

**HAMBURGER BEITRÄGE
ZUR MATHEMATIK**

Heft 340

**Locally finite graphs with ends:
a topological approach**

Reinhard Diestel

July 2009

Reinhard Diestel

This paper serves both as an introduction and as a comprehensive survey of the current state of the art in a newly emerging field: a topological approach to the study of locally finite graphs, which crucially incorporates their ends. Topological arcs and circles, which may pass through ends, assume the role played in finite graphs by paths and cycles.

This approach has made it possible to extend to locally finite graphs many classical theorems of finite graph theory that do not extend verbatim. The shift of paradigm it proposes is thus as much an answer to old questions as a source of new ones; many concrete problems of both types are suggested in the paper.

This paper attempts to provide an entry point to this field for readers that have not followed the literature that has emerged in the last 10 years or so. It takes them on a quick route through what appear to be the most important lasting results, introduces them to key proof techniques, identifies the most promising open problems, and offers pointers to the literature for more detail.

Introduction

This paper describes a topological framework in which many well-known theorems about finite graphs that appear to fail for infinite graphs have a natural infinite analogue. It has been realised in recent years that many such theorems, especially about paths and cycles, do work in a slightly richer setting: not in the (locally finite) graph itself, but in its compactification obtained by adding its *ends*. For example, the plane graph G in Figure 0 has three ends. When we add these, we obtain a compact space $|G|$ in which the fat edges form a *circle*—a subspace homeomorphic to the standard topological circle S^1 . Allowing such circles as ‘infinite cycles’, and allowing topological arcs through ends as ‘infinite paths’, we can restore the truth of many well-known theorems about finite graphs whose infinite analogues would fail if we allowed only the usual finite paths and cycles familiar from finite graphs.

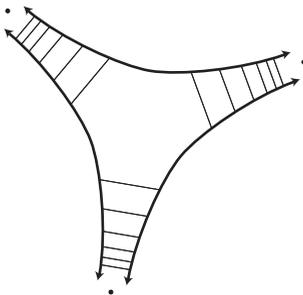


FIGURE 0. A circle through three ends

The aim of this paper is to provide a reasonably complete but readable introduction to this new approach, offering a fast track to its current state of the art. It describes all the fundamental concepts, all the main results, and many open problems. Given this aim, there will not be the space for many proofs. However, there is a complementary source [32, Ch. 8.5] that offers detailed proofs of all the most basic facts (which are here included in the narrative but not all proved), and which the interested reader may wish to consult early for some more ‘feel’ for the subject. For readers already familiar with those basic techniques, the current paper also describes some more advanced but fundamental proof techniques that cannot be found explicitly in [32].

We begin in Section 1 with the definition of the space $|G|$ and an overview of its *basic properties*. These properties mostly concern the topological analogues of familiar finite concepts involving paths and cycles, such as spanning trees, degrees (of ends), connectivity and so on. It is these analogues that we shall need in the place of their finite counterparts when we wish to extend theorems from finite graph theory to infinite graphs whose naive extension fails.

We continue in Section 2 with some *basic theory*, a body of results that are not yet applications of the new topological concepts but relate them to each other, much in the way their finite counterparts are related. Most results in this section concern the homology of a graph, i.e., its *cycle* and *cut space* and they way they interact. These homology aspects have, so far, been the prime field of application for those topological notions involving ends. But there are other applications too, and reading Section 2 will not be a prerequisite for reading the rest of the paper.

Section 3 explains some *proof techniques*, e.g. for the construction of topological arcs and circles as limits of finite paths and cycles, that have evolved over the past 10 years. These years have seen some considerable simplification of the techniques required to deal with the problems one usually encounters in this area. The aim of Section 3 is to describe the state of the art here, so as to equip those new to the field quickly with the main techniques now available.

Section 4 is devoted to *applications*: theorems that involve our new topological concepts but answer questions that could be asked without them. Often, these questions simply ask for an extension to infinite graphs of well known finite theorems whose naive infinite extensions are either false or trivially true. Examples include all the standard planarity criteria in terms of the cycle space, some classical theorems about Euler tours and Hamilton cycles, connectivity results such as the tree-packing theorem of Nash-Williams and Tutte, and infinite electrical networks. The consideration of arcs and circles instead of paths and cycles has made it possible to ask extremal-type questions about the interaction of graph invariants that were so far meaningful only for finite graphs. Some of these are indicated in Section 4; more are given in Stein’s paper [78] in this volume.

Finally, there is an *outlook* in Section 5 to new horizons: extensions to graphs that are not locally finite, implications of our findings in related fields

such as geometric group theory and the homology of locally compact spaces, and applications to infinite matroids.

Open problems are not collected at the end but interspersed within the text. There are plenty of these, as well as the overall quest to push the general approach further: to identify more theorems about paths and cycles in finite graphs that do not extend naively, and to find the correct topological analogue that does extend.

1. Concepts and Basic Theory

Throughout this section, let G be a fixed infinite, locally finite, connected graph. This section serves to introduce the concepts on which our topological approach to the study of such graphs is based: the space $|G|$ formed by G and its ends; topological paths, circles and spanning trees in this space; notions of connectivity in $|G|$. The style will be descriptive and informal, aiming for overall readability; should any technical points remain unclear, the reader is referred to [32, Ch. 8.5] for more formal definitions of the concepts introduced here, and to [32] in general for graph-theoretic terms and notation.

Terms such as ‘path’ or ‘connected’, which formally have different meanings in topology and in graph theory, will be used according to context: in the graph-theoretical sense for graphs, and in the usual topological sense for topological spaces. If the context is ambiguous, the two meanings will probably coincide, making a formal distinction unnecessary.

We call 1-way infinite paths *rays*, and 2-way infinite paths *double rays*. An *end* of G is an equivalence class of rays in G , where two rays are considered equivalent if no finite set of vertices separates them in G . The graph shown in Figure 0 has three ends; the $\mathbb{Z} \times \mathbb{Z}$ grid has only one, the infinite binary tree has continuum many. We write $\Omega(G)$ for the set of ends of G .

Topologically, we view G as a cell complex with the usual topology. Adding its ends compactifies it, with the topology generated by the open sets of G and the following additional basic open sets. For every finite set S of vertices and every end ω , there is a unique component C of $G - S$ in which every ray of ω has a tail. We say that ω *lives in* C and write $C =: C(S, \omega)$. Now for every such S and every component C of $G - S$, we declare as open the union \hat{C} of C with the set of ends living in C and with all the ‘open’ S - C edges of G (i.e., without their endpoints in S). We denote the space just obtained by $|G|$.¹

Theorem 1.1. [31, 32] *The space $|G|$ is compact and metrizable.*

¹ The open neighbourhoods of ends are defined slightly more generally in [32], but for locally finite graphs the two definitions are equivalent. Topologies for graphs with infinite degrees are discussed in Section 5.

The space $|G|$ is known as the *Freudenthal compactification* of G . The main feature of its topology is that rays in G converge as they should: to the end of which they are an element. One can show [34, Prop. 4.5] that every Hausdorff topology on $G \cup \Omega(G)$ with (essentially) this feature (and which induces the 1-complex topology on G) refines the topology of $|G|$. This identifies $|G|$ as the unique most powerful Hausdorff topology on $G \cup \Omega(G)$, in a sense that can be made quite precise [34].²

Of the many natural aspects of this topology let us mention just two more, which relate it to better-known objects. Consider the binary tree T_2 , and think of its rays from the root as 0–1 sequences. The resulting bijection between the ends of T_2 and these sequences is a homeomorphism between $\Omega(T_2)$, as a subspace of $|T_2|$, and $\{0, 1\}^{\mathbb{N}}$ with the product topology. Identifying pairs of ends whose sequences specify the same rational (one sequence ending on zeros, the other on 1s) turns this bijection into a homeomorphism from the resulting identification space of $\Omega(T_2)$ to $[0, 1]$. Without such identification, on the other hand, $\Omega(G)$ is always a subset of a Cantor set.

