

Locally finite graphs with ends: a topological approach.

III. Fundamental group and homology

Reinhard Diestel and Philipp Sprüssel

Abstract

This paper is the last part of a comprehensive survey of a newly emerging field: a topological approach to the study of locally finite graphs that crucially incorporates their ends. Topological arcs and circles, which may pass through ends, assume the role played in finite graphs by paths and cycles. The first two parts of the survey together provide a suitable entry point to this field for new readers; they are available in combined form from the ArXiv [2].

The topological approach indicated above has made it possible to extend to locally finite graphs many classical theorems of finite graph theory that do not extend verbatim. While the first part [3] of this survey introduces the theory as such and the second part [4] is devoted to those applications, this third part looks at the theory from an algebraic-topological point of view.

The results surveyed here include both a combinatorial description of the fundamental group of a locally finite graph with ends and the homology aspects of this space.

1 Introduction

The survey [2] describes a topological framework in which many well-known theorems about finite graphs that appear to fail for infinite graphs do have a natural infinite analogue. It has been realised in recent years that many such theorems, especially about paths and cycles, work in a slightly richer setting: not in the (locally finite) graph G itself, but in its compactification $|G|$ obtained by adding its *ends*.¹ In this setting, the traditional cycle space of a graph is replaced by its *topological cycle space*. The topological cycle space $\mathcal{C} = \mathcal{C}(G)$ of a locally finite graph G is based on (the edge sets of) *topological circles* in $|G|$, homeomorphic images of the unit circle S^1 , allowing infinite sums as long as they are *thin*, that is, every edge appears in only finitely many summands. Since the topological cycle space $\mathcal{C}(G)$ was introduced [5, 6], it has proved surprisingly successful; see [2, 4] for numerous applications.

Given the success of \mathcal{C} for graphs, it seems desirable to recast its definition in homological terms that make no reference to the one-dimensional character of $|G|$ (e.g., to circles), to obtain a homology theory for similar but more general spaces (such as non-compact CW complexes of any dimension) that implements the ideas and advantages of \mathcal{C} more generally. This approach has been pursued

¹For a formal definition of $|G|$ see [2].

in [7, 9, 8]. In this paper we present its main ideas, results and examples. For simplicity, all our coefficients will be taken from \mathbb{F}_2 .

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

The next candidate for the desired description of \mathcal{C} in terms of homology is singular homology. Indeed, \mathcal{C} is built from circles in $|G|$, and circles are singular 1-cycles that generate the first singular homology group $H_1(|G|)$ of $|G|$, so both groups are built from similar elements. On the face of it, it is not clear whether \mathcal{C} might in fact be isomorphic, even canonically, to $H_1(|G|)$. However, it will turn out that it is not: in [9] we prove that \mathcal{C} is always a natural quotient of $H_1(|G|)$, and this quotient is proper unless G is essentially finite. This may seem surprising, since \mathcal{C} is defined via (thin) infinite sums while all sums in the definition of $H_1(|G|)$ are finite, which suggests that \mathcal{C} might be larger than $H_1(|G|)$.

Our approach for the comparison of \mathcal{C} and $H_1(|G|)$ will be to define a homomorphism from $Z_1(|G|)$ to the edge space \mathcal{E} that counts how often the edges of G are traversed by the simplices of a 1-cycle z , and maps z to the set of those edges that are traversed an odd number of times. It will turn out that this homomorphism vanishes on boundaries and that its image is precisely \mathcal{C} . Hence it defines an epimorphism $f: H_1(|G|) \rightarrow \mathcal{C}(G)$. However, we will show that f is not normally injective. Indeed, there will be loops that traverse every edge evenly often (even equally often in either direction), but which can be shown not with some effort to be null-homologous. Thus, \mathcal{C} is a genuinely new object, also from a topological point of view.

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

One of the main aims of this project, therefore, became to develop a combinatorial description of the fundamental group of the space $|G|$ for an arbitrary connected locally finite graph G . In [7] 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 with arbitrary spanning trees but only with *topological spanning trees*. Moreover, we will have to allow infinite words of any countable order type, and likewise allow the reduction sequences cancelling adjacent inverse letters to have arbitrary countable order type. However, these reductions can also 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 as a subgroup 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_∞ .

