

HAMBURGER BEITRÄGE ZUR MATHEMATIK

Heft 277

On the homology of locally finite graphs

R. Diestel, Hamburg

P. Sprüssel, Hamburg

This is an extended version of our paper

The homology of locally finite graphs with ends

It differs from that paper only in that it offers proofs for Lemmas 2, 3, 4 and 9, as well as a new footnote in Section 5.

September 2007

Revised January 21st, 2008

The homology of locally finite graphs with ends

Reinhard Diestel and Philipp Sprüssel

Abstract

We show that the topological cycle space of a locally finite graph is a canonical quotient of the first singular homology group of its Freudenthal compactification, and we characterize the graphs for which the two coincide. We construct a new singular-type homology for non-compact spaces with ends, which in dimension 1 captures precisely the topological cycle space of graphs but works in any dimension.

1 Introduction

Graph homology is traditionally, and conveniently, simplicial: a graph G is viewed as a 1-complex, and one considers its first simplicial homology group. In graph theory, coefficients are typically taken from a field such as \mathbb{Z}_2 , \mathbb{R} or \mathbb{C} , which makes the group into a vector space called the *cycle space* of G .

For reasons to become apparent later we denote this space as $\mathcal{C}_{\text{fin}} = \mathcal{C}_{\text{fin}}(G)$. For the moment it will suffice to take our coefficients from \mathbb{Z}_2 and interpret the elements of \mathcal{C}_{fin} as sets of edges. For finite graphs G , there are a number of classical theorems relating $\mathcal{C}_{\text{fin}}(G)$ to other properties of G , such as planarity. (Think of MacLane’s or Whitney’s theorem, or the Kelmans-Tutte planarity criterion.) The cycle space \mathcal{C}_{fin} has thus become one of the standard aspects of finite graphs used in their structural analysis.

When G is infinite, however, the space \mathcal{C}_{fin} no longer adequately describes the homology of G . Most of the theorems describing the interaction of \mathcal{C}_{fin} with other properties of G —including all those cited above—fail when G is infinite. However, the traditional role of the cycle space in these cases can be restored by defining it slightly differently: when G is locally finite, one takes as generators not the edge sets of the (finite) cycles in G —as one would to generate \mathcal{C}_{fin} —but the (possibly infinite) edge sets of all its *topological circles*, the homeomorphic images of the circle S^1 in the Freudenthal compactification $|G|$ of G by its ends. (One also has to allow infinite sums in the generating process; for these to be well-defined, each edge may occur in only finitely many terms.) We denote this more general space, the *topological cycle space* of G , by $\mathcal{C} = \mathcal{C}(G)$.

The space \mathcal{C} had not been considered in graph theory before [13] appeared, and it has been surprisingly successful at extending the classical cycle space theory of finite graphs to locally finite graphs [3, 4, 5, 8, 22, 34]. However, a question raised in [9] but still unanswered is how new, from a topological viewpoint, is the homology described implicitly by \mathcal{C} . Since topological circles are (images of simplices representing) singular 1-cycles in $|G|$, it is natural to ask how closely $\mathcal{C}(G)$ is related to the first singular homology group not of G itself but of its compactification $|G|$. Indeed it is not clear whether the two

coincide by some natural canonical isomorphism, so that $\mathcal{C}(G)$ would be just another way of looking at $H_1(|G|)$.

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

For the proof that f has a non-trivial kernel we have to go some way towards the solution of another problem (solved fully in [18]): to find a combinatorial description of the fundamental group of the space $|G|$ for an arbitrary connected locally finite graph G .¹ We describe $\pi_1(|G|)$, as for finite G , in terms of reduced words in the oriented chords of a spanning tree. However, when G is infinite this does not work for arbitrary spanning trees; we have to allow infinite words of any countable order type; and reduction by cancelling adjacent inverse sequences of letters does not suffice. However, the kind of reduction we need can be described in terms of word reductions in the free groups F_I on all the finite subsets I of chords, which enables us to embed the group F_∞ of infinite reduced words in the inverse limit of those F_I , and handle it in this form. On the other hand, mapping a loop in $|G|$ to the sequence of chords it traverses, and then reducing that sequence (or word), turns out to be well defined on homotopy classes and hence defines an embedding of $\pi_1(|G|)$ as a subgroup in F_∞ . This combinatorial description of $\pi_1(|G|)$ then enables us to define an invariant on 1-chains in $|G|$ that can distinguish some elements of the kernel of f from boundaries of singular 2-chains, completing the proof that f need not be injective.

Our second aim, then, is to reconcile these different treatments of homology for non-compact spaces between topology and graph theory. We construct a singular homology theory which, in principle, works for any topological space with a subset of points called ends² and which, when applied to graphs, captures precisely their topological cycle space. Essentially, we shall allow infinite sums of cycles and boundaries when building their respective groups, but start from *finite* chains with zero boundary as generators. Thus, topological circles are 1-cycles, as desired. But if G is a 2-way infinite path, then its edges form an infinite 1-chain with zero boundary that is *not* a 1-cycle, because it is not a (possibly infinite) sum of finite 1-cycles. Our new homology thus lies between the usual singular homology and the ‘open homology’ that is built from arbitrary locally finite chains without any further restriction (see, eg, [21]).

Our hope with this paper is to stimulate further work in two directions. One is that its new topological guise makes the cycle space \mathcal{C} accessible to topological methods that might generate some windfall for the study of graphs. And conversely, that since the new topological framework of \mathcal{C} is defined for more general spaces (and higher dimensions), the proven usefulness of \mathcal{C} for graphs might find some interesting topological analogues.

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

²The specification of ends is crucial. They will play a role that differs from that of ordinary points, and it is the use of this role that makes this homology more flexible for spaces where ends occur naturally, such as end compactifications of locally compact spaces.

2 Terminology and basic facts

In this section we briefly run through any non-standard terminology we use. We also list without proof a few easy lemmas that we shall need, and use freely, later on.

For graphs we use the terminology of [10], for topology that of Hatcher [26]. We reserve the word ‘boundary’ for homological contexts and use ‘frontier’ for the closure of a set minus its interior. Our use of the words ‘path’, ‘cycle’ and ‘loop’, where these terminologies conflict, is as follows. The word *path* is used in both senses, according to context (such as ‘*path in X*’, where X was previously introduced as a graph or as a topological space). Note that while topological paths need not be injective, graph-theoretical paths are not allowed to repeat vertices or edges. The term *cycle* will be used in the topological sense only, for a (usually 1-dimensional) singular chain with zero boundary. When we do need to speak about graph-theoretic cycles (i.e., about finite connected graphs in which every vertex has exactly two incident edges) we shall instead refer to the edge sets of those graphs, which we shall call *circuits*. Our graphs may have multiple edges but no loops. This said, we shall from now on use the term *loop* topologically: for a topological path $\sigma: [0, 1] \rightarrow X$ with $\sigma(0) = \sigma(1)$. This loop is *based at* the point $\sigma(0)$. Given any path $\sigma: [0, 1] \rightarrow X$, we write $\sigma^-: s \mapsto \sigma(1 - s)$ for the inverse path. An *arc* in a topological space is a subspace homeomorphic to $[0, 1]$.

Lemma 1 ([25]). *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

All the CW-complexes we consider will be *locally finite*: every point has an open neighbourhood meeting only finitely many cells. Note that a compact subset of such a complex can meet the closures of only finitely many cells, and that locally finite CW-complexes are first-countable [27, 31].

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 [19, 20]. More topological background on this can be found in [2]; for applications to groups see e.g. [32, 33, 36, 37].