Instead of paths and cycles in G we can now consider *arcs* and *circles* in $|G|$: homeomorphic images of the real interval $[0, 1]$ and of the complex unit circle S^1 . While paths and cycles are examples of arcs and circles, Figure 0 shows a circle that is not a cycle. Arcs and circles that are not paths or cycles must contain ends. An arc containing uncountably many ends always induces the ordering of the rationals on a subset of its vertices [16]. Such arcs, and circles containing them, are called *wild* (Figure 1.1), but they are quite common.

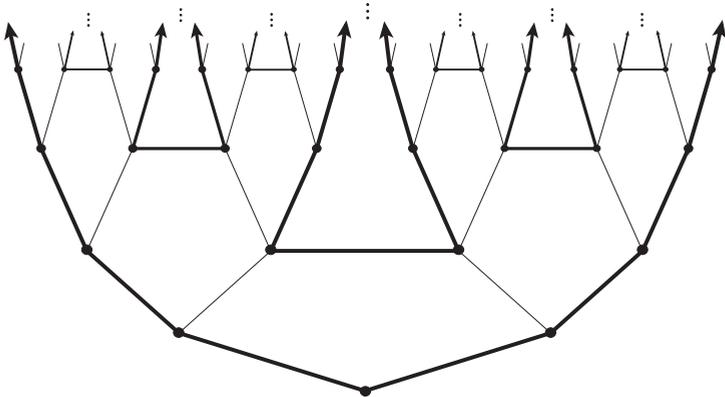


FIGURE 1.1. The heavy edges form a *wild circle*.

² The identity to $|G|$ from $G \cup \Omega(G)$ with any finer topology is continuous, so in $|G|$ there will be at least as many arcs and circles (and possibly more). Arcs and circles in $|G|$ will be our main tools.

Arcs and circles are examples of a natural type of subspace of $|G|$: subspaces that are the closure in $|G|$ of some subgraph of G .³ We call such a subspace X of $|G|$ a *standard subspace*, and write $V(X)$ and $E(X)$ for the set of vertices or edges it contains. Note that the ends in X are ends of G , not of the subgraph that gave rise to X ; in particular, ends in X need not have a ray in that subgraph. Which ends of G are in X is determined just by $V(X)$: they are precisely the ends that are limits of vertices in X .

Given a standard subspace X and an end $\omega \in X$, the maximum number of arcs in X that end in ω but are otherwise disjoint is the (*vertex-*) *degree* of ω in X ; the maximum number of edge-disjoint arcs in X ending in ω is its *edge-degree* in X . Both maxima are indeed attained, but it is non-trivial to prove this [22]. End degrees behave largely as expected; for example, the connected standard subspaces in which every vertex and every end has (vertex-) degree 2 are precisely the circles. (Use Lemma 1.2 to prove this.)⁴ In Section 4.1 we shall define a third type of end degrees, their *relative degree*, which is useful for the application of end degrees to extremal-type problems about infinite graphs.

Standard subspaces have the important property that connectedness and arc-connectedness are equivalent for them. This will often be convenient: while connectedness is much easier to prove (see Lemma 1.4), it is usually arc-connectedness that we need.

Lemma 1.2. [32, 36, 83] *Connected standard subspaces of $|G|$ are locally connected and arc-connected.*

The proof that a connected standard subspace X is locally connected is not hard: an open neighbourhood $\hat{C} \cap X$ of an end ω will be connected if we choose the set S in its definition so as to minimize the number of C - S edges in X . But local connectedness is not a property we shall often use directly. Its role here is that it offers a convenient stepping stone towards the proof of arc-connectedness.⁵ Direct proofs that X is arc-connected can be found in [36, 42], and we shall indicate one in Section 3. Connected subspaces of $|G|$ that are neither open nor closed need not be arc-connected [42].

Lemma 1.2 has the important corollary that arc-components of standard subspaces are closed (because the closure of an arc-component, which is obviously connected and standard, is arc-connected by the lemma.)

³ It takes a line of proof that arcs and circles do indeed have this property, i.e., that the union of their edges is dense in them. This is because any arc between distinct ends must contain an edge incident with finite set S that separates them; this will follow from Lemma 1.3 below.

⁴ However, there may also be surprises. The wild circle of Figure 1.1, for example, is a path-connected standard subspace, but it is also a disjoint union of arcs: the closures of the fat double rays. (These arcs are *really* disjoint, also at ends!)

⁵ By general ‘continuum theory’ [67], compact, connected and locally connected metric spaces are arc-connected.

The edge set $E(C)$ of any circle C will be called a *circuit*. Given any set F of edges in G , we write \overline{F} for the closure of $\bigcup F$ in $|G|$, and call \overline{F} the standard subspace *spanned by* F . (This is with slight abuse of our usual notation, in which we write \overline{X} for the closure in $|G|$ of a subset $X \subseteq |G|$.) Similarly, we write $\overset{\circ}{F}$ for the set of all inner points of edges in F (while usually we write $\overset{\circ}{X}$ for the interior of a subset $X \subseteq |G|$).

The set of edges of G across a partition $\{V_1, V_2\}$ of $V(G)$ is a *cut* of G ; the sets V_1, V_2 are the *sides* of this cut. A minimal non-empty cut is a *bond*.

The following lemma is one of our basic tools for handling arcs. It says that an arc cannot ‘jump across’ a *finite* cut without containing an edge from it:

Lemma 1.3. (Jumping Arc Lemma [32])

Let F be a cut of G with sides V_1, V_2 .

- (i) F is finite if and only if $\overline{V_1} \cap \overline{V_2} = \emptyset$.
- (ii) If F is finite, there is no arc in $|G| \setminus \overset{\circ}{F}$ with one endpoint in V_1 and the other in V_2 . If F is infinite, there may be such an arc.

The proof of Lemma 1.3 for F finite is straightforward from the definition of the topology of $|G|$: deleting the edges of a finite cut splits $|G|$ into two disjoint open sets. When F is infinite, an intersection point ω as in (i) can be obtained as the limit of two vertex sequences, one in V_1 and the other in V_2 , that are joined by infinitely many cut edges.

Although the ‘jumping arc’ is a nice way of remembering Lemma 1.3, its main assertion (i) is not about arcs but about connectedness. It implies that connectedness for standard subspaces can be characterized in graph-theoretical terms alone, without any explicit mention of ends or the topology of $|G|$:

Lemma 1.4. [32] *A standard subspace of $|G|$ is connected if and only if it contains an edge from every finite cut of G of which it meets both sides.*

We shall say that a standard subspace X of $|G|$ is *k-edge-connected* if the deletion of fewer than k edges will not make it disconnected. Similarly, X is *k-vertex-connected* if $V(X) > k$ and the deletion of fewer than k vertices and their incident edges does not leave a disconnected space. Note that for $k = 1$ both notions coincide with ordinary topological connectedness, and that for $X = |G|$ the space X is *k-edge-connected* or *k-vertex-connected* if and only if the graph G is *k-edge-connected* or *k-connected* (by Lemma 1.4).