Having proved that \mathcal{C} is usually a proper quotient of $H_1(|G|)$, the last aim of this project then was to define a variant of singular homology that works in more general spaces, and which for graphs captures precisely \mathcal{C} . First steps in this direction were taken in [9]; it was completed in [8]. Our hope with this translation was to stimulate further work in two directions. One is that its new topological guise should make the cycle space accessible to topological methods that might generate some windfall for the study of graphs. And conversely, that as the approach that gave rise to \mathcal{C} is made accessible for more general spaces—in particular, for CW complexes of higher dimensions—its proven usefulness for graphs might find some more general topological analogues.

The key to the definition of \mathcal{C} , and to its success, is that it treats ends differently from other points. To preserve this feature, our new homology theory is constructed for locally compact Hausdorff spaces X with a fixed Hausdorff compactification \hat{X} , in which the compactification points play the role of ends.

2 Čech homology

The Čech homology of a space is an alternative to singular homology for spaces that are not simplicial complexes. For a general space X , the n th Čech homology group $\check{H}_n(X)$ is the inverse limit of the homology groups of simplicial complexes induced by open covers of X .² In the case of $X = |G|$, one can compute the groups $\check{H}_n(X)$ more directly. To do so, fix a normal spanning tree T of G , with root r say, and denote the subtree of T induced by the first i levels by T_i . Let G_i be the finite graph obtained from G by contracting each component of $G - T_i$; then $\check{H}_n(X)$ is the inverse limit of the family $(H_n(G_i), \leq)_{i \in \mathbb{N}}$. Since $\mathcal{C}(G)$ is the inverse limit of the groups $H_1(G_i)$, we have

Theorem 2.1 ([9]). *For a locally finite graph G we have a canonical isomorphism $\check{H}_1(|G|) \simeq \mathcal{C}(G)$.*

Theorem 2.1 shows that one can describe the topological cycle space in terms of the Čech homology. However, although $\check{H}_1(|G|)$ is isomorphic to $\mathcal{C}(G)$ as a group, it does not sufficiently reflect the combinatorial properties of $\mathcal{C}(G)$, its interaction with the combinatorial structure of G . To make this precise, note that a number of classical results about the cycle space say which circuits generate it—as do the non-separating chordless circuits in a 3-connected graph, say. In the Čech homology, however, it is not possible to decide whether a given homology class in $\check{H}_1(|G|)$ corresponds to a circuit. Indeed, the obvious relation between $\check{H}_1(|G|)$ and the combinatorial structure of G is that every homology class $c \in \check{H}_1(|G|)$ corresponds to a family (c_n) of homology classes in the groups $H_1(G_n)$. One might think that the class c should correspond to a circuit in $|G|$ if and only if every c_n with sufficiently large n corresponds to a circuit in G_n . But this is not the case: the limit of a sequence of cycle space elements in

²See [9] for a formal definition.

the G_n can be a circuit even if the elements of the sequence are not circuits in the G_n .




Figure 1: The graph G (drawn twice) with a normal spanning tree T and a circuit c .

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




Figure 2: The edge sets c_4 in G_4 and c_{10} in G_{10} .

The edge set c from Figure 1 is a circuit, but each edge set c_n it induces on a contracted graph G_n with $n = 6k + 4$ is not a circuit (Figure 2). Indeed, each G_{6k+4} consists of $G[V(T_{6k+4})]$, for each i with $1 \leq i \leq k$ a vertex v_i^{6k+4} corresponding to a contracted tail of the ladder L_i , and a vertex v_0^{6k+4} corresponding to the contracted tail of the wide ladder and all ladders L_j with $j > k$. The edge set c_{6k+4} is not a circuit since it has degree 4 at v_0^{6k+4} . Therefore, c is a circuit although it is the limit of the non-circuits c_{6k+4} .

One can easily manipulate the example so that no c_n with n large enough is a circuit by attaching copies H_1, \dots, H_5 of G to G by connecting the vertices of the first rung of the wide ladder in H_i to some suitable vertices of L_i .

3 Singular homology

A more subtle approach than Čech homology, which has been pursued in [9], is to see to what extent $\mathcal{C}(G)$ can be captured by the singular homology of $|G|$. After all, $\mathcal{C}(G)$ was defined via (the edge sets of) circles in $|G|$, which are just injective singular loops. Can we extend this correspondence between injective loops and circuits to one between $H_1(|G|)$ (singular) and $\mathcal{C}(G)$?