For graphs, ends and the Freudenthal compactification are more usually defined combinatorially, as follows [10, 24, 29]. Let G be a connected locally finite graph. A 1-way infinite 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 [12] that this combinatorial definition of an end coincides with the topological one given earlier for locally finite complexes. The Freudenthal 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 .³ Note that the frontier of $\hat{C}(S, \omega)$ in $|G|$ is precisely S , the frontier of any spanning subgraph of G is precisely $\Omega(G)$, and every ray converges to the end containing it.

The end structure of G is best reflected by a *normal spanning tree*; such trees exist in every connected countable graph [28]. A spanning tree T of G with root r is *normal* if the vertices u, v of every edge $e = uv$ of G are comparable in the order \leq which (T, r) induces on $V(G)$. (Recall that $u \leq v$ if u lies on the unique r - v path rTv of T between r and v .) A key property of normal spanning trees is that the intersection of the down-closures of two vertices separates them in G . This implies that every end of G is represented by a unique ray in T starting at r , and hence that adding all the ends of G to T does not create any circles. More generally, it is not hard to prove the following:

³The definition given in [10] is formally more general, but equivalent to the simpler definition given here when G is locally finite. Generalizations are studied in [30, 35].

Lemma 3. *Let T be a normal spanning tree of G , and let $\overline{T} := T \cup \Omega(G)$ denote its closure in $|G|$. Then for every closed connected set $X \subseteq \overline{T}$ and $x \in X$ there is a deformation retraction of X onto x . In particular, X is contractible.*

Proof. Let X be a closed connected subset of \overline{T} , and let $x \in X$. Then X is also closed in $|G|$ and hence arc-connected [15, Theorem 2.6]. For every $y \in X$ there is a unique x - y arc $x\overline{T}y$ in \overline{T} [10, Theorem 8.5.7] which hence lies in X . The space \overline{T} is metrizable so that every edge between levels n and $n+1$ has length $1/2^{n+1}$ and hence every end has distance 1 from the root [11]. X inherits this metric d , note that $d(x, y) \leq 2$ for all $y \in X$. Further, if $z \in y\overline{T}y'$ for some $y, y' \in \overline{T}$ we have $d(y, y') = d(y, z) + d(z, y')$. We construct a homotopy F in \overline{T} from the identity on \overline{T} to the map $\overline{T} \rightarrow \{x\}$ so that $F(y, t) \in x\overline{T}y \subseteq X$ for every $y \in X$ and $t \in [0, 1]$; then $F \upharpoonright (X \times [0, 1])$ will be the desired homotopy for X . For every $y \in \overline{T}$ and $t \in [0, 1]$ let $F(y, t)$ be the unique point on $x\overline{T}y$ at distance $(1-t)d(x, y)$ 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 $y, y' \in \overline{T}$ and $t \in [0, 1]$; then for every $\varepsilon > 0$ and every $y, y' \in \overline{T}$ with $d(y, y') < \varepsilon/3$ and $t, t' \in [0, 1]$ with $|t - t'| < \varepsilon/3$ we have

$$\begin{aligned} d(F(y, t), F(y', t')) &\leq d(F(y, t), F(y', t)) + d(F(y', t), F(y', t')) \\ &\leq d(y, y') + |t - t'| \cdot d(x, y') \\ &< \varepsilon/3 + (\varepsilon/3) \cdot 2 = \varepsilon. \end{aligned}$$

As $x\overline{T}y$ and $x\overline{T}y'$ are closed, there is a last z point on $x\overline{T}y$ that is also in $x\overline{T}y'$. As \overline{T} contains a unique arc between any two points in \overline{T} , we have $y\overline{T}y' = y\overline{T}z \cup z\overline{T}y'$ and hence $d(y, y') = d(y, z) + d(z, y')$. If $F(y, t) \in z\overline{T}y$ and $F(y', t) \in z\overline{T}y'$, 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 both $F(y, t)$ and $F(y', t)$ are contained in $x\overline{T}y$ or in $x\overline{T}y'$. In particular, one of $F(y, t), F(y', t)$ lies on the arc between the other and x . Then $d(F(y, t), F(y', t)) = |d(x, F(y, t)) - d(x, F(y', t))| = (1-t) \cdot |d(x, y) - d(x, y')| \leq d(y, y')$. \square

Lemma 3 implies that \overline{T} contains no topological circle. Equivalently: for any two points $x, y \in \overline{T}$ there is a unique arc in \overline{T} between x and y . We denote this arc by $x\overline{T}y$. The uniqueness of $x\overline{T}y$ implies that none of its inner points can be an end. (Every arc containing an end also contains a vertex, and any two vertices of T can also be joined by an arc in T itself.)

When T is a normal spanning tree of G , every end ω in $|G|$ has a neighbourhood basis consisting of open sets $\hat{C} = \hat{C}(S, \omega)$ such that S is closed downwards, i.e. where $s' \leq s \in S$ implies $s' \in S$. We call these sets \hat{C} the *basic open neighbourhoods* of the ends⁴ of G (given T). An important property of these sets is that for any two points $x, y \in \hat{C}$ we also have $x\overline{T}y \subseteq \hat{C}$.

Now let $S' = S \cup N(S)$, the (finite) set of vertices in S and their neighbours. We call the subset $\hat{C}(S', \omega)$ of $C(S, \omega)$ the *inside of $\hat{C}(S, \omega)$ around ω* . Note that the neighbours v of vertices $u \in C(S', \omega)$, as well as the edges uv , also lie in $C(S, \omega)$.

More background on normal spanning trees, including an existence proof, can be found in [10, 16, 17].

⁴The *basic open neighbourhoods* of a point $x \in G$ are the connected open neighbourhoods of x containing no vertex other than possibly x .

Let us now introduce the topological cycle space \mathcal{C} of G . This is usually defined over \mathbb{Z}_2 (which suffices for its role in graph theory), but we wish to prove our main results more generally with integer coefficients. (The \mathbb{Z}_2 case will follow, but it should be clear right away that the non-injectivity of our homomorphism $H_1 \rightarrow \mathcal{C}$ is not just a consequence of a wrong choice of coefficients.) We therefore need to speak about orientations of edges.

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*.

Given a set E of edges in G , we write \vec{E} for the set of their orientations, two for every edge in E . Given a partition (U, V) of the vertex set of G , we write $\vec{E}(U, V)$ for the set of all its oriented edges (e, u, v) with $u \in U$ and $v \in V$, and call this set an *oriented cut* of G .

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 4. *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 ($n = 1, 2, \dots$). As $[0, 1]$ is compact, the sequence s_1, s_2, \dots has an accumulation point x , which is also an accumulation point of t_1, t_2, \dots . But now σ fails to be continuous at x , because $\{\sigma(s_n), \sigma(t_n)\} = \{u, v\}$ for each n but each of u and v has a neighbourhood not containing the other. \square

A loop that is injective on $[0, 1)$ is a *circle* in $|G|$. (In most of our references, the term *circle* is used for the image of such a loop.) The set of all edges traversed by a circle is a *circuit*. It is easy to show that the image of a circle is uniquely determined by its circuit C , being the closure of $\bigcup C$ in $|G|$.

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

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

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

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

We remark that $\vec{\mathcal{C}}$ is closed also under infinite thin sums [13], but this is neither obvious nor generally true for thin spans of subsets of $\vec{\mathcal{E}}$ [6]. We remark further that composing the functions in $\vec{\mathcal{C}}$ with the canonical homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}_2$ yields the usual *topological cycle space* $\mathcal{C}(G)$ of G as studied in [3, 4, 5, 8, 7, 13, 14, 15, 22, 23, 34], the \mathbb{Z}_2 vector space of subsets of E obtained as thin sums of (unoriented) circuits.