How about deleting ends as well as vertices and/or edges? For $X = |G|$, this will never help to disconnect X : if deleting a finite set U of vertices and any set of ends disconnects $|G|$, then so does the deletion of U alone, and similarly for edges. For arbitrary standard subspaces X , however, deleting ends can make sense. It will normally result in a subspace that is no longer standard, but the

main reason for primarily considering standard subspaces, that connectedness in them is equivalent to arc-connectedness (Lemma 1.2), is preserved.⁶

So let us call a subspace $X \subseteq |G|$ *substandard* if its closure in $|G|$ is standard (so that X contains no partial edges), and *k-connected* if after the deletion of fewer than k vertices, edges or ends it will still be arc-connected. This makes sense in the context of Menger’s theorem, which Thomassen and Vella [83] proved for topological spaces X that include all subspaces of $|G|$: given any two points $a, b \in X$ and $k \in \mathbb{N}$, if for every set $S \subseteq X$ of fewer than k points there is an a - b arc in $X \setminus S$, then X contains k arcs from a to b that pairwise meet only in a and b .⁷ Hence in a k -connected standard or substandard subspace any two vertices or ends can be linked by k independent arcs: a useful property that can fail in standard subspaces that are merely k -vertex-connected.

To work with arcs and circles in a way that resembles finite graph theory, we need one more addition to our topological toolkit: the notion of a topological spanning tree. A *topological spanning tree* of a connected standard subspace X of $|G|$ is an (arc-) connected standard subspace $T \subseteq X$ of $|G|$ that contains every vertex (and hence every end) of X but contains no circle. A topological spanning tree of $|G|$ will also be called a *topological spanning tree of G* .

The closure \overline{T} of an ordinary spanning tree T of G is not normally a topological spanning tree: as soon as T contains disjoint rays from the same end, \overline{T} will contain a circle. Conversely, the subgraph of G underlying a topological spanning tree need not be a graph-theoretical tree: it will be acyclic, of course, but it need not be connected. Figure 1.2 shows examples of both these phenomena in the double ladder.



FIGURE 1.2. An ordinary spanning tree of G (left), and a topological spanning tree of G (right)

Ordinary spanning trees whose closures are topological spanning trees do always exist, however: all *normal* spanning trees have this property, and all countable connected graphs have normal spanning trees. (A spanning tree T of G is *normal* if, for a suitable choice of a root, the endvertices of every edge of G are comparable in the tree-order of T . See [32].) Often, therefore, normal spanning trees are the best choice of a spanning tree for our purposes.

More generally, we have the following existence lemma:

⁶ Georgakopoulos’s characterization of the subspaces that are connected but not arc-connected [42] implies that any such space has uncountably many arc-components consisting of one end only. One clearly cannot obtain such a space by deleting finitely many ends from a connected standard subspace.

⁷ There is also version of Menger’s theorem for disjoint arcs between sets $A, B \subseteq X$ in [83], which is easier to prove.

Lemma 1.5. *Let X be a connected standard subspace of $|G|$, and let $Z \subseteq X$ be a standard subspace not containing a circle. Then X has a topological spanning tree containing Z .*

Proof. We begin by enumerating the edges in $E(X) \setminus E(Z)$. We then go through these edges one by one, considering each for deletion from X . We delete an edge if this does not disconnect the space $X \setminus \overset{\circ}{F}$, where F is the set of edges already deleted. Having considered every edge in $E(X) \setminus E(Z)$, we are left with a standard subspace T that contains $V(X)$ but contains no circle: this would have an edge in $E(X) \setminus E(Z)$, which was considered for deletion and should have been deleted. The space T is connected (cf. Lemma 1.4), and hence arc-connected (Lemma 1.2). By construction, $Z \subseteq T \subseteq X$ as desired. \square

Unlike in finite graphs, it is considerably harder to construct a topological spanning tree ‘from below’ (maintaining acyclicity) than, as we did just now, ‘from above’ (maintaining connectedness, and using the non-trivial Lemma 1.2). A proof ‘from below’ will be indicated in the proof of Lemma 3.2.

A topological spanning tree T of G features analogues of all the usual properties of spanning trees in finite graphs: adding any edge creates a circle, deleting any edge disconnects it, and it contains a unique arc between any two of its points. Hence every chord $e \in E(G) \setminus E(T)$ creates a well-defined *fundamental circuit* C_e in $T \cup e$, while every edge $f \in E(T)$ lies in a well-defined *fundamental cut* D_f of G (whose sides are the vertex sets of the two arc-components of $T \setminus \overset{\circ}{f}$).

Fundamental circuits and cuts of topological spanning trees are subject to the same duality as for ordinary spanning trees:

Lemma 1.6. *Let C_e be a fundamental circuit and D_f a fundamental cut, both of the same topological spanning tree of G . Then $e \in D_f \Leftrightarrow f \in C_e$.*

However, the sizes of fundamental cuts and circuits behave in opposite ways:

Lemma 1.7. *The following statements hold in any locally finite graph.*

- (i) *The fundamental circuits of an ordinary spanning tree are finite, but its fundamental cuts may be infinite.*
- (ii) *The fundamental circuits of a topological spanning tree may be infinite, but its fundamental cuts are finite.*
- (iii) *The fundamental circuits and cuts of normal spanning trees are finite.*

For (ii), note that if a fundamental cut D_f of a topological spanning tree T was infinite, the two arc-components of $T \setminus \overset{\circ}{f}$ (which are closed by Lemma 1.2) would meet, by Lemma 1.3 (i).

The fact that fundamental cuts of topological spanning trees are finite implies by Lemma 1.3 (ii) that they are in fact bonds.

2. The Topological Cycle Space

The way in which cycles and cuts interact in a graph can be described algebraically: in terms of its ‘cycle space’, its ‘cut space’, and the duality between them. In this section we show how the cycle space theory of finite graphs extends to locally finite graphs in a way that encompasses infinite circuits. The fact that this can be done, that our topological circuits, cuts and spanning trees interact in the same way as ordinary cycles, cuts and spanning trees do in a finite graph, is by no means clear but rather surprising. For example, there is nothing visibly topological about a finite cut in an infinite graph,⁸ so the fact that the edge sets orthogonal to its finite cuts are precisely its topological circuits and their sums (Theorem 2.6) comes as a pleasant surprise: it provides a natural answer to a natural question, but not by design—it was not ‘built into’ the definition of a circle.

As it turns out, extending finite cycle space theory in this way is not only possible but also necessary: it is the ‘topological cycle space’ of a locally finite graph, not its usual finitary cycle space, that interacts with its other structural features, such as planarity, in the way we know it from finite graphs. We shall discuss this in some depth in Section 4. In Section 5 we shall see how our theory integrates with the broader topological context of (singular) homology in more general spaces.

As before, let G be a fixed infinite, connected, locally finite graph. We start by defining the ‘topological cycle space’ $\mathcal{C}(G)$ of G in analogy to the mod-2 (or ‘unoriented’) cycle space of a finite graph:⁹ its elements will be sets of edges (that is to say, maps $E(G) \rightarrow \mathbb{Z}_2$, or formal sums of edges with coefficients in \mathbb{Z}_2) generated from circuits by taking symmetric differences of edge sets. These edge sets, the circuits, and the sums may be infinite.