There are two things to notice about $H_1(|G|)$. The first is that we can subdivide a 1-simplex (or concatenate two 1-simplices into one by the inverse procedure) by adding a boundary. Indeed, if $\sigma: [0, 1] \rightarrow |G|$ is a path in $|G|$ from x to y , say, and z is a point on that path, there are paths σ' from x to z and σ'' from z to y such that $\sigma' + \sigma'' - \sigma$ is the boundary of a singular 2-simplex ‘squeezed’ on to the image of σ . The second fact to notice is that inverse paths cancel in pairs: if σ^+ is an x - y path in $|G|$, and σ^- an y - x path with the same image as σ^+ , then $[\sigma^+ + \sigma^-] = 0 \in H_1$.³ These two facts together imply that every homology class in H_1 is represented by a single loop: given any 1-cycle, we first add pairs of inverse paths between the endpoints of its simplices to make its image connected in the right way, and then use Euler’s theorem to concatenate the 1-simplices of the resulting chain into a single loop σ . Moreover, we may assume that this loop is based at a vertex.

To establish the desired correspondence between $H_1(|G|)$ and $\mathcal{C}(G)$, we would like to assign to a homology class in $H_1(|G|)$, represented by a single loop σ , an edge set $f([\sigma]) \in \mathcal{C}(G)$. Intuitively, we do this by counting for each edge e of G how often σ traverses it entirely (which, since the domain of σ is compact, is a finite number of times), and let $f([\sigma])$ be the set of those edges e for which this number is odd. Using the usual tools of homology theory, one can make this precise in such a way that f is clearly a well defined homomorphism $H_1(|G|) \rightarrow \mathcal{E}(G)$,⁴ and whose image is easily seen to be $\mathcal{C}(G)$. What is not clear at once is whether f is 1–1 and onto.

Surprisingly, f is indeed surjective—and this is not even hard to show. Indeed, let an edge set $D \in \mathcal{C}(G)$ be given. Our task is to find a loop σ that traverses every edge in D an odd number of times, and every other edge of G an even number of times. As a first approximation, we let σ_0 be a path that traverses every edge of some fixed normal spanning tree of G exactly twice, once in each direction; see [2, Sec. 3.3] for how to construct such a loop. Moreover, we construct σ_0 in such a way that it pauses at every vertex v —more precisely, so that $\sigma_0^{-1}(v)$ is a union of finitely many closed intervals at least one of which is non-trivial. Next, we write D as a thin sum $D = \sum_i C_i$ of circuits; such a representation of D exists by definition of $\mathcal{C}(G)$. For each of these C_i we pick a vertex $v_i \in \overline{C_i}$, noting that no vertex of G gets picked infinitely often, because it has only finitely many incident edges and the C_i form a thin family. Finally, we turn σ_0 into the desired loop σ by expanding the pause at each vertex v to a loop going once round every $\overline{C_i}$ with $v = v_i$. It is not hard to show that σ is continuous [9], and clearly it traverses every edge of G the desired number of times.

³To see that this sum is a boundary, subtract the constant 1-simplex σ with value x : there is an obvious singular 2-simplex of which $\sigma^+ + \sigma^- - \sigma$ is the boundary. Subtracting σ is allowed, since $\sigma = \sigma + \sigma - \sigma$, too, is a boundary: of the constant 2-simplex with value x .

⁴For each edge e , let $f_e: |G| \rightarrow S^1$ be a map wrapping e once round S^1 and mapping all of $|G| \setminus \dot{e}$ to one point of S^1 . Let π denote the group isomorphism $H_1(S^1) \rightarrow \mathbb{F}_2$. Given $h \in H_1(|G|)$, let $f(h) := \{e \mid (\pi \circ (f_e)_*)(h) = 1 \in \mathbb{F}_2\}$. See [9] for details.

