we shall need another lemma:

Lemma 21. *Let σ be a path in H and let $w_\sigma: S \rightarrow A$ 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 $w_\tau = w_\sigma \upharpoonright S'$. Moreover, τ can be chosen so that $\tau \upharpoonright [a, b] = \sigma \upharpoonright [a, b]$ 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 15(i): As w_σ is the trace of a path, all its monotonic subwords converge in H . Hence also all monotonic subwords of $w_\sigma \upharpoonright S'$ 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 explicitly.

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 $\sigma(a)$ is an end, and we let $\tau(a) := \sigma(a)$. Clearly $w_\tau = w_\sigma \upharpoonright S'$ and $\tau \upharpoonright [a, b] = \sigma \upharpoonright [a, b]$ for each $[a, b] \in S'$; it remains to show that τ is 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 easy to prove, using Lemma 5 and the fact that σ is continuous. \square

Proof of Lemma 14. The proof that $\alpha \sim \beta$ implies $r(w_\alpha) = r(w_\beta)$ was shown already 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 five 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 (\bar{e}, \bar{e}) by a local homotopy $f_t \rightarrow f_{t'}$ that retracts a small segment of f_t running through \bar{e} and back through \bar{e} without traversing any other chords. These pairs of passes have to be nested in the right way, and finding this pairing will be the second part of our proof. In the third part we determine a (time) order in which to cancel those nested pairs. This order 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 fourth step we smooth out those discontinuities by inserting further local homotopies $f_t \rightarrow f_{t'}$ at those countably many times t . In part five of the proof we finally show that the homotopy F thus defined is indeed continuous.

We now embark on part one of the proof, simplifying 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 2 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 $\text{Im } \alpha$ to a straight path α' that has the same passes as α . (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 $\alpha \upharpoonright [a, b]$ and $\alpha' \upharpoonright [a, b]$ are a pass of their respective path or neither is, and if both are then $\alpha \upharpoonright [a, b] = \alpha' \upharpoonright [a, b]$.

By (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 themselves straight.

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 $\text{Im } \alpha \cup \text{Im } \beta$. If β is constant, this homotopy can be chosen time-injective. (6)

To prove (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 $\gamma: [a, b] \rightarrow T$ of α or β that are maximal with $\gamma((a, b)) \subseteq T \setminus xTy$, 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 $\text{Im } \gamma$ (Lemma 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 7, there

is a time-injective homotopy in $\text{Im } \gamma$ from γ to the constant map $[a, b] \rightarrow \{\gamma(a)\}$ ($= \{\gamma(b)\}$). Combining all these homotopies, one for every γ , we obtain time-injective homotopies (in $\text{Im } \alpha$ and $\text{Im } \beta$) 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 2 to apply (6) to segments between passes, we obtain the following generalization:

If paths α and β in H with identical endpoints have precisely the same passes (i.e., they agree on each domain of a pass of α or β), then there is a homotopy between them in $\text{Im } \alpha \cup \text{Im } \beta$. (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 (7), we may assume that

If $(a, b) \subseteq [0, 1]$ 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 $[a, b]$. (8)

This completes part one of the proof.

In order to construct our desired homotopy F between σ and τ , we have to use 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 \bar{e}_i, \bar{e}_i . The positions of w_σ , in this case, are the domains of the passes of σ . Let us write

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

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. Let $[a, a^-]$ be the domain of π and $[b^-, b]$ that of ρ , and assume that $a < a^- \leq b^- < b$. Then $\pi(a) = \rho(b) =: z$ and $\pi(a^-) = \rho(b^-) =: z^-$ are the two vertices of e . If $\alpha \upharpoonright [a^-, b^-]$ traverses no chord, and if β is the path obtained from α by replacing its segment $\alpha \upharpoonright [a, b]$ with the constant map $[a, b] \rightarrow \{z\}$, we say that β is obtained from α by *cancelling the pair p of passes*. Since a, a^-, b, b^-, z, 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 $\sigma(a(p)) = \sigma(b(p)) = z(p)$ and $\sigma(a^-(p)) = \sigma(b^-(p)) = z^-(p)$. These points are vertices and hence lie in T . (9)

This completes the second part of our proof.

In the third part, we now wish to determine the order in which our homotopy F has to 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 $(a(p), b(p))$ and $(a(p'), b(p'))$ are either disjoint or nested; accordingly, we call p and p' *parallel* or *nested*. If p and p' are nested and $[a(p), b(p)]$ contains $[a(p'), b(p')]$, we say that p *surrounds* p' and write $p \geq p'$. This is 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. (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 (arbitrary) linear extension of \leq defined by the reduction R .