Let us make this more precise. Let the *edge space* $\mathcal{E}(G)$ of G be the \mathbb{Z}_2 -vector space of all maps $E(G) \rightarrow \mathbb{Z}_2$, which we think of as subsets of $E(G)$ with symmetric difference as addition. The *vertex space* $\mathcal{V}(G)$ is defined likewise. Call a family $(D_i)_{i \in I}$ of elements of $\mathcal{E}(G)$ *thin* if no edge lies in D_i for infinitely many i . Let the (*thin*) *sum* $\sum_{i \in I} D_i$ of this family be the set of all edges that lie in D_i for an odd number of indices i . Given any subset $\mathcal{D} \subseteq \mathcal{E}(G)$, the edge sets that are sums of sets in \mathcal{D} form a subspace of $\mathcal{E}(G)$. The (*topological*) *cycle space* $\mathcal{C}(G)$ of G is the subspace of $\mathcal{E}(G)$ consisting of the sums of circuits. The *cut space* $\mathcal{B}(G)$ of G is the subspace of $\mathcal{E}(G)$ consisting of all the cuts in G and the empty set. (Unlike the circuits, these already form a subspace.) We sometimes call the elements of $\mathcal{C}(G)$ *algebraic cycles* in G .

⁸ With hindsight, of course, Lemma 1.4 shows that this impression is wrong.

⁹ One can also define ‘oriented’ versions $\vec{\mathcal{C}}(G)$ of $\mathcal{C}(G)$, with integer or real coefficients [38]. Some of the theorems listed below have obvious oriented analogues, with very similar proofs. But the differences have not yet been investigated systematically and may well be worth further study; see Sections 2.3–4 and 4.3 for some good problems.

2.1 Generating sets

The sums of elements of $\mathcal{D} \subseteq \mathcal{E}(G)$, and the subspace of $\mathcal{E}(G)$ consisting of all those sums, are said to be *generated* by \mathcal{D} . For example, the cycle space of the graph in Figure 0 is generated by its central hexagon and its squares, or by the infinite circuit consisting of the fat edges and all the squares. Bruhn and Georgakopoulos [19] proved that if \mathcal{D} is thin, the space it generates is closed under thin sums. As we shall see, this applies to both $\mathcal{C}(G)$ and $\mathcal{B}(G)$.

As in Lemma 1.7, the duality between $\mathcal{C}(G)$ and $\mathcal{B}(G)$ —which we address more thoroughly later—is reflected by a switch between topological and ordinary spanning trees:

Theorem 2.1. [32]

- (i) *Given an ordinary spanning tree of G , its fundamental cuts generate $\mathcal{B}(G)$, but its fundamental circuits need not generate $\mathcal{C}(G)$.*
- (ii) *Given a topological spanning tree of G , its fundamental circuits generate $\mathcal{C}(G)$, but its fundamental cuts need not generate $\mathcal{B}(G)$.*
- (iii) *Given a normal spanning tree of G , its fundamental circuits generate $\mathcal{C}(G)$, and its fundamental cuts generate $\mathcal{B}(G)$.*

To prove the first assertion in Theorem 2.1 (ii), one shows that a given set $D \in \mathcal{C}(G)$ equals the sum $\sum C_e$ taken over all chords e of the topological spanning tree such that $e \in D$. One has to show that these C_e form a thin family (use Lemmas 1.6 and 1.7 (ii)), but also that $D = \sum C_e$. While for finite G one just notes that $D + \sum C_e$ consists of tree edges and deduces that $D + \sum C_e = \emptyset$ (yielding $D = \sum C_e$), this last implication is now non-trivial: it is not clear that the tree, which by assumption contains no circuit, cannot contain a sum of circuits. The proof in [32] deduces this, i.e. that D and $\sum_e C_e$ also agree on tree edges, from Theorem 2.5 (i) \Rightarrow (iv) below, which in turn is an easy consequence of the jumping arc lemma.

For the second assertion in (ii), recall that the edge set of a topological spanning tree can miss an infinite cut (as in Figure 1.2). Such a cut will not be a sum of fundamental cuts D_f , because any such sum contains all those tree-edges f .

The proof of assertion (i) is analogous to that of (ii), though in the first statement one also has to show that the fundamental cuts do not generate *more* than $\mathcal{B}(G)$: that a thin sum of cuts is again a cut. (See the discussion after Theorem 2.2.) For the second statement of (i), recall that the edges of an ordinary spanning tree may contain a circuit, which will not be a sum of fundamental circuits C_e because any such sum contains all those chords e .

Assertion (iii) follows from (i) and (ii).

By Lemmas 1.6 and 1.7, the fundamental circuits and cuts of normal spanning trees form thin families. Hence by the result of [19] cited earlier, the sub-

spaces they generate in $\mathcal{E}(G)$ are closed under thin sums. By Theorem 2.1 (iii) these are the spaces $\mathcal{C}(G)$ and $\mathcal{B}(G)$:

Theorem 2.2.

- (i) $\mathcal{C}(G)$ is generated by finite circuits and is closed under infinite thin sums.
- (ii) $\mathcal{B}(G)$ is generated by finite bonds and is closed under infinite thin sums.

Since we already used the closure statement of (ii) for cuts in our proof of Theorem 2.1, our proof of this part of Theorem 2.2 (ii) has been circular. But one does not need Theorem 2.1, or indeed [19], to prove that a sum of cuts is again a cut, a set of edges across a vertex partition. One way to do this directly is to construct that partition explicitly. The simplest way, however, is indirect: to show first that the cuts are precisely the edge sets that meet every finite circuit in an even number of edges,¹⁰ and then to observe that the set of those edge sets is closed under taking sums.

But also the closure statement of (i), for $\mathcal{C}(G)$, is non-trivial. Indeed, if $\sum D_i$ is a thin sum of sets $D_i = \sum_j C_i^j \in \mathcal{C}(G)$, where the C_i^j are circuits, the combined sum $\sum_{i,j} C_i^j$ need not be thin. By Theorem 2.3 below one can assume for each i that its own C_i^j are disjoint, in which case the combined family (C_i^j) will be thin (because the family (D_i) is thin by assumption). But Theorem 2.3 is hard, and too big a tool for this proof. As for cuts, the quickest way to prove that $\mathcal{C}(G)$ is closed under thin sums is to use Theorem 2.7 (i): that the elements of $\mathcal{C}(G)$ are precisely those sets of edges that meet every finite cut evenly, a class of edge sets that is obviously closed under sums. See also the discussion after Theorem 2.7.

A bond is *atomic* if it consists of all the edges at some fixed vertex. Following Tutte, let us call a circuit in $|G|$ *peripheral* if it contains every edge between its incident vertices (i.e., has no chord in G) and the set of these vertices does not separate G .

Theorem 2.3.

- (i) Every element of $\mathcal{C}(G)$ is a disjoint union of circuits.
- (ii) If G is 3-connected, its peripheral circuits generate $\mathcal{C}(G)$.
- (iii) Every element of $\mathcal{B}(G)$ is a disjoint union of bonds.
- (iv) The atomic bonds of G generate $\mathcal{B}(G)$.

Statement (i) was first established in [35], with a long and involved proof. As techniques developed, simpler and more elegant proofs were found [44, 83, 85]. We shall discuss these techniques, and the different proofs they lead to, in Section 3. Statement (ii), which extends a theorem of Tutte for finite graphs (see [32]), is due to Bruhn [11]. It fails if we restrict the peripheral circuits

¹⁰ This is Theorem 2.7 (ii); its proof will be independent of what we are proving here.

or the sums used to finite circuits or sums. Statements (iii) and (iv) are easy, with proofs as for finite graphs.