Equally surprisingly, perhaps, f is usually not injective (see below). In summary, therefore, the topological cycle space $\mathcal{C}(G)$ of G is related to the first singular homology group of G as follows:

Theorem 3.1 ([9]). *The map $f: H_1(|G|) \rightarrow \mathcal{E}(G)$ is a group homomorphism onto $\mathcal{C}(G)$, which has a non-trivial kernel if and only if G contains infinitely many (finite) circuits.*

An example of a non-null-homologous loop in $|G|$ whose homology class maps to the empty edge set $\emptyset \in \mathcal{C}(G)$ is easy to describe. Let G be the one-way infinite ladder L (with its end on the right), and define a loop ρ in L , as follows. We start at time 0 at the top-left vertex, v_0 say, and begin by going round the first square of G in a clockwise direction. This takes us back to v_0 . We then move along the horizontal edge incident with v_0 , to its right neighbour v_1 . From here, we go round the second square in a clockwise direction, back to v_1 and on to its right neighbour v_2 . We repeat this move until we reach the end ω of G on the right, say at time $\frac{1}{2} \in [0, 1]$. So far, we have traversed the first vertical edge and every bottom horizontal edge once (in the direction towards v_0), every other vertical edge twice (once in each direction), and every top horizontal edge twice in the direction towards the end. From there, we now use the remaining half of our time to go round the infinite circle formed by the first vertical edge and all the horizontal edges one and a half times, in such a way that we end at time 1 back at v_0 and have traversed every edge of L equally often in each direction. Clearly, f maps (the homology class of) this loop ρ to $0 \in \mathcal{C}(L)$.




Figure 3: The loop ρ is not null-homologous, but $f([\rho]) = \emptyset$.

The loop ρ is indeed not null-homologous [9], but it seems non-trivial to show this. To see why this is hard, let us compare ρ to a loop winding round a finite ladder in a similar fashion, traversing every edge once in each direction. Such a loop σ is still not null-homotopic, but it is null-homologous. To see this, we subdivide it into single edges: we find a finite collection of 1-simplices σ_i , four for every edge on the top and two for every other edge, such that $[\sigma] = [\sum_i \sigma_i]$ and every σ_i just traverses its edge. Next, we pair up these σ_i into cancelling pairs: if σ_i and σ_j traverse the same edge e (in opposite directions), then $[\sigma_i + \sigma_j] = 0$. Hence $[\sigma] = [\sum_i \sigma_i] = 0$, as claimed. But we cannot imitate this proof for ρ and the infinite ladder, because homology classes in $H_1(|G|)$ are still finite chains: we cannot add infinitely many boundaries to subdivide ρ infinitely often.

As it happened, the proof of the seemingly simple fact that ρ is not null-homologous took a detour via the solution of a much more fundamental problem: the problem of understanding the fundamental group of $|L|$, or more generally, of $|G|$ for a locally finite graph G . In order to distinguish ρ from boundaries, we looked for a numerical invariant Λ of 1-chains that was non-zero on ρ but both linear and additive (so that $\Lambda(\sigma_1 \sigma_2) = \Lambda(\sigma_1 + \sigma_2) = \Lambda(\sigma_1) + \Lambda(\sigma_2)$ for concatenations of 1-simplices σ_1, σ_2) and invariant under homotopies (so that

$\Lambda(\sigma_1\sigma_2) = \Lambda(\sigma)$ when $\sigma \sim \sigma_1\sigma_2$). Then, given a 2-simplex τ with boundary $\partial\tau = \sigma_1 + \sigma_2 - \sigma$, we would have $\Lambda(\partial\tau) = \Lambda(\sigma_1\sigma_2) - \Lambda(\sigma) = 0$, so Λ would vanish on all boundaries but not on ρ . We did not quite find such an invariant Λ , but a collection of similar invariants which, together, can distinguish loops like ρ from boundaries.

4 The fundamental group of $|G|$

In this section we will sketch the combinatorial description of $\pi_1(|G|)$ given in [7]. Our description involves infinite words and their reductions in a ‘continuous’ setting, and embedding the group they form as a subgroup of a limit of finitely generated free groups.

Let G be a locally finite connected graph, fixed throughout this section, and let T be a topological spanning tree of $|G|$. When G is finite, then $\pi_1(|G|) = \pi_1(G)$ is the free group F on the set $\{e_0, \dots, e_n\}$ of chords of any fixed spanning tree. The standard description of F is given in terms of reduced words of those oriented chords, where reduction is performed by cancelling adjacent inverse pairs of letters such as $\vec{e}_i\vec{e}_i$ or $\vec{e}_i\vec{e}_i$. The map assigning to a path in $|G|$ the sequence of (oriented) chords it traverses defines the canonical group isomorphism between $\pi_1(|G|)$ and F ; in particular, reducing the words obtained from homotopic paths yields the same reduced word.