To define \preceq , we start by enumerating the elements of \mathcal{P} arbitrarily. Then, recursively for $i = 0, 1, \dots$, let M_i be a maximal chain in $\mathcal{P} \setminus (M_0 \cup \dots \cup M_{i-1})$ 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 (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 . (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 $[s(p), t(p)] \subseteq [0, 1]$ in which F can cancel p . These intervals will reflect \preceq in that

$$t(p) < s(p') \text{ whenever } p \prec p'; \tag{12}$$

in particular, they will be disjoint for different p . While t runs through $[s(p), t(p)]$, the path f_t will change only on $[a(p), b(p)]$, 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 $\alpha \upharpoonright [a, b]$ with the constant map $[a, b] \rightarrow \{z\}$. We say that a homotopy from α to β *retracts* $\alpha \upharpoonright [a, b]$ to z if it is relative to $[0, a] \cup [b, 1]$, time-injective, and maps $[a, b] \times [0, 1]$ to $\alpha([a, b])$. If $\alpha([a, b]) \subseteq X \subseteq H$, we may also say that this retraction is performed *in* X .

Our paths f_t will satisfy the following assertions:

$$\begin{aligned} & \text{The passes of } f_t \text{ are also passes of } \sigma; \text{ thus, } f_t \upharpoonright [c, d] = \sigma \upharpoonright [c, d] \\ & \text{for every } t \in [0, 1] \text{ and every domain } [c, d] \text{ of a pass of } f_t. \end{aligned} \quad (13)$$

$$\begin{aligned} & \text{For every } p \in \mathcal{P}, \text{ the passes of } f_{s(p)} \text{ are exactly those passes of} \\ & \sigma \text{ that are contained in pairs } p' \succeq p; \text{ in particular, } p \text{ is a pair of} \\ & \text{passes of } f_{s(p)}. \end{aligned} \quad (14)$$

$$\begin{aligned} & \text{For every } p \in \mathcal{P}, \text{ the path } f_{t(p)} \text{ is obtained from } f_{s(p)} \text{ by can-} \\ & \text{celling the pair } p. \text{ This is achieved by a homotopy } (f_t)_{t \in [s(p), t(p)]} \\ & \text{that first retracts } f_{s(p)} \upharpoonright [a^-(p), b^-(p)] \text{ to } z^-(p) \text{ in } T \text{ and then} \\ & \text{retracts the resulting path } f_t \upharpoonright [a(p), b(p)] \text{ to } z(p) \text{ in } e(p). \end{aligned} \quad (15)$$

The first part of the homotopy in (15) will be obtained by applying (6), with $\alpha = f_{s(p)} \upharpoonright [a^-(p), b^-(p)]$ and $\beta: [a^-(p), b^-(p)] \rightarrow \{z^-(p)\}$. This will turn $f_{s(p)}$ into a path f_t mapping $[a(p), b(p)]$ onto the chord $e(p)$, to which the second part of the homotopy is then applied.

Let us note for later use:

$$\text{The path } f_{t(p)} \text{ maps } [a(p), b(p)] \text{ to the vertex } z(p). \quad (16)$$

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

In the fourth part, we now turn to the reason why we have not chosen the intervals $[s(p), t(p)]$ 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 . 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

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

Since $P_{i(r)}^+$ is countable, it has an (infinite) coinital sequence $p_0^+ > p_1^+ > \dots$. Then $\lim a(p_n^+) = a^+$. As σ is continuous, the vertices $z(p_n^+) = \sigma(a(p_n^+))$ converge in H to $\sigma(a^+)$. As only finitely many of the vertices $z(p_n^+)$ can coincide

(Lemma 9), $\sigma(a^+)$ 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 $x \in (a(q), b(q))$ also for some $q \in P^-$. Then

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

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

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

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

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

The reader may wish to construct some examples of such σ at this point.²

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 (12)–(15), but apart from this arbitrarily. Since $p_0^+ \succ p_1^+ \succ \dots$ we then have $t_0^+ > t_1^+ > \dots$, where $t_n^+ := t(p_n^+)$ is the time at which F has just cancelled the pair p_n^+ . Let

$$t^+ := \lim_n t_n^+.$$

For every $x \in [a(p_n^+), b(p_n^+)]$, statement (16) yields $f_{t_n^+}(x) = z(p_n^+) = \sigma(a(p_n^+))$. Every $x \in [a^+, b^+]$ satisfies this for all n , so for such x the continuity of F and σ imply

$$f_{t^+}(x) = \lim_n f_{t_n^+}(x) = \lim_n \sigma(a(p_n^+)) = \sigma(a^+) = z(r). \quad (17)$$