Theorem 2.3 implies that, as in finite graphs, the circuits in $|G|$ are the minimal non-empty elements of $\mathcal{C}(G)$, while the bonds in G are the minimal non-empty elements of $\mathcal{B}(G)$.

Finally, there is a generating theorem for the cycle space of a finite graph whose extension to $|G|$ and $\mathcal{C}(G)$ requires an interesting twist. When G is finite, its cycle space is generated by the edge sets of its *geodetic* cycles: those that contain a shortest path in G between any two of their vertices [32]. This is still true for locally finite graphs (with ‘arc’ instead of ‘path’), but only if we measure length of an arc in the right way: rather than by counting its edges, we have to assign lengths to the edges of G , and then measure the length of an arc by adding up the lengths of its edges. By giving edges shorter lengths if they lie ‘far out near the ends’, it is possible to do this in such a way that the resulting metric on $|G|$ induces its original topology.¹¹ And then, no matter how exactly we choose the edge lengths, Georgakopoulos and Sprüssel proved that the geodetic circles do what they should:

Theorem 2.4. [48] *Given any edge-length function $\ell: E(G) \rightarrow [0, \infty)$ whose resulting metric on $|G|$ induces the original topology of $|G|$, the circuits of the geodetic circles in $|G|$ generate $\mathcal{C}(G)$.*

2.2 Characterizations of algebraic cycles

There are various equivalent ways to describe the elements of $\mathcal{C}(G)$ and of $\mathcal{B}(G)$, each extending a similar statement about finite graphs. Let us list these now, beginning with $\mathcal{C}(G)$.

A closed topological path in a standard subspace X of $|G|$, based at a vertex, is a *topological Euler tour* of X if it traverses every edge in X exactly once. One can show that if $|G|$ admits a topological Euler tour it also has one that traverses every end at most once [45]. For arbitrary standard subspaces this is false: consider the closure of two disjoint double rays in the $\mathbb{Z} \times \mathbb{Z}$ grid.

Recall that, given a set D of edges, \overline{D} denotes the closure of the union of all the edges in D , the standard subspace of $|G|$ *spanned by* D .

Theorem 2.5. *The following statements are equivalent for sets $D \subseteq E(G)$:*

- (i) $D \in \mathcal{C}(G)$, that is to say, D is a sum of circuits in $|G|$.
- (ii) Every component of \overline{D} admits a topological Euler tour.
- (iii) Every vertex and every end has even (edge-) degree in \overline{D} .
- (iv) D meets every finite cut in an even number of edges.

¹¹ This will be discussed in detail in Section 5.2.

The equivalence of (i) and (ii) was proved in [34] for $D = E(G)$ and extended to arbitrary D by Georgakopoulos [44]; we shall meet the techniques needed for the proof in Section 3. The equivalence with (iii) is a deep theorem, due to Bruhn and Stein [22] for $D = E(G)$ and to Berger and Bruhn [8] for arbitrary D . Note that (iii) assumes that end degrees have a parity even when they are infinite. Finding the right way to divide ends of infinite edge-degree into ‘odd’ and ‘even’ was one of the major difficulties to overcome for this characterization of $\mathcal{C}(G)$.

The equivalence of (i) with (iv), again from [34], is one of the cornerstones of topological cycle space theory: its power lies in the fact that the finitary statement in (iv) is directly compatible with compactness proofs (see Section 3). Its implication (i)→(iv) follows from the jumping arc lemma, applied to the circles whose circuits sum to the given set $D \in \mathcal{C}(G)$. For the converse implication one compares a given set D as in (iv) with the sum $\sum C_e$ of fundamental circuits of a topological spanning tree taken over all chords $e \in D$. It is clear that D agrees with this sum on chords. Using Lemmas 1.6 and 1.7 one proves that they also agree on tree edges.

2.3 Cycle-cut orthogonality

Given a set $\mathcal{F} \subseteq \mathcal{E}(G)$, let us write \mathcal{F}_{fin} for the set of finite elements of \mathcal{F} . Call two sets $D, F \subseteq E(G)$ *orthogonal* if $|D \cap F|$ is finite and even, and put

$$\mathcal{F}^\perp := \{ D \subseteq E(G) \mid D \text{ is orthogonal to every } F \in \mathcal{F} \}.$$

Let us abbreviate $\mathcal{V} := \mathcal{V}(G)$, $\mathcal{E} := \mathcal{E}(G)$, $\mathcal{C} := \mathcal{C}(G)$ and $\mathcal{B} := \mathcal{B}(G)$.

The following theorem describes the duality between the topological cycle and cut space of G in terms of orthogonal sets. Its four assertions show some interesting symmetries: they can be summarized neatly as ‘all equations containing each of the symbols $=, \mathcal{C}, \mathcal{B}, \text{fin}$ and $^\perp$ exactly once (and no others)’.

Theorem 2.7.

- (i) $\mathcal{C} = \mathcal{B}_{\text{fin}}^\perp$.
- (ii) $\mathcal{B} = \mathcal{C}_{\text{fin}}^\perp$.
- (iii) If G is 2-connected, then $\mathcal{C}^\perp = \mathcal{B}_{\text{fin}}$.
- (iv) $\mathcal{B}^\perp = \mathcal{C}_{\text{fin}}$.

Theorem 2.7 (i) is just a reformulation of Theorem 2.5 (i)⇔(iv).

Assertion (ii) is the dual (see Section 5.4) finitary characterization of the cut space. An interesting feature of its seemingly trivial inclusion $\mathcal{B} \subseteq \mathcal{C}_{\text{fin}}^\perp$ is that it requires an application of Theorem 2.3 (i). Indeed, any cut F is clearly orthogonal to every finite circuit. But an arbitrary finite element D of \mathcal{C} might come as an infinite sum circuits (finite or infinite), and even if F is orthogonal to every term in the sum it need not be orthogonal to D . However, the circuits

into which D decomposes *disjointly* by Theorem 2.3 (i) will be finite and yield D in a finite sum, which preserves orthogonality. The proof of $\mathcal{B} \supseteq \mathcal{C}_{\text{fin}}^\perp$ is the same as for finite graphs [32]: contracting the edges not in a given set $D \in \mathcal{C}_{\text{fin}}^\perp$ leaves a graph with edge set D , whose finite circuits have to be even by definition of D , because they expand to finite circuits of G . This contracted bipartite graph defines a partition of $V(G)$ showing that D is a cut.

The inclusion $\mathcal{C}^\perp \supseteq \mathcal{B}_{\text{fin}}$ in (iii) is equivalent to $\mathcal{C} \subseteq \mathcal{B}_{\text{fin}}^\perp$ of (i). The converse inclusion is due to Richter and Vella [73]. Its interesting part is that an edge set orthogonal to every circuit cannot be infinite. (It will be a cut by (ii), since in particular it is orthogonal to every finite circuit.) It is here that 2-connectedness is needed; assuming 2-edge-connectedness, for example, is easily seen not to be enough.