The topological cycle space $\mathcal{C}(G)$ can be characterized as the set of those subsets of E that meet every *finite* cut of G in an even number of edges [13, 10]. The characterization has an oriented analogue:

Theorem 5. *An element φ of $\vec{\mathcal{E}}$ lies in $\vec{\mathcal{C}}$ if and only if $\sum_{\vec{e} \in \vec{F}} \varphi(\vec{e}) = 0$ for every finite oriented cut \vec{F} of G .*

The proof of Theorem 5 is not completely trivial. But it adapts readily from the unoriented proof given e.g. in [10], which we leave to the reader to check if desired.

When G is fixed, we write C_n for the group of singular n -chains in $|G|$ (with coefficients in \mathbb{Z} unless otherwise mentioned), $Z_n = \text{Ker } \partial_n$ and $B_n = \text{Im } \partial_{n+1}$ for the corresponding groups of cycles and boundaries, and $H_n = Z_n/B_n$. We view all singular 1-simplices as maps from the real interval $[0, 1]$ to $|G|$. The homology class of a cycle z is denoted by $[z]$.

A cycle that can be written as a sum of 1-simplices no two of which share their first point is an *elementary cycle*. Every 1-cycle is easily seen to be a sum of elementary 1-cycles, a decomposition which is not normally unique.

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

Lemma 6. *Let σ be a singular 1-simplex in $|G|$, and let $s \in (0, 1)$. Write σ' and σ'' for the 1-simplices obtained from the restrictions of σ to $[0, s]$ and to $[s, 1]$ by reparametrizing linearly. Then $\sigma' + \sigma'' - \sigma \in B_1$. \square*

When σ is a summand in a cycle $z \in Z_1$, we shall say that the equivalent cycle z' obtained by replacing σ with $\sigma' + \sigma''$ in the sum arises by *subdividing* σ (at s or at $\sigma(s)$). A frequent application of Lemma 6 is the following:

Corollary 7. *Every non-zero element of $H_1(|G|)$ is represented by a sum of loops each based at a vertex.*

Proof. Pick a cycle representing a given homology class, and decompose it into elementary cycles. Use Lemma 6 to concatenate their simplices into a single loop. If such a loop α passes through a vertex, we can subdivide it there and suppress its original boundary point, obtaining a homologous loop based at that vertex. If α does not pass through a vertex, then $\text{Im } \alpha \subseteq \vec{e}$ for some edge e (since non-trivial sets of ends are never connected), so α is null-homotopic and $[\alpha] = 0$. \square

3 Comparing $H_1(|G|)$ with $\vec{\mathcal{C}}(G)$

Let $G = (V, E)$ be a connected locally finite graph. Our aim in this section is to compare the first singular homology group H_1 of $|G|$ (with integer coefficients)

with the oriented topological cycle space $\vec{\mathcal{C}}(G)$ of G , the group of thin sums of oriented circuits in G . When G is finite then $|G| = G$, and all circuits and their thin sums are finite. Hence in this case $\vec{\mathcal{C}}$ is just the first simplicial homology group of G , so the two groups are indeed the same.

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

Our approach will be to define a homomorphism $f: H_1 \rightarrow \vec{\mathcal{E}}(G)$ that counts for a given homology class h how often the 1-simplices of a cycle representing h , when properly concatenated, traverse a given edge \vec{e} ; we then let $f(h) \in \vec{\mathcal{E}}(G)$ map \vec{e} to this number.⁵ We shall prove that $f(h)$ always lies in $\vec{\mathcal{C}}$ and, perhaps surprisingly, that f maps H_1 onto $\vec{\mathcal{C}}$. However, we find that f is not normally injective. Our first main result characterizes the graphs for which it is:

Theorem 8. *The map $f: H_1(|G|) \rightarrow \vec{\mathcal{E}}(G)$ is a group homomorphism onto $\vec{\mathcal{C}}(G)$, which has a non-trivial kernel if and only if G contains infinitely many (finite) circuits.*

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

We remark that the last condition in Theorem 8 can be rephrased in various natural ways: that G has a spanning tree with infinitely many chords; that every spanning tree of G has infinitely many chords; or that G contains infinitely many *disjoint* (finite) circuits [10, Ex. 37, Ch. 8]. The remainder of this section and the next two sections will be devoted to the proof of Theorem 8.

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

The following lemma is easy to prove using homotopies in S^1 , combined by Lemma 2:

Lemma 9. *Let $\alpha: [0, 1] \rightarrow |G|$ be a loop based at a vertex. If α traverses e exactly k times in its natural direction and exactly ℓ times in the opposite direction, then $\pi([f_e \circ \alpha]) = k - \ell$.*

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

The domains of distinct passes of α through e are closed subintervals of $[0, 1]$ meeting at most in their boundary points. The rest of $[0, 1]$ is a finite disjoint union of open intervals (s, t) (or $(s, 1]$ or $[0, t)$). Each of these is in turn a disjoint union, possibly infinite, of open intervals (s', t') which α maps to \vec{e} and

⁵The precise definition of f will be given shortly.

closed intervals which $f_e \circ \alpha$ maps to $1 \in \mathbb{C}$. Since $[s, t]$, by definition, contains no pass through e , α always maps s' and t' to the same endvertex of e . Then $\alpha \upharpoonright [s', t']$ is homotopic to the constant map to that vertex, and $(f_e \circ \alpha) \upharpoonright [s', t']$ is homotopic to the constant map to 1. By Lemma 2, these homotopies combine to a homotopy of $(f_e \circ \alpha) \upharpoonright [s, t]$ to the constant map with value 1.

We deduce that $f_e \circ \alpha$ is homotopic to a concatenation $\sigma_1 \cdot \dots \cdot \sigma_n$ of loops in S^1 of which (after reparametrization) k are equal to η_1 and ℓ are equal to the inverse loop $\overline{\eta_1}: \lambda \mapsto 1 - \eta_1(\lambda)$, and the rest are constant with value 1. The result follows. \square

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

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

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

Lemma 10. $\text{Im } f \subseteq \vec{\mathcal{C}}(G)$.

Proof. By Theorem 5 it suffices to show that for every finite oriented cut \vec{F} of G and every $h \in H_1(|G|)$ we have $\sum_{\vec{e} \in \vec{F}} f(h)(\vec{e}) = 0$. Let $\vec{F} = \vec{E}(U, U')$ and h be given, let $F = \{e \mid \vec{e} \in \vec{F}\}$, and assume for simplicity that the orientations $\vec{e} \in \vec{F}$ of these edges are their natural orientations. Since f is a homomorphism, we may assume that h is represented by an elementary 1-cycle, which we may choose by Corollary 7 to consist of a loop α based at a vertex. We shall prove that α traverses the edges in F as often from U to U' as it does from U' to U . Then

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

by Lemma 9.

Let $[s_1, t_1], \dots, [s_n, t_n]$ be the domains of the passes of α through edges of F . In order to prove that as many of these passes are from a vertex in U to one in U' as vice versa, it suffices to show that each of the segments $\beta = \alpha \upharpoonright [t_i, s_{i+1}]$ has either all its vertices in U or all its vertices in U' , assuming for simplicity that α is based at $t_n =: t_0$. If the starting vertex $\beta(t_i)$ of β lies in U , say, put

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

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

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

Lemma 11. $\text{Im } f \supseteq \vec{\mathcal{C}}(G)$.