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

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

As $q \prec p$ for all $q \in P^-$ and $p \in P^+$, we have $t_x^- \leq t^+$ by (12). As earlier, the fact that $x \in [a(q_n), b(q_n)]$ for all n implies

$$f_{t_x^-}(x) = \lim_n f_{t(q_n)}(x) \stackrel{(16)}{=} \lim_n z(q_n) = \lim_n \sigma(a(q_n)) = \sigma(a_x^-) =: z_x. \quad (18)$$

²Here is one example: Suppose that ω is a non-trivial end of H , and let e'_0, e'_1, \dots 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 $\bar{e}'_2, \bar{e}'_3, \dots$ and reaches ω at time $1/3$. For $x \in [2/3, 1]$, we let $\sigma(x) := \sigma(1-x)$. Thus in $[2/3, 1]$, σ returns from ω to v , traversing $\dots, \bar{e}'_3, \bar{e}'_2$. Between time $1/3$ and $2/3$ it traverses $\bar{e}'_0, \bar{e}'_1, \bar{e}'_1, \bar{e}'_0$. Let the domains of these passes be $[0.35, 0.36]$, $[0.4, 0.41]$, $[0.5, 0.51]$, and $[0.64, 0.65]$. In \mathcal{P} , the unique pass through \bar{e}'_i forms a pair with the unique pass through \bar{e}'_i . In the order \prec , the pairs p_0, p_1 that consist of passes through e'_0 and through e'_1 , respectively, are smaller than all 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 $(1/3, 2/3)$ outside $(0.35, 0.65)$ are not critical with respect to r .

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 (17) and (18), we shall 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 one of these 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

$$P_{i(r)}^- \neq \emptyset \text{ for every bad partition } r = (P^-, P^+), \text{ indeed for every} \quad (19)$$

critical partition that has a critical point.

To prove (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 $q < p$. By the maximality of M_i , this means that $q \not\leq p$ for some $p \in M_i$. Since $q \leq p$ for all $p \in P_i^+$ (as shown earlier), it follows that $P_i^- = M_i \setminus P_i^+ \neq \emptyset$, as claimed.

By (19), the nested union

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

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 $(a(p), b(p))$ 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.) 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 $x \in (a(q), b(q)) \subseteq (a(p), b(p))$ 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 (9), we obtain:

$$\text{For every } r \in \mathcal{R} \text{ we have } \sigma(a(r)) = \sigma(b(r)) \quad (= z(r) \in \Omega(G) \subseteq T). \quad (20)$$

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 $[s(q), t(q)]$ for all $q \in \mathcal{P} \cup \mathcal{R}'$, extending (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 (17), we shall have to have $f_{t_r^+}$ map all of $[a(r), b(r)]$ to $\{z(r)\}$. For bad r , this will require the insertion of a local homotopy $f_{s(r)} \rightarrow f_{t(r)}$ analogous to (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 $[a(q), b(q)] \subseteq [a(q'), b(q')]$. 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:

$$\text{Given } q, q' \in \mathcal{P} \cup \mathcal{R} \text{ with } q \preceq q', \text{ either } q \leq q' \text{ or the intervals } (a(q), b(q)) \text{ and } (a(q'), b(q')) \text{ are disjoint.} \quad (21)$$

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 (19):

$$\text{For every } r \in \mathcal{R} \text{ that has a critical point, in particular for every } r \in \mathcal{R}', \text{ there is a pair } p \in \mathcal{P} \text{ such that } p \leq r. \text{ Given any } r' \in \mathcal{R} \text{ with } r' \prec r, \text{ this } p \text{ can be chosen so that } r' \prec p < r. \quad (22)$$

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

A partition $(P^-, P^+) \in \mathcal{R}$ precedes all pairs in P^+ and succeeds all pairs in P^- . Two partitions $r = (P^-, P^+)$ and $\tilde{r} = (\tilde{P}^-, \tilde{P}^+)$ satisfy $r \preceq \tilde{r}$ if and only if $\tilde{P}^+ \subseteq P^+$. (23)

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

$$t(q) < s(q') \text{ whenever } q \prec q'. \quad (24)$$

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

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

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

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

For every $x \in [0, 1]$, the set

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

is a \leq -chain, by (21). If $Q_x^t \neq \emptyset$, we set

$$(a_x^t, b_x^t) := \bigcup_{q \in Q_x^t} (a(q), b(q))$$

and define $f_t(x) := \sigma(a_x^t) = \sigma(b_x^t) \in T$. (Recall (9) and (20).) If $Q_x^t = \emptyset$, we call x *unchanged at time t* , define $f_t(x) := \sigma(x)$, and put $a_x^t := x =: b_x^t$. Thus,

$$\text{for all } t \notin C \text{ and all } x: f_t(x) = \sigma(a_x^t) = \sigma(b_x^t). \quad (25)$$

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:

For all $t \notin C$ and all x with $Q_x^t \neq \emptyset$, the function f_t is constant on (a_x^t, b_x^t) with value $f_t(x) \in T$. (26)