In (iv), the weaker inclusion of $\mathcal{B}^\perp \subseteq \mathcal{C}$ follows from (i). To see that an edge set orthogonal to all cuts cannot be infinite, assume it is, pick infinitely many independent edges, and obtain a contradiction by extending these to a cut. The inclusion $\mathcal{B}^\perp \supseteq \mathcal{C}_{\text{fin}}$ is equivalent to the inclusion $\mathcal{B} \subseteq \mathcal{C}_{\text{fin}}^\perp$ of (i).

Theorem 2.7 (ii) has a nice corollary (which we already used in our proofs of Theorems 2.1 and 2.2): that $\mathcal{C}(G)$ and $\mathcal{B}(G)$ are closed under infinite sums. While this is not immediate from the definition of either $\mathcal{C}(G)$ or $\mathcal{B}(G)$, the proof that $\mathcal{B}_{\text{fin}}^\perp (= \mathcal{C}(G))$ and $\mathcal{C}_{\text{fin}}^\perp (= \mathcal{B}(G))$ are closed under thin sums is immediate from the definition of these sets.

We remark further that, unlike in finite graphs, there can be edge sets that are orthogonal to all circuits but not to all elements of \mathcal{C} , and edge sets that are orthogonal to all bonds but not to all cuts [32] (even when G is 2-connected).

For finite G it is well known that $\dim(\mathcal{E}) = \dim(\mathcal{C}) + \dim(\mathcal{B})$ (see e.g. [32]), so $\mathcal{E}/\mathcal{B} \simeq \mathcal{C}$ and $\mathcal{E}/\mathcal{C} \simeq \mathcal{B}$. All this is still true for infinite G and our topological cycle space \mathcal{C} . But it is more instructive to find canonical such isomorphisms: to define an epimorphism $\sigma: \mathcal{E} \rightarrow \mathcal{C}$ with kernel \mathcal{B} , and an epimorphism $\tau: \mathcal{E} \rightarrow \mathcal{B}$ with kernel \mathcal{C} . This can be done with Theorem 2.7. Indeed, let T be a normal spanning tree¹² of G , and write C_e and D_f for its fundamental circuits and cuts. By Lemma 1.7, all these are finite. Given $E \subseteq E(G)$, define

$$\sigma(E) := \sum \{ C_e : e \in E(G) \setminus E(T) \text{ and } |E \cap C_e| \text{ is odd} \}$$

and

$$\tau(E) := \sum \{ D_f : f \in E(T) \text{ and } |E \cap D_f| \text{ is odd} \}.$$

Corollary 2.8.

- (i) *The map σ is an epimorphism $\mathcal{E} \rightarrow \mathcal{C}$ with kernel \mathcal{B} .*
- (ii) *The map τ is an epimorphism $\mathcal{E} \rightarrow \mathcal{B}$ with kernel \mathcal{C} .*

¹² For each of σ and τ defined below we shall need that T is both an ordinary and a topological spanning tree.

Proof. (i) To see that σ is onto, let $D \in \mathcal{C}$ be given and let E consist of those chords e whose fundamental circuits C_e sum to D . The kernel of σ consists of those sets E that meet every C_e evenly. These are precisely the sets E that meet every finite circuit evenly, since these are generated by finitely many C_e . By Theorem 2.7 (ii), these sets E are precisely those in \mathcal{B} .

(ii) is analogous, using Theorem 2.7 (i). \square

By Theorem 2.7, subsets of sets in \mathcal{C} cannot contain an odd cut, and subsets of sets in \mathcal{B} cannot contain an odd circuit. These properties are the only obstructions to being a subset of a set in \mathcal{C} or \mathcal{B} :

Theorem 2.9. *Let $E \subseteq E(G)$ be any set of edges.*

- (i) *E extends to an element of \mathcal{B} if and only if E contains no odd circuit.*
- (ii) *E extends to an element of \mathcal{C} if and only if E contains no odd bond.*

Proof. We prove the backward implications.

(i) Pick a spanning tree in each component of the graph $(V(G), E)$. Adding edges of $E(G) \setminus E$, extend the union of these trees to a spanning tree T of G . Consider the sum

$$F := \sum_{f \in E \cap E(T)} D_f \in \mathcal{B}$$

of fundamental cuts of T . (The sum is thin by Lemmas 1.6 and 1.7.) We claim that $E \subseteq F$. Let $e \in E$ be given. If $e \in E(T)$ then $e \in F$, because D_e is the only fundamental cut of T containing e . If e is a chord of T , then its fundamental circuit C_e lies in E , by construction of T . As C_e is finite and hence even (by assumption), Lemma 1.6 implies that $e \in D_f$ for an odd number of edges $f \in E \cap E(T)$, giving $e \in F$ as desired.

(ii) In every (arc-) component of the standard subspace $\overline{E(G) \setminus E}$ pick a topological spanning tree (Lemma 1.5). Use Lemma 1.5 again to extend their union to a topological spanning tree T of G ; note that the additional edges will be in E , since adding them did not create a circle. Consider the sum

$$D := \sum_{e \in E \setminus E(T)} C_e \in \mathcal{C}$$

of fundamental circuits of T . (The sum is thin by Lemmas 1.6 and 1.7.) We claim that $E \subseteq D$. Let $d \in E$ be given. If $d \notin E(T)$ then $d \in D$, because C_d is the only fundamental circuit of T containing d . If $d \in E(T)$, then its fundamental cut D_d lies in E , by construction of T . As D_d is finite (Lemma 1.7) and hence even (by assumption; recall that fundamental cuts are in fact bonds), Lemma 1.6 implies that $d \in C_e$ for an odd number of edges $e \in E \setminus E(T)$, giving $d \in D$ as desired. \square

2.4 Orthogonal decomposition

If we can write a set $E \subseteq E(G)$ as a sum $E = D + F$ with $D \in \mathcal{C}$ and $F \in \mathcal{B}$, we call this a *decomposition* of E . If the sets D and F are orthogonal (which for infinite sets is not automatic, their intersection could be infinite or even odd; cf. Lemma 1.3), we call this decomposition of E *orthogonal*.

Problem 2.10.

- (i) For which G does every set $E \in \mathcal{E}(G)$ have a decomposition?
- (ii) For which G does every set $E \in \mathcal{E}(G)$ have an orthogonal decomposition?
- (iii) For which G does every set $E \in \mathcal{E}(G)$ have a decomposition $E = D + F$ with either $D \in \mathcal{C}(G)$ and $F \in \mathcal{B}_{\text{fin}}(G)$ or $D \in \mathcal{C}_{\text{fin}}(G)$ and $F \in \mathcal{B}(G)$?
- (iv) When are such decompositions unique?

Note that (iii) \Rightarrow (ii) \Rightarrow (i) for G fixed, by Theorem 2.7.

If G is finite, a look at the dimensions of the subspaces involved shows that Problem 2.10 has a clear solution: the answer is ‘yes’ on all counts if $\mathcal{C}(G) \cap \mathcal{B}(G) = \{\emptyset\}$, while otherwise it is ‘no’ on all counts.

For G finite, non-empty elements of $\mathcal{C}(G) \cap \mathcal{B}(G)$ —edge sets that are both (algebraic) cycles and ‘cocycles’ (cuts)—are called *bicycles*. If G has no bicycle then, naturally, it is *pedestrian*; Figure 2.1 shows a pedestrian graph¹³ from [17]. We shall adopt these definitions verbatim for infinite graphs, but remark that only the finite bicycles are orthogonal to themselves. In view of Theorem 2.7 we would, ideally, hope for a positive answer to Problem 2.10 (ii) for all G that have no finite bicycles. As we shall see, this is too much to ask—but what exactly is possible is an open problem.