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

Consider the following example. Let G be the infinite ladder, with a topological spanning tree T consisting of one side of the ladder, all its rungs, and its unique end ω (Figure 4). The path running along the bottom side of the ladder and back is a null-homotopic loop. Since it traces the chords $\vec{e}_0, \vec{e}_1, \dots$ all the way to ω and then returns the same way, the infinite word $\vec{e}_0\vec{e}_1\dots\vec{e}_1\vec{e}_0$ should reduce to the empty word. But it contains no cancelling pair of letters, such as $\vec{e}_i\vec{e}_i$ or $\vec{e}_i\vec{e}_i$.




Figure 4: The infinite ladder and its topological spanning tree T (bold edges)

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

Example 4.1. We construct a locally finite graph G and a null-homotopic loop σ in $|G|$ whose trace of chords contains no cancelling pair of subsequences, of any order type.

Let T be the binary tree with root r . Like in [2, pp. 30–31] we can construct a loop σ in $|T|$ that traverses every edge of T once in each direction, see Figure 5.




Figure 5: A loop running twice through each edge of the binary tree.

The loop σ is easily seen to be null-homotopic. It is also easy to check that no sequence of passes of σ through the edges of T is followed immediately by the inverse of this sequence.

The edges of T are not chords of a topological spanning tree, but this can be achieved by changing the graph: just double every edge.⁵ The new edges together with all vertices and ends then form a topological spanning tree in the resulting graph G , whose chords are the original edges of our tree T , and σ is still a (null-homotopic) loop in $|G|$.

Example 4.1 shows that there is no hope of capturing homotopies of loops in terms of word reduction defined recursively by cancelling pairs of inverse adjacent subwords, finite or infinite. We shall therefore define the reduction of infinite words differently, though only slightly. We shall still cancel inverse letters in pairs, even only one at a time, and these reduction ‘steps’ will be ordered linearly (rather unlike the simultaneous dissolution of all the chords by the homotopy in the example). However, the reduction steps will not be well-ordered.

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

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

⁵And subdivide the new edges once, in case you prefer to obtain a simple graph instead of a graph with multiple edges.

A *word* is a map

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

(the letter \bar{e}_i being the inverse of \bar{e}_i), where S is a totally ordered (countable) set, the set of *positions* of (the letters used by) w , and every letter has only finitely many preimages in S . A *reduction* of a word w is a totally ordered set R of disjoint pairs of positions of w such that the positions in each pair are mapped to inverse letters and are adjacent in the word obtained from w by deleting all (positions of) letters contained in earlier pairs in R . We say that w *reduces to* the word $w \upharpoonright (S \setminus \bigcup R)$. If w has no nonempty reduction, we call it *reduced*. Note that neither the set S of positions of a word w nor a reduction of w have to be well-ordered.

It was shown in [7] that every word w reduces to a unique word $r(w)$ ⁶ and hence the reduced words form a group F_∞ . It was also shown that F_∞ embeds canonically in the inverse limit of the groups F_n , the free groups on the sets $\{e_0, \dots, e_n\}$.

On the other hand, the fundamental group of $|G|$ embeds in F_∞ : Mapping a homotopy class $\langle \alpha \rangle$ to the word $r(w_\alpha)$, where w_α is the *trace* of α , the word induced by the passes of α through the chords of T (with their natural order given by α), turns out to be well-defined; in other words, the traces of homotopic loops reduce to the same word. The harder part is to show the converse: that two loops are homotopic whenever their traces reduce to the same word. In [7], it was shown that the homotopy can even be chosen so that it contracts pairs of passes, one at a time, like known from finite graphs.

The map $\langle \alpha \rangle \mapsto r(w_\alpha)$ is not normally surjective. For example, $\bar{e}_0 \bar{e}_1 \dots$ will always be a reduced word, but no loop in $|G|$ can pass through these chords in precisely this order if they do not converge to an end. Hence if there is a non-converging sequence of chords—which is the case whenever there are two ends of G with no contractible neighbourhood in $|G|$ —then the reduced word $\bar{e}_0 \bar{e}_1 \dots$ lies outside the image of our map $\langle \alpha \rangle \mapsto r(w_\alpha)$.