To prove (26), pick $y \in (a_x^t, b_x^t)$. We show that $Q_y^t \neq \emptyset$ and $(a_y^t, b_y^t) = (a_x^t, b_x^t)$; 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 $(a(q), b(q))$ contains both x and y , giving $q \in Q_x^t \cap Q_y^t$. As Q_x^t is a \leq -chain, we have $q \leq q'$ for all large enough $q' \in Q_x^t$. Since $y \in (a(q), b(q)) \subseteq (a(q'), b(q'))$ implies $q' \in Q_y^t$, all large enough $q' \in Q_x^t$ lie in Q_y^t , giving $(a_x^t, b_x^t) \subseteq (a_y^t, b_y^t)$. Likewise, $(a_y^t, b_y^t) \subseteq (a_x^t, b_x^t)$ and hence $(a_x^t, b_x^t) = (a_y^t, b_y^t)$. This completes the proof of (26), and with it the proof that the f_t defined so far are continuous.

We have just shown that $(a_y^t, b_y^t) = (a_x^t, b_x^t)$ for all x and y with $y \in (a_x^t, b_x^t)$. Therefore, for any x, y the intervals (a_x^t, b_x^t) and (a_y^t, b_y^t) are either identical or

disjoint: if (a_x^t, b_x^t) meets (a_y^t, b_y^t) , in a point z say, we have $(a_x^t, b_x^t) = (a_z^t, b_z^t) = (a_y^t, b_y^t)$. An immediate consequence of this is the following:

For all $t \in [0, 1] \setminus C$ and $x \in [0, 1]$, the points a_x^t and b_x^t are unchanged at time t . (27)

It remains to define f_t for $t \in (s(q), t(q))$ 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 $t \in (s(q), t(q))$ as a homotopy between $f_{s(q)}$ and $f_{t(q)}$. When $q \in \mathcal{R}'$, this homotopy should retract $f_{s(q)} \upharpoonright [a(q), b(q)]$ to $z(q)$ in T by a direct application of (6). When $q \in \mathcal{P}$, our plan is to follow (15) and achieve the same result in two stages. We first wish to use (6) to retract $f_{s(q)} \upharpoonright [a^-(q), b^-(q)]$ to $z^-(q)$ in T . This should turn $f_{s(q)} \upharpoonright [a(q), b(q)]$ 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 (6) and implement (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 $[a(q), b(q)]$ to $z(q)$;
- if $q \in \mathcal{P}$: that $f_{s(q)}$ agrees with σ on $D(q) := (a(q), a^-(q)) \cup (b^-(q), b(q))$ and maps $[a^-(q), b^-(q)]$ to T ;
- if $q \in \mathcal{R}'$: that $f_{s(q)}$ maps $[a(q), b(q)]$ 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 $(a(q'), b(q'))$ 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 $\sigma(a(q)) = \sigma(b(q)) = z(q)$. (Recall (9) and (20).)

Next, we prove the second statement for $x \in (a(q), b(q))$. At time $t(q)$, none of these x was unchanged, since $q \in Q_x^{t(q)}$ for all these x . In fact, q is the greatest element of each of the \leq -chains $Q_x^{t(q)}$. Therefore $(a_x^{t(q)}, b_x^{t(q)}) = (a(q), b(q))$, and hence $f_{t(q)}(x) = \sigma(a(q)) = z(q)$ by (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 $[a(q), b(q)]$, but not in $D(q)$ if $q \in \mathcal{P}$. We have to show that $f_{s(q)}(x) \in T$. This follows from (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 $p \leq q$. If $p = q$ this means that $x \in D(q)$, contradicting our choice of x . Hence $p < q$, and $t(p) < s(q)$ by (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 (6) as outlined earlier to choose f_t for all $t \in (s(q), t(q))$ as follows:

For every $r \in \mathcal{R}'$, the paths $(f_t)_{t \in [s(r), t(r)]}$ form a homotopy retracting $f_{s(r)} \upharpoonright [a(r), b(r)]$ to $z(r)$. (28)

For every $p \in \mathcal{P}$, the paths $(f_t)_{t \in [s(p), t(p)]}$ form a homotopy that first retracts $f_{s(p)} \upharpoonright [a^-(p), b^-(p)]$ to $z^-(p)$ in T and then retracts the resulting path on $[a(p), b(p)]$ to $z(p)$ in $e(p)$. (29)

For our later proof that F is continuous, let us note an important property of the homotopies in (28) and (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 5 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 (28) and (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 $x \in (a(q), b(q))$, 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 $t \in [s(q), t(q)]$. (30)

The definition of F is now complete. Note finally that

$$\text{Im } F = \text{Im } \sigma. \quad (31)$$

This completes part four of our proof.