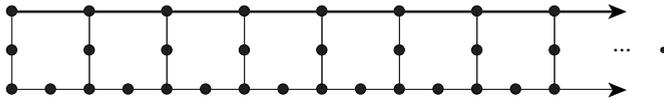


FIGURE 2.1. An infinite pedestrian graph

So what can we say when G is infinite? In general, neither (ii) nor (iii) of Problem 2.10 need hold even for pedestrian graphs. Indeed [12], consider an orthogonal decomposition $E = D + F$ of a set $E \subseteq E(G)$. Then $D \cap F$ is finite (and even), and E the disjoint union of a set $E_C := D \setminus (D \cap F)$ and a set $E_B := F \setminus (D \cap F)$ that can be turned into elements of $\mathcal{C}(G)$ or $\mathcal{B}(G)$, respectively, by adding a finite set (the set $D \cap F$). Moreover, if E is infinite then at least one of E_C and E_B must be infinite. Hence, for G to be a counterexample to Problem 2.10 (ii), all it takes is an infinite set $E \subseteq E(G)$ that has no infinite subset E_C that can be turned into an algebraic cycle by adding

¹³ To see it’s pedestrian, consider the leftmost edge of a hypothetical bicycle, and check on which sides of the cut that it defines the vertices near it come to lie.

finitely many edges, and no infinite subset $E_{\mathcal{B}}$ that can be turned into a cut by adding finitely many edges. It is easy to see that the edge set of the upper horizontal ray in the pedestrian graph of Figure 2.1 is such a set E in that graph.

Some alternative approaches to Problem 2.10 remain: to settle for (i); to prove (ii) and (iii) for specific graphs G ; or to prove the assertions from (ii) and (iii) for specific edge sets $E \subseteq E(G)$ in arbitrary graphs G .

The entire set $E = E(G)$, for example, always has an orthogonal decomposition, even into disjoint sets: this is a theorem of Gallai for finite graphs, whose extension to locally finite graphs [17] has an easy compactness proof of the kind we shall discuss in Section 3.

For singleton sets $\{e\}$, the following is known:

Theorem 2.11. [20, 25] *Let e be any edge of G .*

- (i) *Either $e \in E$ for some $E \in \mathcal{C} \cap \mathcal{B}$, or there are $D \in \mathcal{C}_{\text{fin}}$ and $F \in \mathcal{B}_{\text{fin}}$ such that $\{e\} = D + F$, but not both.*
- (ii) *Either $e \in E$ for some $E \in \mathcal{C}_{\text{fin}} \cap \mathcal{B}_{\text{fin}}$, or there are $D \in \mathcal{C}$ and $F \in \mathcal{B}$ such that $\{e\} = D + F$, but not both.*

Theorem 2.11 (i) implies a positive answer to Problem 2.10 (i) for pedestrian graphs G : the sets from \mathcal{C}_{fin} and \mathcal{B}_{fin} needed to generate the singleton sets $\{e\}$ form a thin family [20, Prop. 13], and their sum over all $e \in E$ is exactly E .

Corollary 2.12. *If $\mathcal{C}(G) \cap \mathcal{B}(G) = \emptyset$, then every set $E \subseteq E(G)$ has a decomposition $E = D + F$ with $D \in \mathcal{C}(G)$ and $F \in \mathcal{B}(G)$.*

What can we say if all we assume is that G has no *finite* bicycle? Then every singleton set $\{e\}$ has a decomposition $\{e\} = D + F$ as in Theorem 2.11 (ii). But unlike for pedestrian graphs, this cannot always be chosen orthogonal: there can be edges e all whose decompositions $\{e\} = D + F$ are such that both D and F (and hence also $D \cap F$) are infinite [20]. Moreover, the family of *all* the D and F that can be used in such singleton decompositions will not be thin, as soon as G has an infinite bicycle [12]. However we might still hope that for every concrete $E \subseteq E(G)$ we can choose these decompositions, one for every $e \in E$, so that the sets D and F used form a thin family. This would yield the result, stronger than Corollary 2.12, that every edge set in a graph without a finite bicycle has a decomposition.

Finally, uniqueness. For pedestrian graphs, any orthogonal decomposition $E = D + F$ of an edge set E will be unique: if $E = D' + F'$ was another, then $D - D' = F' - F$ would be a bicycle. For those pedestrian graphs that do have orthogonal decompositions, we thus have linear projections $\mathcal{E} \rightarrow \mathcal{C}$ and $\mathcal{E} \rightarrow \mathcal{B}$ mapping E to D and to F , respectively.

Problem 2.13. *Are there explicit descriptions of these projections?*

For finite graphs G there are such descriptions [9], but these involve counting the number of spanning trees of G —not overly helpful when G is infinite.

Finally, homology theory appears to suggest the following related question:

Problem 2.14. *For which G is there a natural ‘boundary operator’ $\partial: \mathcal{E}(G) \rightarrow \mathcal{V}(G)$ with kernel $\mathcal{C}(G)$, and a natural ‘coboundary operator’ $\delta: \mathcal{V}(G) \rightarrow \mathcal{E}(G)$ with image $\mathcal{B}_{\text{fin}}(G)$?*

I suspect that this simplicial attempt to marry our topological cycle space to the usual homological setup is in fact inadequate, and that the way to do this lies not in changing the boundary operators but in restricting the chains. A singular approach that does capture the topological cycle space as its first homology group, just as in finite graphs, and where some infinite chains are allowed but not every subset (or sum) of edges is a legal 1-chain, has been constructed in [40]. It allows for some infinite chains, but not all. In particular, its 1-chains do not include all subsets of $E(G)$, so it does not imply a solution to Problem 2.13 as posed above. See Section 5 for more.

It may also be instructive to study how the answer to the questions in this section, in particular to Problem 2.10, changes if we replace our coefficient ring \mathbb{Z}_2 with \mathbb{Z} or \mathbb{R} . There will be no bicycles then, which simplifies things. But other problems occur. With integer coefficients, for example, 1-dimensional simplicial (co-) chains even of a finite graph need not decompose into an algebraic cycle and a cut (both oriented): as one can readily check, the triangle with one oriented edge mapped to 1 and the others to 0 is a counterexample.

For infinite graphs and real coefficients, let $\vec{\mathcal{E}}(G)$ consist of only those functions $\psi: \vec{E}(G) \rightarrow \mathbb{R}$ for which $(\psi(\vec{e}) = \psi(\bar{e})$ for all $\vec{e} \in \vec{E}(G)$ and

$$\sum_{e \in E(G)} \psi^2(e) < \infty,$$

where $\psi^2(e) := \psi(\vec{e})^2 (= \psi(\bar{e})^2)$. For such $\psi, \psi' \in \vec{\mathcal{E}}(G)$ we may define

$$\langle \psi, \psi' \rangle := \sum_{e \in E(G)} \psi(\vec{e})\psi'(\vec{e}) \quad (< \infty),$$

where \vec{e} is the *natural orientation* of e , one of its two orientations that was picked and fixed once and for all when G was first defined. Now ψ and ψ' can be called *orthogonal* if $\langle \psi, \psi' \rangle = 0$, and we can study Problem 2.10.