Proof. Let $\varphi = \sum_{\alpha \in A} \varphi_\alpha \in \vec{\mathcal{C}}(G)$ be an arbitrary thin sum of oriented circuits, where each α is a circle in $|G|$. Ignoring any φ_α that are constant with value 0, we may assume that each α is based at a vertex $v(\alpha)$. (Recall that if the image of α contains no vertex it must lie inside an edge, because non-trivial sets of ends cannot be connected.) We shall construct a loop τ in $|G|$ such that $f([\tau]) = \varphi$.

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

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

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

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

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

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

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

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

To complete the proof of Theorem 8 it remains to show that f has a non-trivial kernel if and only if G contains infinitely many circuits. The forward implication of this is easy. Indeed, suppose that G contains only finitely many circuits, and let T be a normal spanning tree of G . Then T has only finitely many chords, so $|G|$ is homotopy equivalent to a finite graph (Lemma 3). Hence, as is well known, $H_1(|G|)$ equals the first simplicial homology group of G viewed as a 1-complex, which in turn is clearly isomorphic to $\vec{\mathcal{C}}(G)$. Therefore f must be injective.

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

The next section, therefore, contains an interlude in which we describe the fundamental group $\pi_1(|G|)$ of an arbitrary locally finite connected graph G combinatorially. We shall then complete the proof of Theorem 8 in Section 5.

4 A combinatorial description of $\pi_1(|G|)$

Our aim in this section is to prove some aspects of a combinatorial description of $\pi_1(|G|)$ that we need for our proof of Theorem 8. In [18], we give a more comprehensive such description; see Theorem 17 below.

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. 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 .

Our description of $\pi_1(|G|)$ for infinite G will be similar in spirit, but more complex. We shall start not with an arbitrary spanning tree but with a normal spanning tree. (The trees that work are precisely the *topological spanning trees* defined in [15] or [10], which include the normal spanning trees.) 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 reductions of words in a way that captures homotopy of paths.

Consider the following example. Let G be the infinite ladder, with a spanning tree T consisting of one side of the ladder and all its rungs (drawn bold in 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 be reducible. 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, one can construct graphs which, for any suitable spanning tree, contain null-homotopic loops whose trace of chords contains no such cancelling subsequences (of any order type).⁶

We shall therefore define the reduction of infinite words differently, in a non-recursive way: just as a homotopy can shrink a loop simultaneously (rather than recursively) in many places at once, our reduction ‘steps’ will be ordered linearly but not be well-ordered. This definition is less straightforward, but it has an important property: as for finite G , our notion of reduction will be purely combinatorial and make no reference to the topology of $|G|$.

The main step then will be to show that $\pi_1(|G|)$ embeds as a subgroup in the group of reduced words, and how. We shall see that, as in the finite case, the map assigning to a path in $|G|$ its trace of chords and reducing that trace is well defined on homotopy classes. In [18] we shall prove that the map it induces on these classes is injective. Then $\pi_1(|G|)$ can be viewed as a subgroup of the group of reduced infinite words, which in turn can be viewed as a subgroup of the inverse limit of the free groups with generators the finite sets of oriented chords of any fixed normal spanning tree (Theorem 17).

Let us make all this precise. Let G be a locally finite connected graph, fixed throughout this section. Let T be a fixed normal spanning tree of G , and write \bar{T} for its closure $T \cup \Omega(G)$ in $|G|$. 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 . If T has only finitely many chords, then $|G|$ is homotopy equivalent to a finite graph, by Lemma 3. Let us therefore assume that T has infinitely many chords. Enu-

⁶For example, consider for the binary tree T_2 the loop σ constructed in the proof of Lemma 11, which traverses every edge of T_2 exactly once in each direction. This loop is null-homotopic in $|T_2|$ (Lemma 3), but no sequence of edges, of any order type, is followed immediately by the inverse of that sequence. The edges of T_2 aren’t chords of a spanning tree, but this can be achieved by changing the graph: just double every edge and subdivide the new edges once. The new edges then form a normal spanning tree in the resulting graph G , whose chords are the original edges of our T_2 , and σ is still a (null-homotopic) loop in $|G|$.

merate these as e_0, e_1, \dots , and denote their natural orientations as $\vec{e}_0, \vec{e}_1, \dots$. The circles in $T + e_i$ are the *fundamental circles* of e_i . (Up to orientation and reparametrization, there is a only one such circle for every chord.)

Let us call the elements of the set

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

letters, and two letters \vec{e}_i, \bar{e}_i *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* of (the letters used by) w , such that $w^{-1}(a)$ is finite for every $a \in A$. The only property of S relevant to us is its order type, so two words $w: S \rightarrow A$ 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

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

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 \vec{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 $w = \vec{e}_1 \bar{e}_1 \vec{e}_1$ then $r(w) = \vec{e}_1$, but this letter \vec{e}_1 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 $\vec{e}_0, \vec{e}_1, \dots$ 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 $\{\vec{e}_i \mid i \in I\}$.

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:

Any result of first reducing and then restricting a word can also be obtained by first restricting and then reducing it. (2)

Our aim is to define a reduction map r also for infinite words, so that $(w_1, w_2) \mapsto r(w_1 w_2)$ makes the set of reduced (finite or infinite) words into a group F_∞ . This group F_∞ will contain $\pi_1(|G|)$ as a subgroup by an embedding $\langle \alpha \rangle \mapsto r(w_\alpha)$, where w_α is the word of chords traced out by α .

To define such a reduction map r , we need to show that every infinite word reduces to a unique reduced word. Existence is immediate:

Lemma 12. *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. 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 13. *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 [10] 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:

For all $m \leq n$, we have $R_m = \{\{t, t'\} \in R_n \mid t, t' \in S_m\}$, and the ordering of R_m on this set agrees with that induced by R_n . (3)

Let $s' \in S$ be such that $\{s, s'\} \in R_n$ for some n ; then $\{s, s'\} \in R_n$ for every $n \geq N$, by (3).

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 $p' \not< p$. 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 $\bar{e}_1 \bar{e}_1 \bar{e}_1$ is again an example.

Lemma 13 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 13 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 14. *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 12, 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 [10] 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 13 every position of v and of v' is permanent. \square

With Lemma 14 we are now able to prove:

Lemma 15. *Every word reduces to a unique reduced word.*

Proof. By Lemma 12, 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 14, 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 14, 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_∞ , with multiplication defined as $(w_1, w_2) \mapsto r(w_1 w_2)$, identity the empty word \emptyset , and the inverse w^- of $w: S \rightarrow A$ defined as the map on the same S , but with the inverse ordering, that satisfies $\{w(s), w^-(s)\} = \{\bar{e}_i, \bar{e}_i\}$ for some i for every $s \in S$. (Thus, w^- is w taken backwards and with reverse chord orientations.) Note that the proof of associativity requires an application of Lemma 15.

In [18], we prove that F_∞ embeds canonically in the inverse limit of the groups F_I (which by (2) form an inverse system with respect to restriction), and that $\pi_1(|G|)$ embeds canonically in F_∞ . All we need here is the existence of that particular homomorphism $\pi_1(|G|) \rightarrow F_\infty$, which is defined as follows.

Every path σ in $|G|$ 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 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 .