It remains to show that the set $(f_t)_{t \in [0, 1]}$ 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 $\text{Im } f_1 \subseteq T$; then (6) and our assumption that $\text{Im } \tau \subseteq T$ imply $f_1 \sim \tau$, which is good enough. For a proof that $\text{Im } f_1 \subseteq T$ consider any $x \in [0, 1]$. If x is not unchanged at time $t = 1$, then $f_1(x) \in T$ by (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 $p \in Q_x^1$, 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 $F((x - \varepsilon, x + \varepsilon), (t - \varepsilon, t + \varepsilon)) \subseteq U$, proceeding in two steps.

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

For every $\varepsilon > 0$, let

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

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 $F((x - \varepsilon_0, x + \varepsilon_0), (t - \varepsilon_0, t]) \subseteq U$.

If Q_ε is never empty but finite for some ε , there exists $q \in \mathcal{P} \cup \mathcal{R}'$ such that $t \in (s(q), t(q)]$. Since $(f_t)_{t \in [s(q), t(q)]}$ was defined as a homotopy, in (28) or (29), we have $F((x - \varepsilon_0, x + \varepsilon_0), (t - \varepsilon_0, t]) \subseteq U$ 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 (19). Hence $Q_{\varepsilon_0} \subseteq \hat{M}_i$, so the intervals $(a(q), b(q))$ with $q \in Q_{\varepsilon_0}$ are nested; put

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

From (9), (20) and Lemma 9 we know that $\sigma(a)$ is the limit of an infinite sequence of ends or distinct vertices $\sigma(a(q)) = z(q)$, 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 (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 (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 (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 $(x - \varepsilon_1, x + \varepsilon_1) \subseteq (a(q), b(q))$.

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 (25) yields $F(x', t') = \sigma(a_{x'}^{t'})$. These points lie in U' , since $a \leq a_{x'}^{t'} \leq a(q_0) < a + \delta$.

- (ii) Suppose now that $x \notin (a, b)$, say $x \leq a$. 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 (27), hence $f_t(x)$ is an end.) As $f_{t(q_0)}$ is continuous, there is an $\varepsilon_1 \leq t - t(q_0) (\leq \varepsilon_0)$ such that $t - \varepsilon_1 \notin C$ and $f_{t(q_0)}((x - \varepsilon_1, x + \varepsilon_1)) \subseteq U'$. 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

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

As $a_{x'}^{t'}$ is unchanged at time t' (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 (25). This last point lies in U' , by the above inequality and the choice of ε_1 .

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 $t' \in (s(q), t(q))$ with $q \in Q_{\varepsilon_1}$, we have $f_{t'}(x') = f_{t(q)}(x')$ for

these same x' , because $x' \notin (a, b) \supseteq (a(q), b(q))$ and hence $f_{t'}(x')$ remained constant throughout the homotopy defined in (28) or (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 $F((x - \varepsilon, x + \varepsilon), [t, t + \varepsilon]) \subseteq U$.

For every $\varepsilon > 0$, let

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

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 (22) we know that P^+ meets every Q_ε . So P^+ has no least element in \preceq , and

$$t = \inf\{t(p) \mid p \in P^+\}. \quad (32)$$

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.

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

$$f_t \text{ maps } (a(r), b(r)) \text{ to } z(r). \quad (33)$$

We shall then use (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 $[a(q), b(q)]$ with $q \in Q_{\varepsilon_0}$ are nested, and

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

by definition of $a(r)$ and $b(r)$, and (22). By our assumption for Case 2a, $x \in [a(r), b(r)]$.

If r is bad, then r lies in \mathcal{R}' and precedes all pairs in P^+ . By (24) and (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 (22), while $t(p) \leq t(r') \leq t$ (by (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_y^t for every $y \in (a(r), b(r))$, giving $a_y^t = a(r)$. As $t \notin C$, (25) yields

$$f_t(y) = \sigma(a_y^t) = \sigma(a(r)) \stackrel{(17)}{=} z(r),$$

completing the proof of (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 $y \in (a(r), b(r))$. As before, we have $f_t(y) = \sigma(a_y^t)$ by (25). If y is critical, then $Q_y^t \neq \emptyset$, and hence $a_y^t = a_y^-$ by (22) and the definitions of r , a_y^- and a_y^t . Thus

$$f_t(y) = \sigma(a_y^t) = \sigma(a_y^-) = z_y = z(r)$$

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

$$Q_y^t \cap \mathcal{P} = \emptyset. \quad (34)$$

We first prove that y is unchanged at time t . Indeed, otherwise $Q_y^t \neq \emptyset$, and by (34) there exists an $\tilde{r} = (\tilde{P}^-, \tilde{P}^+) \in \mathcal{R}'$ such that $y \in (a(\tilde{r}), b(\tilde{r}))$ and $t(\tilde{r}) \leq t$. By (24), \tilde{r} precedes all pairs in P^+ , which by (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 $y \in (a(\tilde{r}), b(\tilde{r})) \subseteq (a(p), b(p))$ and $t(p) \leq t$, contradicting (34).