In order to describe the image of this map precisely, let us call a word $w: S \rightarrow A$ *monotonic* if there is an enumeration s_0, s_1, \dots of S such that either $s_0 < s_1 < \dots$ or $s_0 > s_1 > \dots$. Let us say that w *converges* if the sequence of chords corresponding to its sequence $w(s_0), w(s_1), \dots$ of letters converges. If w is the trace of a loop in $|G|$, then by the continuity of this path all the monotonic subwords of w —and hence those of $r(w)$ —converge. It was shown in [7] that the converse is also true: A reduced word is the trace of a loop in $|G|$ if and only if all its monotonic subwords converge.

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

Theorem 4.2 ([7]). *Let G be a locally finite connected graph, let T be a topological 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 $\{\bar{e}_0, \bar{e}_1, \dots\} \cup \{\bar{e}_0, \bar{e}_1, \dots\}$. Its image consists of those reduced words whose monotonic subwords all converge in $|G|$.*

⁶Unique as an abstract word, not as a restriction of w : The word $\bar{e}_0 \bar{e}_0 \bar{e}_0$, for example, reduces to \bar{e}_0 , but this letter can have the first or the last position in the original word

- (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.)*

Theorem 4.2 provides an interesting interaction between the topological cycle space of G and the fundamental group of $|G|$: It is a well-known fact that the first (singular) homology group of a space is the abelianization of its fundamental group. For graphs, this yields that the (classical) cycle space of G is the abelianization of $\pi_1(G)$. Theorem 4.2 implies an analogous result for the topological cycle space: It is the *strong abelianization* of $\pi_1(|G|)$ [11, Theorem 6.19], the quotient of $\pi_1(|G|)$ obtained by factoring out all words in which every letter appears as often as its inverse.

5 An ad-hoc homology for locally compact spaces

In this section we take up the thread of defining $\mathcal{C}(G)$ in terms of homology. We have seen that Čech homology—although its first group is isomorphic to the topological cycle space—fails to properly reflect its relation to the combinatorial structure of G . For this reason, we shall keep at our singular approach to define \mathcal{C} in terms of homology. Since by Theorem 3.1 standard singular homology is not the right theory to capture $\tilde{\mathcal{C}}$, we shall define a singular-type homology that does so.

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

Loops like the one in Figure 3 suggest that our homology should allow to subdivide a 1-simplex infinitely often: Then, every 1-chain in $|G|$ will be homologous to the sum of its passes through edges of G , and hence it will be null-homologous if and only if it lies in the kernel of f . The idea is thus to define the homology so that we obtain essentially the same 1-cycles as in standard singular homology but more boundaries.

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

This suggests to require a similar property from the chains in our homology: They will have to be locally finite in G but not at ends.⁷ This will enable us to subdivide paths in $|G|$ infinitely often, but the required local finiteness in G will keep us from obtaining undesired cycles, such as the edges of a double-ray

⁷The formal definition of ‘locally finite’ will be given shortly.

(all directed the same way), which has zero boundary but does not correspond to an element of the cycle space. In the ad-hoc homology we shall define in this section we will rule out such cycles by imposing an additional condition on cycles. This will lead to the desired result in dimension 1, i.e. our first homology group will be \mathcal{C} , but generate problems elsewhere. More precisely, this homology will fail to satisfy the Eilenberg-Steenrod axioms for homology, which is caused precisely by this restriction on cycles.

In [8] we thus change our approach slightly: Instead of restricting the group of cycles we define chains differently, so as to obtain 1-cycles that are essentially finite and 2-cycles that allow us to subdivide 1-simplices infinitely often. This homology theory then satisfies the axioms [8]. On the other hand, the proof that this homology theory specializes in dimension 1 to yield \mathcal{C} relies on the corresponding result for the ad-hoc homology defined in this section. Moreover, it introduces some of the main ideas from [8] in a technically simpler setting.

Let X be a locally compact Hausdorff space and let \hat{X} be a Hausdorff compactification of X . (See e.g. [1] for more on such spaces.) Note that every locally compact Hausdorff space is Tychonoff, and thus has a Hausdorff compactification. Although we do not make any assumptions on the type of the compactification, apart from being Hausdorff, we will call the points in $\hat{X} \setminus X$ *ends*, even if they are not ends in the usual, more restrictive, sense.

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

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

Note that as X is locally compact, (i) is equivalent to asking that every compact subspace of X meets the image of σ_i for only finitely many i . Condition (ii), like (i), underscores that ends are not treated on a par with the points in X : we allow them to occur on infinitely many σ_i (which (i) forbids for points of X), but not in the fundamental role of images of 0-faces: all simplices must be ‘rooted’ in X .