Lemma 16. *The traces of homotopic paths in $|G|$ reduce to the same word.*

Proof. We first consider homotopic loops $\alpha \sim \beta$ in $|G|$, based at the same vertex or end. We wish to show that $r(w_\alpha) = r(w_\beta)$. By Lemma 14 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 $|G|$ by attaching a 2-cell to $|G|$ for every $j \notin I$, by an injective attachment map from the boundary of the 2-cell onto the unique image of the fundamental circles of e_j . This space deformation retracts onto $\bar{T} \cup \{e_i \mid i \in I\}$, and hence by Lemma 3 onto a finite graph G_I consisting of a subtree of T and the chords e_i with $i \in I$. Composing α and β with the map $|G| \rightarrow G_I$ in which this retraction ends yields homotopic loops α' and β' in G_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 finite graphs, we deduce that

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

This completes the proof of the lemma for loops. The general case follows, since paths joining the same two vertices or ends can be made into loops by appending a path in \bar{T} joining their endpoints, which does not change their traces. \square

Lemma 16 implies in particular that the map $\langle \alpha \rangle \mapsto r(w_\alpha)$ from $\pi_1(|G|)$ to F_∞ is well defined. By (2), it is a homomorphism. In [18] we prove that it is injective—the converse of Lemma 16—and determine its image. All in all, the fundamental group of $|G|$ can be described combinatorially by canonical group embeddings $\pi_1(|G|) \rightarrow F_\infty \rightarrow \varprojlim F_I$ as follows (see [18] for details):

Theorem 17 ([18]). *Let G be a locally finite connected graph. Let T be a normal spanning tree of G , 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 $\pi_1(|G|)$ to the group F_∞ of reduced finite or infinite words in $\{\bar{e}_0, \bar{e}_1, \dots\} \cup \{\bar{e}_0, \bar{e}_1, \dots\}$, with image the set of reduced words whose monotonic subwords 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.)*

5 Distinguishing boundaries from other chains

In this section we wind up our proof of Theorem 8, picking up the thread from Section 3. There we had defined a group homomorphism $f: H_1(|G|) \rightarrow \bar{\mathcal{C}}(G)$, shown that its image is $\bar{\mathcal{C}}(G)$ (Lemmas 10 and 11), and seen that f is injective if G contains only finitely many circuits. The assertion left to prove is that f is not injective if G contains infinitely many circuits, which we now assume.

Let T be a normal spanning tree of G . Each of the infinitely many circuits in G is a finite sum (mod 2) of distinct fundamental circuits of T [10]. Therefore T has infinitely many chords, e_0, e_1, \dots say. Since $|G|$ is compact, there is a sequence $\bar{e}_{i_0}, \bar{e}_{i_1}, \dots$ of chords whose first points converge to an end ω of G . There exists a loop ρ in $|G|$, based at a vertex, that traverses $\bar{e}_{i_0}, \bar{e}_{i_1}, \dots, \bar{e}_{i_0}, \bar{e}_{i_1}, \dots$ in this order and runs otherwise along \bar{T} . (Thus, ρ starts with passes through $\bar{e}_{i_0}, \bar{e}_{i_1}, \dots$, interspersed with finite segments of T between the endpoints of these passes, until it reaches ω , from where it returns along \bar{T} to the starting vertex of \bar{e}_{i_0} ; it then traverses $\bar{e}_{i_0}, \bar{e}_{i_1}, \dots$ interspersed with connecting segments of T to reach ω a second time, and finally returns from there along \bar{T} to its starting vertex. Note that the convergence of e_{i_0}, e_{i_1}, \dots is essential for ρ to be a path: there is no path in $|G|$ through an ω -sequence of chords that does not converge.) Since ρ traverses the chords of T equally often in both directions, Theorem 5 and Lemmas 9 and 10 imply that ρ also traverses the edges of T equally often in both directions.

Hence $f([\rho]) = 0 \in \bar{\mathcal{C}}(G)$. To complete the proof that f is not injective, and thereby the proof of Theorem 8, we show that $[\rho] \neq 0$, i.e. that ρ is not the boundary of any 2-chain. In order to do so, we shall use our results from Section 4 to define an invariant of 1-chains that can distinguish ρ from boundaries. As in Section 4 we consider only paths whose boundary points are vertices or ends, so our invariant will be defined only for chains of 1-simplices with this property. However, it is easy to see that this entails no loss of generality.⁷

We need some more notation. Given $k \in \mathbb{N}$ and a reduced word $w: S \rightarrow A$ (where $A = \{\bar{e}_0, \bar{e}_1, \bar{e}_2, \dots\}$ as before), write $n^+(w, k)$ for the number of

⁷Indeed, if $\rho = \sum \lambda_n \partial \tau_n$ for 2-simplices τ_n , we can modify each τ_n into another 2-simplex τ'_n whose 0-faces are vertices or ends, and such that $\rho = \sum \lambda_n \partial \tau'_n$. For every inner point x of an edge $e_x = u_x v_x$ in $|G|$ pick a fixed path π_x from x to v_x (say). Then append to every 1-simplex σ occurring in the boundary of a τ_n and ending in such a point x the path π_x after x , turning σ into a path σ' between two vertices by appending at most two such paths π_x . Then if $\partial \tau_n = \sigma_0 - \sigma_1 + \sigma_2$, say, it is easy to see that also $\sigma'_0 - \sigma'_1 + \sigma'_2$ is the boundary of a 2-simplex τ'_n . And clearly $\rho = \sum \lambda_n \partial \tau_n$ implies that also $\rho = \sum \lambda_n \partial \tau'_n$, since we modified only 1-simplices that cancelled out anyway in this sum.

intervals in S which can be written as $\{s_0, s_1, \dots\}$ with $s_0 < s_1 < \dots$ and $w(s_j) = \bar{e}_{i_{k+j}}$ for every $j \in \mathbb{N}$. (This number exists: there are at most $|w^{-1}(\bar{e}_{i_k})|$ such intervals, and this number is finite by our definition of ‘word’.) Put

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

(Recall that w^- is w backwards with inverse letters, so $n^+(w^-, k)$ counts the intervals in S which can be written as $\{s_0, s_1, \dots\}$ with $s_0 > s_1 > \dots$ and $w_\sigma(s_j) = \bar{e}_{i_{k+j}}$ for every $j \in \mathbb{N}$.)

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

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

Unlike $N(\varphi, k)$ for fixed k , the function N is not a homomorphism. Nevertheless, it will help us distinguish our special path ρ from boundaries: we shall prove that N vanishes on boundaries, while clearly $N(\rho) = 1$. (Indeed, the word w_ρ is easily seen to be reduced (cf. Lemma 13); hence $N(\rho, k) = n(w_\rho, k) = 1$ for all k , since $n^+(w_\rho, k) = 1$ while $n^+(w_\rho^-, k) = 0$.)

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

$$\begin{aligned} &\text{If } w = w_1 w_2 \text{ is a reduced word, there exists a } k \in \mathbb{N} \text{ such that} \\ &n(w, \ell) = n(w_1, \ell) + n(w_2, \ell) \text{ for all } \ell \geq k. \end{aligned} \quad (4)$$

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

For our proof that N vanishes on boundaries φ , it suffices to show that every 2-simplex τ satisfies $N(\partial\tau, k) = 0$ for large enough k : then $N(\varphi, k) = 0$ for some (large) k , and hence $N(\varphi) = 0$ as claimed. So consider a 2-simplex τ , with boundary $\partial\tau = \sigma_2 - \sigma_1 + \sigma_0$ denoted so that σ_2 ends at the starting vertex of σ_0 . Write $w_i := r(w_{\sigma_i})$ for the words to which the traces of the σ_i reduce ($i = 1, 2, 3$), and $w_{20} := r(w_{\sigma_{20}})$, where $\sigma_{20} := \sigma_2 \sigma_0$ is the path consisting of σ_2 followed by σ_0 .