Thus y is unchanged at time t . Then (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 (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 $[a, b] \subseteq [a(r), b(r)]$. 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 $(a(q), b(q))$ meets $[a, b]$. As y is unchanged at time t we have $(a(q), b(q)) \not\supseteq [a, b]$ 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 $[a(q), a^-(q)]$ 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 $[a(q), b(q)] = \bigcap_{p \in \tilde{P}_q^+} [a(p), b(p)]$ there is a pair $p \in \tilde{P}_q^+$ with $a(p) \in (a, a(q)) \subseteq (a, b)$, with a similar contradiction. Hence every point in $[a, b]$ is unchanged at time t and thus $f_t \upharpoonright [a, b] = \sigma \upharpoonright [a, b]$.

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 (8). We prove $\sigma(a) = z(r)$; the proof that $\sigma(b) = z(r)$ is analogous. If $a = a(r)$, then $\sigma(a) = \sigma(a(r)) = z(r)$, by definition of $z(r)$. If $a \neq a(r)$ then, by the choice of $[a, b]$, there is a sequence $y_0 < y_1 < \dots$ of points in $[a(r), b(r)]$ 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. As a is unchanged at time t , and hence $a \notin (a(q), b(q))$ for every $q \in Q^t$, we have $y_n \leq b_{y_n}^t \leq a$ and hence also $\lim_n b_{y_n}^t = a$. As σ is continuous,

the points $\sigma(b_{y_n}^t)$ converge in H to $\sigma(a)$. Since each y_n is critical but not bad, we have $\sigma(b_{y_n}^t) = z(r)$ for every n , and thus $\sigma(a) = z(r)$. This completes the proof of (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 $Q_x^t \neq \emptyset$ would imply $P^- \neq \emptyset$ by (22) and (24). Thus, $f_t = \sigma$. Consider the \leq -chain Q_x^1 of all $q \in \mathcal{P} \cup \mathcal{R}'$ with $x \in (a(q), b(q))$. By our assumption for Case 2a,

$$Q_x^1 \cap Q_\varepsilon \neq \emptyset \quad \text{for all } \varepsilon > 0. \quad (35)$$

Since $Q_x^t = \emptyset$ this means that Q_x^1 , like Q_ε , has no least element, and by (22) neither does $Q_x^1 \cap \mathcal{P} = Q_x^1 \cap P^+$. Therefore

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

Pick $p_0, p_1, \dots \in Q_x^1 \cap P^+$ with $\lim_n a(p_n) = a(r)$. As σ is continuous, $\lim_n z(p_n) = \sigma(a(r)) =: z(r)$. (Recall that $z(p_n) = \sigma(a(p_n))$.) By (9) and Lemma 9, $z(r)$ is an end.

To complete the proof of (33), we show that the interval $(a(r), b(r))$ avoids every domain of a pass of σ : then it is maximal with this property, and (33) follows from (8) and the fact that $f_t = \sigma$. (For the application of (8) note that $\sigma(b(r)) = \sigma(a(r))$, by taking limits in (9).) Suppose that $(a(r), b(r))$ 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 $y \notin (a(p'), b(p'))$. As this contradicts the choice of y as a point in $(a(r), b(r))$, we have $p \notin Q_x^1$. Thus, $x \notin (a(p), b(p))$, and hence $p \not\leq q$ for all $q \in Q_x^1$. But p is nested with every such q , since $y \in (a(p), b(p)) \cap (a(q), b(q))$. Therefore $p < q$, and hence

$$t < t(p) < s(q) \quad \text{for all } q \in Q_x^1,$$

by (24) and since $p \notin P^- = \emptyset$. But then for $\varepsilon < t(p) - t$ we have $Q_\varepsilon \cap Q_x^1 = \emptyset$, contradicting (35). This completes the proof of (33).

We have thus shown for both sets of definitions that f_t maps $(a(r), b(r))$, and hence also $[a(r), b(r)]$, to the end $z(r)$. As $x \in [a(r), b(r)]$, 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 $(a(r) - \delta, b(r) + \delta)$ to U' . By the definition of $[a(r), b(r)]$ (in either case) we can find a pair $p_0 \in P^+$ such that $a(p_0) \in (a(r) - \delta, a(r))$ and $b(p_0) \in (b(r), b(r) + \delta)$. Choose ε so that $t + \varepsilon \notin C$, and small enough that $t + \varepsilon < t(p_0)$ as well as $(x - \varepsilon, x + \varepsilon) \subseteq (a(p_0), b(p_0))$. 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' (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 (25).