When $(\sigma_i \mid i \in I)$ is an admissible family of n -simplices, any formal linear combination $\sum_{i \in I} \lambda_i \sigma_i$ with all $\lambda_i \in \mathbb{Z}$ is an *n -sum in X* .⁸ We regard n -sums $\sum_{i \in I} \lambda_i \sigma_i$ and $\sum_{j \in J} \mu_j \tau_j$ as *equivalent* if for every n -simplex ρ we have $\sum_{i \in I, \sigma_i = \rho} \lambda_i = \sum_{j \in J, \tau_j = \rho} \mu_j$. Note that these sums are well-defined since an n -simplex can occur only finitely many times in an admissible family. We write $C_n(X)$ for the group of *n -chains*, the equivalence classes of n -sums. The elements of an n -chain are its *representations*. Clearly every n -chain c has a unique representation whose simplices are pairwise distinct—which we call the *reduced representation* of c —, but we shall consider other representations too. The subgroup of $C_n(X)$ consisting of those n -chains that have a finite representation is denoted by $C'_n(X)$.

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

⁸In standard singular homology, one does not usually distinguish between formal sums and chains. It will become apparent soon why we have to make this distinction.

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 it has a representation $\sum_{i \in I} \lambda_i \sigma_i$ for which I admits a partition into finite sets I_j ($j \in J$) such that, for every $j \in J$, the n -chain $z_j \in C'_n(X)$ represented by $\sum_{i \in I_j} \lambda_i \sigma_i$ lies in $Z'_n(X)$. Any such representation of φ as a formal sum will be called a *standard representation of φ as a cycle*.⁹ 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 $\sum_{j \in J} \lambda_j z_j$ where each z_j is the (finite) boundary of a singular $(n+1)$ -simplex. We therefore have homology groups

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

as usual.

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

Let us look at an example which might indicate whether we obtain the desired cycles in order to capture the topological cycle space. Consider the *double ladder*. This is the 2-ended graph G with vertices v_n and v'_n for all integers n , and with edges e_n from v_n to v_{n+1} , edges e'_n from v'_n to v'_{n+1} , and edges f_n from v_n to v'_n . The 1-simplices corresponding to these edges, oriented in their natural directions, are θ_{e_n} , $\theta_{e'_n}$, and θ_{f_n} , see Figure 6.




Figure 6: The 1-chains φ and φ' in the double ladder.

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

⁹Since the σ_i need not be distinct, φ has many representations by formal sums. Not all of these need admit a partition as indicated—an example will be given later in the section.

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

By contrast, one can show that $\varphi \notin Z_1(G)$ if \hat{G} is the Freudenthal compactification of G . This is proved in [9], but is not obvious. For example, one might try to represent φ as $\varphi = \sum_{n=1}^{\infty} z'_n$ with $z'_n := \theta_{e_{-n}} + \theta_{n-1} + \theta_{e_n} - \theta_n$, where $\theta_n : [0, 1] \rightarrow e_{-n} \cup \dots \cup e_n$ maps 0 to v_{-n} and 1 to v_{n+1} , see Figure 7.




Figure 7: Finite cycles summing to φ —by an inadmissible sum.

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

This homology indeed captures the cycle space [9]. To see this, note that since infinite chains are allowed, we can add infinitely many boundaries to a loop like in Figure 3 so as to subdivide it into its edge passes. Note that the family of boundaries we add has to be locally finite and it is not obvious that this can always be satisfied. (See [9] for how to choose the boundaries.) Therefore, two chains are homologous if both of them traverse each edge of G the same number of times. Together with the fact that the homomorphism f from the first singular homology group $H_1(|G|)$ to $\mathcal{C}(G)$ can be extended to a homomorphism $H_1(G) \rightarrow \mathcal{C}(G)$ [9], this implies that $H_1(G)$ and $\mathcal{C}(G)$ are isomorphic.

Theorem 5.1 ([9]). *If G is a locally finite graph and $\hat{G} = |G|$, then $H_1(G)$ is canonically isomorphic to $\mathcal{C}(G)$.*

Note that it does not suffice to require the chains to be locally finite without any further assumptions, as it is the case for the *locally finite homology* defined in [10]: This homology does *not* capture the cycle space. Indeed, applied to $|G|$ it yields the usual singular homology, since every locally finite chain in a compact space is finite. On the other hand, applied to G , the locally finite homology allows for chains like φ above, which do not correspond to an element of the cycle space.