Note that $w_{20} = r(w_2 w_0)$. Indeed, we can reduce $w_{\sigma_{20}}$ by first applying to $w_{\sigma_2} \subseteq w_{\sigma_{20}}$ the reduction that turns w_{σ_2} into w_2 , and then apply to $w_{\sigma_0} \subseteq w_{\sigma_{20}}$ the reduction that turns w_{σ_0} into w_0 . Together this is a reduction of $w_{\sigma_{20}}$ to $w_2 w_0$. Let R be a reduction of $w_2 w_0$ to $r(w_2 w_0)$. Since we started with $w_{\sigma_{20}}$, the reduced word $r(w_2 w_0)$ we end up with is $r(w_{\sigma_{20}}) = w_{20}$ (Lemma 15).

Let us look at what R does. Since w_2 and w_0 are both reduced, every pair of positions in R has one position in w_2 and the other in w_0 . Hence if w denotes the subword of w_2 whose positions are deleted by R , we have found reduced words w, w'_2, w'_0 such that

$$w_2 = w'_2 w \quad \text{and} \quad w_0 = w^- w'_0 \quad \text{and} \quad w_{20} = w'_2 w'_0.$$

By (4), therefore, we have for all large enough k

$$\begin{aligned} n(w_2, k) &= n(w'_2, k) + n(w, k) \\ n(w_0, k) &= n(w^-, k) + n(w'_0, k) \\ n(w_{20}, k) &= n(w'_2, k) + n(w'_0, k). \end{aligned}$$

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

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

for all these k .

Since σ_{20} is homotopic to σ_1 (across τ), Lemma 16 implies that $w_{20} = w_1$. We therefore deduce that

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

for all large enough k , as desired. This completes the proof of Theorem 8.

6 A new homology for non-compact complexes

In this section we introduce a modification of singular homology for any topological space X embedded in a larger space \hat{X} , which assigns a special role to the points in $\hat{X} \setminus X$. The kind of example we have in mind is that X is locally compact and \hat{X} a compactification of X ; see e.g. [1] for more on such spaces. For this reason we shall call the points in $\hat{X} \setminus X$ the *ends* of X ; but formally we make no assumptions other than that $X \subseteq \hat{X}$. For compact $X = \hat{X}$ our homology will coincide with standard singular homology. When X is a graph and \hat{X} its Freudenthal compactification, the first group of our homology will be canonically isomorphic to the group of the topological cycle space of X .

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

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

- (i) $(\sigma_i \mid i \in I)$ is locally finite in X , that is, every $x \in X$ has a neighbourhood in X that meets the image of σ_i for only finitely many i ;
- (ii) every σ_i maps the 0-faces of Δ^n (the standard n -simplex) to X .

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

When $(\sigma_i \mid i \in I)$ is a good family, any formal linear combination $\sum_{i \in I} \lambda_i \sigma_i$ with all $\lambda_i \in \mathbb{Z}$ is an *n-chain in X* . Since only finitely many σ_i in a good family can be equal, we can add up the coefficients of equal terms and thus assume if desired that the σ_i in an n -chain are pairwise distinct; however, we

shall *not* normally assume this. We write $C_n(X)$ for the group of all n -chains in X , and $C'_n(X)$ for its subgroup of finite n -chains. The boundary operator $\partial_n: C_n \rightarrow C_{n-1}$ is defined as usual by extending linearly from $\partial\sigma_i$. Note that ∂ is well defined (i.e., that it preserves the local finiteness required of chains), and $\partial^2 = 0$. Chains in $\text{Im } \partial$ will be called *boundaries*.

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

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

More precisely, an n -chain $\varphi \in C_n(X)$ shall lie in $Z_n(X)$ if we can write it as $\varphi = \sum_{i \in I} \lambda_i \sigma_i$ in such a way that I admits a partition into finite sets I_j ($j \in J$) with $z_j := \sum_{i \in I_j} \lambda_i \sigma_i \in Z'_n(X)$ for every $j \in J$. Any such representation of φ as a formal sum will be called a *standard representation* of φ as a cycle.⁸ We call the elements of $Z_n(X)$ the *n -cycles* of X .

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

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

as usual. Alternatively, we may study the subgroups

$$H'_n(X) := \{[z] : z \in Z'_n\}$$

of $H_n(X)$ formed by the homology classes of finite cycles.

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

To illustrate all this let us look at a simple example,⁹ the *double ladder*. This is the 2-ended graph G with vertices v_n and v'_n for all integers n , and with edges e_n from v_n to v_{n+1} , edges e'_n from v'_n to v'_{n+1} , and edges f_n from v_n to v'_n . The 1-simplices corresponding to these edges, oriented in their natural directions, are θ_{e_n} , $\theta_{e'_n}$ and θ_{f_n} . For the infinite chains $\varphi := \sum \theta_{e_n}$, $\varphi' := \sum \theta_{e'_n}$ and $\psi := \varphi - \varphi'$ we have $\partial\varphi = \partial\psi = 0$, and neither sum as written above contains a finite cycle. But while it can be shown that $\varphi \notin Z_1(G)$, we have $\psi \in Z_1(G)$, because ψ can be rewritten as $\varphi = \sum z_n$ with finite cycles $z_n = \theta_{e_n} + \theta_{f_{n+1}} - \theta_{e'_n} - \theta_{f_n}$. By contrast, the unique representation of ψ as a formal sum in which like terms are combined, $\psi = \varphi - \varphi'$, is not a standard representation of ψ as a cycle. As we shall see in Section 7, we even have $[\psi] \in H'_1$, although this is not obvious now.

There are some immediate questions that arise from the definitions given above. For example, is Z_n closed under locally finite sums? Is $H'_n = H_n$, i.e., is every homology class represented by a finite cycle? (For $n = 1$ we shall see in

⁸Since the σ_i need not be distinct, φ has many representations as a formal sum. Not all of these need admit a partition as indicated—see below.

⁹Another example will be given in (5), which the reader is invited to skip to now.

the next section that it is.) We shall not pursue these questions here, but will explore the properties of our homology (including more fundamental questions than these) in another paper. In the remainder of this paper we show that our homology achieves its aim: for graphs, it captures precisely the oriented topological cycle space.

7 $H_1(G)$ equals $\vec{\mathcal{C}}(G)$

In this section we show that, for graphs G , the groups $H_1(G)$ and $H'_1(G)$ defined in Section 6 coincide, and are canonically isomorphic to the topological cycle space $\vec{\mathcal{C}}(G)$ of G .

In analogy to our notation of Section 3, we shall denote this isomorphism by $f: H_1(G) \rightarrow \vec{\mathcal{C}}(G)$. In our definition of f we shall have to refer to the map which, in Section 3, was denoted as $f: H_1(|G|) \rightarrow \vec{\mathcal{E}}(G)$; this map will now be denoted as $f': H_1(|G|) \rightarrow \vec{\mathcal{E}}(G)$. (Recall that $|G|$ denotes the Freudenthal compactification of G , and that $H_1(|G|)$ is its usual first singular homology group.¹⁰) When G is finite, our new function f will coincide with f' .

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

To define $f([\varphi])(\vec{e})$, we show that for all large enough finite subchains $\varphi' \in Z'_1(G)$ of φ the values of $f'([\varphi'])(\vec{e})$ agree (the homology class $[\varphi']$ being taken in $H_1(|G|)$), and set $f([\varphi])(\vec{e})$ to this common value. Write I_e for the set of those $i \in I$ whose σ_i meets e ; since e is compact and $(\sigma_i \mid i \in I)$ is a good family, I_e is a finite set.

Let $\pi: H_1(S^1) \rightarrow \mathbb{Z}$ and $f_e: |G| \rightarrow S^1$ be defined as in Section 3, and write $(f_e)_\# : C_1(|G|) \rightarrow C_1(S^1)$ for the chain map induced by f_e .