(b) Our second case is that there is an ε_0 such that $x \notin (a(q), b(q))$ 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 $(a(q), b(q))$ 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 (25), since $Q_{x'}^{t'} = Q_{x'}^t$. For $t' \in [t, t + \varepsilon_1] \cap C$, say $t' \in (s(q), t(q))$ with $q \in Q_{\varepsilon_1}$, this implies $F(x', t') = F(x', s(q)) = F(x', t)$ by (28) or (29), since $x' \notin (a(q), b(q))$ 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:

For every ε , the interval $(x - \varepsilon, x + \varepsilon)$ meets $(a(q), b(q))$ for infinitely many $q \in Q_\varepsilon$. (36)

By (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 (9) and (20), the sequence $(\sigma(a_n))_{n \in \mathbb{N}}$ is a sequence of vertices and ends, and every vertex appears only finitely often (Lemma 9). Therefore $\sigma(x)$ is an end.

Let us show that x is unchanged at time t . By (36), any interval $(a(q), b(q))$ that contains x meets some $(a(q_n), b(q_n))$, but is not contained in it since $q_n \in Q_{\varepsilon_0}$ and hence $x \in (a(q), b(q)) \setminus (a(q_n), b(q_n))$. Thus $q > q_n$ by (21), and $t(q) > t(q_n) > t$ by (24), giving $q \notin Q_x^t$ as desired. As $x \notin (a(q), b(q))$ 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 $\sigma((x - \varepsilon, x + \varepsilon)) \subseteq U'$ and $t + \varepsilon \notin C$. Then for every $x' \in [x, x + \varepsilon)$ and $t' \in [t, t + \varepsilon] \setminus C$ we have $x \leq a_{x'}^{t'}$ ($\leq x'$), because x is unchanged at time t' and hence $x \notin (a(q), b(q))$ for every $q \in Q_{x'}^{t'}$. Likewise, for $x' \in (x - \varepsilon, x]$ we have $x' \leq b_{x'}^{t'} \leq x$. Thus for every $x' \in (x - \varepsilon, x + \varepsilon)$ and $t' \in [t, t + \varepsilon] \setminus C$ we have $a_{x'}^{t'} \in (x - \varepsilon, x + \varepsilon)$ or $b_{x'}^{t'} \in (x - \varepsilon, x + \varepsilon)$, say $b_{x'}^{t'} \in (x - \varepsilon, x + \varepsilon)$. By (27), $b_{x'}^{t'}$ is unchanged at time t' and hence also at time $t \leq t'$. We thus have $F(x', t') = \sigma(a_{x'}^{t'}) = \sigma(b_{x'}^{t'}) = f_t(b_{x'}^{t'}) \in U'$ by (25).

This completes the fifth 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 traverses 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 (5) to straighten σ to a path σ' , which we then trimmed further to obtain a path σ'' satisfying (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 $\text{Im } \sigma$; see (5). In the second homotopy we retracted segments $\sigma' \upharpoonright [a, b] \subseteq T$ to paths $\sigma'' \upharpoonright [a, b]$ with image $\sigma'(a)T\sigma'(b)$.

This is the unique $\sigma'(a) - \sigma'(b)$ arc in T , so Lemma 1 implies that the image of $\sigma' \upharpoonright [a, b]$ contains it. The homotopy between $\sigma' \upharpoonright [a, b]$ and $\sigma'' \upharpoonright [a, b]$, which (6) says runs in the union of the images of those two paths, thus in fact runs in $\text{Im } \sigma' \subseteq \text{Im } \sigma$, and hence so does the entire second homotopy $\sigma' \sim \sigma'' = f_0$. The third homotopy, $f_0 \sim f_1$, runs in $\text{Im } f_0 \subseteq \text{Im } \sigma' \subseteq \text{Im } \sigma$ by (31). Similarly, the last homotopy $\tau \sim \tau'$ runs in $\text{Im } \tau$, and the penultimate one, $f_1 \sim \tau'$, runs in $\text{Im } f_1 \cup \text{Im } \tau' \subseteq \text{Im } \sigma \cup \text{Im } \tau$. 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 $\text{Im } \sigma \cup \text{Im } \tau$ between σ and τ . (37)

To complete our proof of Lemma 14, we now consider the case in which both σ and τ traverse chords. By Lemma 21, 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 (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 2: note that $\sigma \upharpoonright [a, b]$ and $\sigma' \upharpoonright [a, b]$ 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 $w_{\sigma \upharpoonright [a, b]}$ to $w_{\sigma' \upharpoonright [a, b]} = \emptyset$. Likewise, we obtain $\tau \sim \tau'$. It thus suffices to prove $\sigma' \sim \tau'$, from our assumptions that

$$w_{\sigma'} = r(w_\sigma) = r(w_\tau) = w_{\tau'} \quad (38)$$

(cf. Lemma 13).