As mentioned before, the ad-hoc homology defined above does not satisfy the Eilenberg-Steenrod axioms for homology. (For an example, as well as a listing of the axioms, see [8].) This is caused by the fact that the cycles are not chosen to be the entire kernel of ∂ but with the additional property that they are a locally finite sum of finite cycles.

For this reason, we develop in [8] a homology that does satisfy the axioms and that is defined without further assumptions on the cycles. Like before, we define this homology for locally compact Hausdorff spaces X with a fixed Hausdorff compactification \hat{X} . For this homology to capture $\mathcal{C}(G)$ we have to allow infinite chains, since chains like (the chain consisting of) the loop in

Figure 3 have to be null-homologous in our homology—as they correspond to the empty edge set in G —but are not the boundary of a finite chain. On the other hand, we cannot allow all locally finite chains, as this would yield the locally finite homology mentioned above. The solution to this dilemma is surprisingly simple: We allow only those simplices to appear infinitely often in a chain that are needed to subdivide a path, or more generally, a simplex. This will enable us to subdivide simplices into their edge passes and the isomorphism between our new homology and $\mathcal{C}(G)$ will follow like for the ad-hoc homology above.

A main feature of the simplices whose boundaries we need to subdivide a path σ is that they are in a sense ‘one-dimensional’: they can be written as the composition of a map $\Delta^2 \rightarrow \Delta^1$ and σ .¹⁰ This leads us to the following definition: Call a singular n -simplex τ in \hat{X} *degenerate* if there is a compact Hausdorff space X_τ of topological dimension less than n such that τ can be written as the composition of continuous maps $\Delta^n \rightarrow X_\tau \rightarrow \hat{X}$.

We would now like to say that we only allow chains (that have a representation) with all but finitely many simplices degenerate. This would not be a proper definition of ‘chain’ since the boundary of a chain would not have to be a chain in this case. This can easily be remedied: Call a chain *good* if it has the above property. We now allow all n -chains that are the sum of a good n -chain and the boundary of a good $(n+1)$ -chain. This homology turns out to satisfy all the Eilenberg-Steenrod axioms [8], and the fact that all 2-simplices in the one-dimensional space $|G|$ are degenerate implies that we indeed obtain the right boundaries. Hence

Theorem 5.2 ([8]). *If G is a locally finite graph and $\hat{G} = |G|$, then the first group $H_1(G)$ of the new homology is canonically isomorphic to $\mathcal{C}(G)$.*

References

- [1] H. Abels and P. Strantzas, *Proper transformation groups*, in preparation.
- [2] R. Diestel, *Locally finite graphs with ends: a topological approach*, <http://arxiv.org/abs/0912.4213>, 2009.
- [3] R. Diestel, *Locally finite graphs with ends: a topological approach. I. Basic theory*, Discrete Math. (to appear).
- [4] R. Diestel, *Locally finite graphs with ends: a topological approach. II. Applications*, Discrete Math. (to appear).
- [5] R. Diestel and D. Kühn, *On infinite cycles I*, Combinatorica **24** (2004), 68–89.
- [6] ———, *On infinite cycles II*, Combinatorica **24** (2004), 91–116.
- [7] R. Diestel and P. Sprüssel, *The fundamental group of a locally finite graph with ends*, arXiv **0910.5647** (2009).
- [8] ———, *On the homology of locally compact spaces with ends*, arXiv **0910.5650** (2009).

¹⁰Note that in general spaces the image of such a 2-simplex does not have to be one-dimensional, since σ could be a space-filling curve.

- [9] ———, *The homology of locally finite graphs with ends*, *Combinatorica* (to appear).
- [10] B. Hughes and A. Ranicki, *Ends of complexes*, Cambridge Univ. Press, 1996.
- [11] P. Sprüssel, *On the homology of infinite graphs with ends*, Ph.D. thesis, Universität Hamburg, 2010, see <http://www.sub.uni-hamburg.de/opus/volltexte/2010/4479/pdf/Dissertation.pdf>.

Reinhard Diestel
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 20146 Hamburg
 Germany

Philipp Sprüssel
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 20146 Hamburg
 Germany

Version 22.03.2010