Lemma 18. *For all finite sets I' such that $I_e \subseteq I' \subseteq I$ and $\varphi' := \sum_{i \in I'} \lambda_i \sigma_i \in Z_1(|G|)$, the values of $f'([\varphi'])(\vec{e})$ agree.*

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

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

for all such φ' , and the result follows.

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

We now define $f: H_1(G) \rightarrow \vec{\mathcal{E}}(G)$ by letting $f([\varphi])$ map an oriented edge \vec{e} to the common value of $f'([\varphi'])(\vec{e})$ for all φ' as in Lemma 18. In order to show

¹⁰We shall use $C_1(G)$, $Z_1(G)$, $B_1(G)$ and $H_1(G)$ to refer to our new homology of $|G|$ that relies on the information of which points of $|G|$ are ends, while $C_1(|G|)$, $Z_1(|G|)$, $B_1(|G|)$ and $H_1(|G|)$ continue to refer to the usual singular homology of the space $|G|$.

that f is well defined, let $\varphi \in Z_1(G)$ and $\psi \in B_1(G)$ be given in any standard representations $\varphi = \sum_{i \in I} \lambda_i \sigma_i$ and $\psi = \sum_{i \in J} \lambda_i \sigma_i$ with $I \cap J = \emptyset$. We show that f assigns the same value to $[\varphi] = [\varphi + \psi]$ no matter whether we base its computation on φ or on $\varphi + \psi$: this proves that $f([\varphi])$ depends neither on the choice of φ as a representative of $[\varphi]$ nor on its representation as $\sum_{i \in I} \lambda_i \sigma_i$.

Given $\vec{e} \in \vec{E}$, let I_e be the set of all $i \in I$ such that σ_i meets e , and define J_e likewise. Let $I' \subseteq I$ and $J' \subseteq J$ be finite sets containing I_e and J_e , respectively, such that $\varphi' := \sum_{i \in I'} \lambda_i \sigma_i \in Z_1(|G|)$ and $\psi' := \sum_{i \in J'} \lambda_i \sigma_i \in B_1(|G|)$; such sets exist since φ and ψ are given in standard representations. Then

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

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

Since f' is a homomorphism with image $\vec{\mathcal{C}}(G)$ (Section 3), Lemma 18 implies that so is f . Indeed, for a proof that $f([\varphi]) \in \vec{\mathcal{C}}(G)$ consider the finite oriented cuts \vec{F} of G , and apply Theorem 5 to any finite subchain φ' of φ containing all the simplices that meet this cut. The proof that f is surjective is the same as in Section 3: every element of $\vec{\mathcal{C}}(G)$ has the form $f([\tau])$ with τ a single loop. Thus in fact,

$$\vec{\mathcal{C}}(G) \subseteq f(H'_1(G)) \subseteq f(H_1(G)) \subseteq \vec{\mathcal{C}}(G)$$

with equality.

The proof of $\text{Im } f \subseteq \vec{\mathcal{C}}(G)$ indicated above uses critically that infinite cycles are of finite character. The following example illustrates that it also depends critically on our rule that chains must be locally finite. Let $G = \mathbb{R}$, with vertex set \mathbb{Z} . For vertices $m < n$ let $\sigma_{m,n} : [0, 1] \rightarrow G$ interpolate linearly between m and n . Now consider the following sum of boundaries, each of the form $(\sigma_{k,m} + \sigma_{m,n} - \sigma_{k,n})$:

$$\begin{aligned} \sum_{i \in \mathbb{Z}} \sigma_{i,i+1} &= \sum_{i \in \mathbb{N}} (\sigma_{-(i+1),i} + \sigma_{i,i+1} - \sigma_{-(i+1),i+1}) \\ &\quad + \sum_{i \in \mathbb{N}} (\sigma_{-(i+2),-(i+1)} + \sigma_{-(i+1),i+1} - \sigma_{-(i+2),i+1}). \end{aligned} \tag{5}$$

The combined sum on the right-hand side of the equation is well defined in that every $\sigma_{m,n}$ occurs at most twice. It also satisfies the requirements for a standard representation of cycles in that it comes partitioned into finite cycles, even boundaries. The left-hand side, which is obtained from the right-hand side by deleting cancelling pairs of simplices, is a well-defined locally finite chain with zero boundary. But it is not (and should not be) an element of $Z_1(G)$. This is not a contradiction only because its apparent ‘standard representation’ as the combined sum on the right is not locally finite, and hence not a legal chain: the point 0, for example, lies in the image of every $\sigma_{-n,m}$.

Our final goal is to show that f is injective. For finite G , the standard proof is to rewrite a given cycle $z \in Z_1(G)$ as a homologous sum of simplexes each traversing exactly one edge. If $[z] \in \text{Ker } f$, every edge is traversed equally often

in both directions, and we can pair up the simplices traversing it accordingly. Each pair is homologous to a boundary, and hence so is z .

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

In our new setup, however, this can indeed be done. In fact, it turns out that our restriction that any boundary chains to be added must be locally finite exactly strikes the balance between being restrictive enough to rule out counterexamples like the one above and being general enough to allow the subdivision into chains of single edges even of complicated cycle like our non-injectivity example from Section 3.

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

Lemma 19. *For every $z \in Z'_1(G)$ there exist a chain $\varphi = \sum_{i \in I} \sigma_i \in Z_1(G)$ and a chain $b \in B_1(G)$ such that $z + b = \varphi$, every σ_i maps $[0, 1]$ homeomorphically to some edge e , and all these edges e as well as the images of the simplices in b are contained in the image of the 1-simplices in z .*

Proof. We may clearly assume that z is an elementary cycle consisting of a single loop τ_0 that is based at a vertex and is not null-homotopic. In particular, τ_0 traverses an edge. Since τ_0 traverses every edge only finitely often (Lemma 4), τ_0 contains only countably many passes through edges, π_1, π_2, \dots say, which we reparametrize as maps from $[0, 1]$.

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

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

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

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

- (i) $\sigma_1, \dots, \sigma_n$ and all τ_j ($j \in J_n$) are paths in $|G|$ between (possibly identical) vertices;
- (ii) if $n \geq 1$, every τ_j ($j \in J_n$) is a segment of some τ_i with $i \in J_{n-1}$;
- (iii) if $n \geq 1$, there exists $j(n)$ such that $J_{n-1} \setminus J_n = \{j(n)\}$ and the finite chain $b_n := \sigma_n - \tau_{j(n)} + \sum_{j' \in J_n \setminus J_{n-1}} \tau_{j'}$ lies in $B_1(G)$;
- (iv) σ_n is homotopic to π_n relative to $\{0, 1\}$;
- (v) suitably reparametrized, $(\pi_{n+1}, \pi_{n+2}, \dots)$ is the family of all edge-passes of the paths τ_j ($j \in J_n$); specifically, the edge-passes in the paths $\tau_{j'}$ with $j' \in J_n \setminus J_{n-1}$ are precisely those in $\tau_{j(n)}$ other than π_n .