Our aim is to use (7) to obtain the desired homotopy $\sigma' \sim \tau'$. But (7) requires that the two paths considered have the same passes, not just the same trace. In order to make (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 corresponding chords.

Every path $\alpha: [0, 1] \rightarrow H$ 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}_α 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 (38) defines an order-preserving bijection $\pi: \mathcal{I}_{\sigma'} \rightarrow \mathcal{I}_{\tau'}$.

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 $\pi_{\sigma'}: \mathcal{I} \rightarrow \mathcal{I}_{\sigma'}$ and $\pi_{\tau'}: \mathcal{I} \rightarrow \mathcal{I}_{\tau'}$ that map open to open intervals bijectively, and trivial to trivial intervals, and which commute with $\pi: \mathcal{I}_{\sigma'} \rightarrow \mathcal{I}_{\tau'}$. We can now define surjective maps $\varphi, \psi: [0, 1] \rightarrow [0, 1]$ 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 (7) yields the desired homotopy $\sigma'' \sim \tau''$, by Lemma 2. \square

References

- [1] H. Abels and P. Strantzalos, *Proper transformation groups*, in preparation.
- [2] H.-J. Baues and A. Quintero, *Infinite Homotopy Theory*, Kluwer Academic Publ., 2001.
- [3] J.W. Cannon and G.R. Conner, *The combinatorial structure of the Hawaiian Earring group*, *Topology Appl.* **106** (2000), 225–271.
- [4] I.M. Chiswell and T.W. Müller, *A class of groups with canonical \mathbb{R} -tree action*, in preparation.
- [5] G.R. Conner and K. Eda, *Fundamental groups having the whole information of spaces*, *Topology Appl.* **146–147** (2005), 317–328.
- [6] R. Diestel, *Graph Theory* (3rd edition), Springer-Verlag, 2005,
Electronic edition available at:
<http://www.math.uni-hamburg.de/home/diestel/books/graph.theory>.
- [7] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, *J. Combin. Theory (Series B)* **87** (2003), 197–206.
- [8] ———, *Topological paths, cycles and spanning trees in infinite graphs*, *Europ. J. Combinatorics* **25** (2004), 835–862.
- [9] R. Diestel and I. Leader, *A proof of the Bounded Graph Conjecture*, *Invent. math.* **108** (1992), 131–162.
- [10] ———, *Normal spanning trees, Aronszajn trees and excluded minors*, *J. London Math. Soc.* **63** (2001), 16–32.
- [11] R. Diestel and P. Sprüssel, *The homology of locally finite graphs with ends*, preprint 2008.
- [12] K. Eda, *Free σ -products and noncommutatively slender groups*, *J. Algebra* **148** (1991), 243–263.
- [13] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, *Math. Zeitschr.* **33** (1931), 692–713.
- [14] ———, *Neuaufbau der Endentheorie*, *Annals of Mathematics* **43** (1942), 261–279.

- [15] R. Halin, *Über unendliche Wege in Graphen*, Math. Annalen **157** (1964), 125–137.
- [16] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.
- [17] A. Hatcher, *Algebraic Topology*, Cambridge Univ. Press, 2002.
- [18] G. Higman, *Unrestricted free products and varieties of topological groups*, J. London Math. Soc. **27** (1952), 73–81.
- [19] B. Hughes and A. Ranicki, *Ends of complexes*, Cambridge Univ. Press, 1996.
- [20] H.A. Jung, *Connectivity in infinite graphs*, Studies in Pure Mathematics (L. Mirsky, ed.), Academic Press, 1971.
- [21] B. Krön, *End compactifications in non-locally-finite graphs*, Math. Proc. Cambridge Phil. Soc. **131** (2001), 427–443.
- [22] R. Möller, *Ends of graphs*, Math. Proc. Cambridge Phil. Soc. **111** (1992), 255–266.
- [23] ———, *Ends of graphs II*, Math. Proc. Cambridge Phil. Soc. **111** (1992), 455–460.
- [24] J. Stillwell, *Classical topology and combinatorial group theory*, Springer-Verlag, 1980.
- [25] C. Thomassen and A. Vella, *Graph-like continua and Menger’s theorem*, preprint 2006.
- [26] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, J. Combin. Theory (Series B) **58** (1993), 248–268.
- [27] S. Willard, *General topology*, Dover Publications, 2004.
- [28] W. Woess, *Random walks on infinite graphs and groups*, Cambridge University Press, 2002.

Reinhard Diestel
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 20146 Hamburg
 Germany

Philipp Sprüssel
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 20146 Hamburg
 Germany