

**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:

$$\begin{aligned} & \text{If } R \text{ is a reduction of } w \text{ then } \{\{s, s'\} \in R \mid w(s) \in A_I\}, \\ & \text{with the ordering induced from } R, \text{ is a reduction of } w \upharpoonright I. \end{aligned} \tag{1}$$

In particular:

$$\begin{aligned} & \text{Any result of first reducing and then restricting a word can} \\ & \text{also be obtained by first restricting and then reducing it.} \end{aligned} \tag{2}$$

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 $\bar{e}_1\bar{e}_0\bar{e}_1\bar{e}_0\bar{e}_2\bar{e}_0\bar{e}_2\bar{e}_0\cdots\bar{e}_i\bar{e}_0\bar{e}_i\bar{e}_0\in F_I$ 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