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

Let us take a look at the simplices τ_j with $j \in J$. By definition of J , we have $j \in J_n$ for all large enough n . By (i) and (ii), τ_j is a segment of τ_0 between two vertices, and by (v) it contains none of the passes π_1, π_2, \dots . So it does not traverse any edge. Hence,

$$\tau_j \text{ is a null-homotopic loop based at a vertex.} \quad (6)$$

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

Next, let us show that the following:

$$\begin{aligned} & \text{The family of all simplices added or deleted in the construction,} \\ & \text{that is, of all } \sigma_i \text{ } (i \in I), \text{ all } \tau_j \text{ } (j \in J) \text{ and all } \tau_{j(n)}, \text{ is locally} \\ & \text{finite and hence good.} \end{aligned} \quad (7)$$

To prove (7), let x be any point in G . If x is a vertex, let E_x be the set of edges at x ; if $x \in \bar{e}$ for an edge e , let $E_x := \{e\}$. Choose an open neighbourhood U of x contained in $\bigcup E_x$. Since τ_0 traverses each edge in E_x only finitely often, only finitely many of the paths σ_i ($i \in I$) meet U . Similarly, any path τ_j with $j \in J$ that meets U must be based at a vertex incident with an edge in E_x . Since there are only finitely many such vertices, and at each only finitely many τ_j are

based, only finitely many τ_j with $j \in J$ meet U . Finally, consider a path $\tau_{j(n)}$. This path traverses an edge (in π_n), so if it meets U it must also traverse an edge in E_x or adjacent to an edge in E_x . Only finitely many of the passes π_k traverse such edges. By (v), any τ_j containing π_k satisfies $j \in J_1 \cup \dots \cup J_{k-1}$, so $j(n) \in J_1 \cup \dots \cup J_{k-1}$ for the largest such k . Since this is a finite set and the map $n \mapsto j(n)$ is injective, only finitely many n are such that $\tau_{j(n)}$ meets U . This completes the proof of (7).

To complete the proof, we show that $z + b = \varphi$ for $b := \sum_{i \in I} b_i - \sum_{j \in J} \tau_j$, and in particular that $b \in B_1(G)$. By (7), the family of all simplices in b is good, so $b \in B_1(G)$ by (iii) and (6). Likewise, the family of all σ_i is good. Since

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

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

We can now easily complete the proof that our function $f: H_1(G) \rightarrow \vec{\mathcal{C}}(G)$ is injective. Consider any $[z] \in \text{Ker } f$. As $z \in Z_1(G)$, it has a standard representation as $z = \sum_{j \in J} z_j$ with all $z_j \in Z'_1(G)$. By Lemma 19, there are $b_j \in B_1(G)$ ($j \in J$) such that $z_j + b_j = \varphi_j$, where $\varphi_j = \sum_{i \in I_j} \sigma_i$ is a chain of simplices each traversing exactly one edge, and these edges as well as the images of the simplices in b_j lie in the image of z_j . The fact that z is a locally finite chain therefore implies that so are

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

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

For $I := \bigcup_{j \in J} I_j$, we thus have $Z_1 \ni z + b = \varphi = \sum_{i \in I} \sigma_i$, with $b \in B_1(G)$. Since $[z] \in \text{Ker } f$, we thus have $[\varphi] \in \text{Ker } f$. Therefore the loops formed by the elementary cycles in $\varphi = \sum_{i \in I} \sigma_i$ traverse, in total, each edge of G equally often in both directions (Lemma 9). Since each of the σ_i traverses precisely one edge, we can thus pair them up into cancelling pairs $\sigma_i + \sigma_{i'} \in B_1(G)$, where σ_i and $\sigma_{i'}$ traverse the same edge but in opposite directions. Hence $\varphi = \sum_{i \in I} \sigma_i \in B_1(G)$, giving $z = \varphi - b \in B_1(G)$ as desired.

We have thus shown that f is an injective group homomorphism from $H_1(G)$ to $\vec{\mathcal{C}}(G)$ whose restriction to $H'_1(G)$ still maps onto $\vec{\mathcal{C}}(G)$. Hence all these groups coincide, which is our second main result:

Theorem 20. *The function f is a group isomorphism between $H_1(G)$ and $\vec{\mathcal{C}}(G)$, as well as between $H'_1(G)$ and $\vec{\mathcal{C}}(G)$. In particular, $H'_1(G) = H_1(G)$.*

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] H. Bruhn, *The cycle space of a 3-connected locally finite graph is generated by its finite and infinite peripheral circuits*, J. Combin. Theory (Series B) **92** (2004), 235–256.
- [4] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. Comput. **15** (2006), 75–90.
- [5] H. Bruhn, R. Diestel, and M. Stein, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, J. Graph Theory **50** (2005), 150–161.
- [6] H. Bruhn and A. Georgakopoulos, *Bases and closures with infinite sums*, preprint 2006.
- [7] H. Bruhn and M. Stein, *On end degrees and infinite circuits in locally finite graphs*, to appear in *Combinatorica*.
- [8] ———, *MacLane’s planarity criterion for locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 225–239.
- [9] R. Diestel, *The cycle space of an infinite graph*, Comb., Probab. Comput. **14** (2005), 59–79.
- [10] ———, *Graph Theory* (3rd edition), Springer-Verlag, 2005,
Electronic edition available at:
<http://www.math.uni-hamburg.de/home/diestel/books/graph.theory>.
- [11] ———, *End spaces and spanning trees*, J. Combin. Theory (Series B) **96** (2006), 846–854.
- [12] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory (Series B) **87** (2003), 197–206.
- [13] ———, *On infinite cycles I*, *Combinatorica* **24** (2004), 68–89.
- [14] ———, *On infinite cycles II*, *Combinatorica* **24** (2004), 91–116.
- [15] ———, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.
- [16] R. Diestel and I. Leader, *A proof of the Bounded Graph Conjecture*, *Invent. math.* **108** (1992), 131–162.
- [17] ———, *Normal spanning trees, Aronszajn trees and excluded minors*, J. London Math. Soc. **63** (2001), 16–32.
- [18] R. Diestel and P. Sprüssel, *The fundamental group of locally finite graphs with ends*, preprint 2008.
- [19] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, *Math. Zeitschr.* **33** (1931), 692–713.
- [20] ———, *Neuaufbau der Endentheorie*, *Annals of Mathematics* **43** (1942), 261–279.

- [21] D.B. Fuchs and O.Ya. Viro, *Topology II*, Springer-Verlag, 2004.
- [22] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, preprint 2007.
- [23] A. Georgakopoulos and P. Sprüssel, *Geodesic topological cycles in locally finite graphs*, preprint 2006.
- [24] R. Halin, *Über unendliche Wege in Graphen*, Math. Annalen **157** (1964), 125–137.
- [25] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.
- [26] A. Hatcher, *Algebraic Topology*, Cambridge Univ. Press, 2002.
- [27] B. Hughes and A. Ranicki, *Ends of complexes*, Cambridge Univ. Press, 1996.
- [28] H.A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), 1–22.
- [29] ———, *Connectivity in infinite graphs*, Studies in Pure Mathematics (L. Mirsky, ed.), Academic Press, 1971.
- [30] B. Krön, *End compactifications in non-locally-finite graphs*, Math. Proc. Cambridge Phil. Soc. **131** (2001), 427–443.
- [31] A.T. Lundell and S. Weingram, *Topology of CW-complexes*, Springer-Verlag, 1969.
- [32] R. Möller, *Ends of graphs*, Math. Proc. Cambridge Phil. Soc. **111** (1992), 255–266.
- [33] ———, *Ends of graphs II*, Math. Proc. Cambridge Phil. Soc. **111** (1992), 455–460.
- [34] M. Stein, *Arboriticity and tree-packing in locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 302–312.
- [35] C. Thomassen and A. Vella, *Graph-like continua and Menger’s theorem*, preprint 2006.
- [36] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, J. Combin. Theory (Series B) **58** (1993), 248–268.
- [37] W. Woess, *Random walks on infinite graphs and groups*, Cambridge University Press, 2002.

Reinhard Diestel
 <diestel@math.uni-hamburg.de>
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 20146 Hamburg
 Germany

Philipp Sprüssel
 <Philipp.Spruessel@gmx.de>
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 20146 Hamburg
 Germany