2.5 Duality

For finite multigraphs G and G^* one calls G^* a *dual* of G if there is a bijection $e \mapsto e^*$ between their edge sets $E := E(G)$ and $E^* := E(G^*)$ such that every set $F \subseteq E$ satisfies

$$F \subseteq \mathcal{C}(G) \text{ if and only if } F^* \in \mathcal{B}(G^*), \quad (\dagger)$$

where $F^* := \{e^* \in E^* \mid e \in F\}$. If G is 3-connected, there is at most one such G^* (up to isomorphism), which is again 3-connected. In this case we think of G^* as being *defined* by (\dagger) , given G . By the finite version of Theorem 2.7, condition (\dagger) implies the corresponding condition with \mathcal{C} and \mathcal{B} swapped. Hence if G^* is a dual of G then G is a dual of G^* , and for 3-connected G we have $G^{**} = G$.

Can such duals exist also for locally finite infinite graphs? Geometric duals of plane graphs suggest not: the geometric dual of the $\mathbb{N} \times \mathbb{N}$ grid,¹⁴ for example, has a vertex of infinite degree. And indeed, Thomassen [80] proved that if G is a locally finite 3-connected graph that (even only) has a locally finite ‘finitary dual’—a graph G^* for which only the finite sets of edges in G and G^* are required to satisfy (\dagger) —then every edge of G must lie in exactly two finite peripheral circuits. In particular, the $\mathbb{N} \times \mathbb{N}$ grid has no such dual.

Hence in order to be able to define duals for all locally finite graphs that should have a dual—that is, for the locally finite planar graphs—we shall have to allow as duals at least some graphs with infinite degrees. And if we want duals to have duals, we have to define duals also for such graphs with infinite degrees. As we shall see in Section 5, allowing arbitrary countable graphs with infinite degrees leads to considerable difficulties. But, fortunately, Thomassen [80] proved also that any graph with (even only) a finitary dual must satisfy another condition:

Every two vertices can be separated by finitely many edges. (*)

Thus, no matter how we define duals: as soon as their finite edge sets are required to satisfy (\dagger) —which is certainly the minimum we shall have to ask—these graphs will have to satisfy $(*)$. On the other hand, the class of graphs satisfying $(*)$ is not closed under taking finitary duals, and finitary duals are not unique [13, 80]. Extending duality to infinite graphs thus seems to be fraught with problems.

In our topological setup, however, the problem has a most elegant solution. Let us say that vertex v *dominates* a ray R in G if G contains an infinite v – R fan (Fig. 5.1). It is easy to see that if v dominates one ray of an end it dominates all its rays; we then also say that it *dominates* that end ω . Assuming that G satisfies $(*)$, we can then define a space \tilde{G} similar to $|G|$, but with the difference that new points at infinity are added only for the undominated ends, while the dominated ends are made to converge to the vertex dominating them.¹⁵ Figure 2.2 shows a graph G whose two ends are both dominated by the vertex v . Thus, \tilde{G} has no added ends, but the double rays ‘bend round’ so that both their tails converge to v .

In this space \tilde{G} , duality works as if by magic. We define $\mathcal{C}(G)$ as the set of circuits (the edge sets of circles) in \tilde{G} , define $\mathcal{B}(G^*)$ as the set of cuts of G^* (as

¹⁴ ... naively defined just for this example; see Problem 2.17 below

¹⁵ Note that, by $(*)$, this vertex is unique. We shall define \tilde{G} more formally in Section 5.

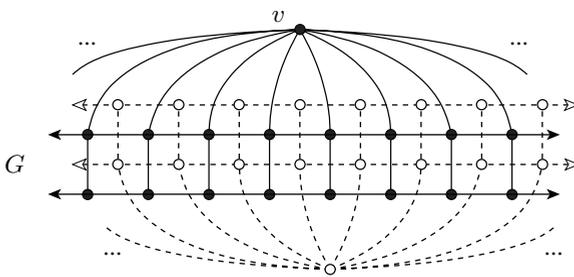


FIGURE 2.2. A self-dual graph with infinite circuits but no added ends

before), and call G^* a *dual* of G if there is a bijection $e \mapsto e^*$ between their edge sets satisfying (\dagger) for all sets $F \subseteq E(G)$, finite or infinite. If G is 3-connected (and satisfies $(*)$), then any dual G^* of G is unique and also satisfies $(*)$, and we have $G^{**} = G$ witnessed by the map $e^* \mapsto e =: e^{**}$ [13].

For the purpose only of the next theorem, let us call a (graph-theoretical) spanning tree of G satisfying $(*)$ *acirclic* if its closure in \tilde{G} contains no circle (and similarly for G^*). Note that unless G has no dominated ends, normal spanning trees of G will *not* be acirclic: they contain a ray from every end, and a ray from a dominated ends that starts at its dominating vertex forms a circle in \tilde{G} . However, one can show that every connected graph satisfying $(*)$ has an acirclic spanning tree [13].

Theorem 2.15. [13] *Let $G = (V, E)$ and $G^* = (V^*, E^*)$ be connected multigraphs satisfying $(*)$, and let $*$: $E \rightarrow E^*$ be a bijection. Then the following two assertions are equivalent:*

- (i) G and G^* are duals of each other, and this is witnessed by the map $*$ and its inverse.
- (ii) Given a set $F \subseteq E$, the graph (V, F) is an acirclic spanning tree of G if and only if $(V^*, E^* \setminus F^*)$ is an acirclic spanning tree of G^* .

(For (ii), define $F^* := \{e^* \in E^* \mid e \in F\}$ as before.)

Can we generalize Theorem 2.15 to topological spanning trees that are not necessarily graph-theoretical trees? A positive answer to this question would be in line with the philosophy of our theory, but the answer is not obvious. Once we use ends as ‘connectors’ (as we do in a topological spanning tree), we should treat them on a par with edges: we should have a bijection between the ends of \tilde{G} and those of \tilde{G}^* , and use each end from this common set of ends in precisely one of the two trees. But no matter how one defines the ‘ends of \tilde{G} ’, whether one takes just the undominated ends or the undominated ends plus the dominating vertices (i.e., the set of all limit points of rays), there need not be a bijection between the ends of \tilde{G} and those of \tilde{G}^* : the $3 \times \mathbb{Z}$ grid, for example,

is a 3-connected graph with two ends, whose dual has no undominated end and only one dominating vertex.

Surprisingly, things work better if we, initially, continue to consider the ends of G and G^* in the original spaces $|G|$ and $|G^*|$. Bruhn and Stein [23] have shown that, if G is 2-connected, there is a homeomorphism between the subspaces Ω of $|G|$ and Ω^* of $|G^*|$ that is *compatible* with the duality map $e \mapsto e^*$ from E to E^* in the sense that, for every set $F \subseteq E$, an end $\omega \in \Omega$ is an accumulation point of $\bigcup F$ in $|G|$ if and only if ω^* is an accumulation point of $\bigcup F^*$ in $|G^*|$. The question, then, could be roughly as follows: is it true that no matter how we form a spanning tree in \tilde{G} from some of the edges and ends of G (taken in $|G|$), the other edges and ends (taken in $|G^*|$) will form a spanning tree in \tilde{G}^* ? And that, conversely, if this is the case for two graphs G and G^* then these form a dual pair?

Figure 2.3 shows a dual pair of graphs illustrating this. The grey tree in G^* needs the ‘shared’ end in order to be connected, and adding the end to it does not create a circle. The dual black tree, however, is already connected (without the end), and adding the end there would create a circle. Hence there is a unique way of assigning the end to one of the two trees, so that both are spanning and acirclic in their respective graphs.

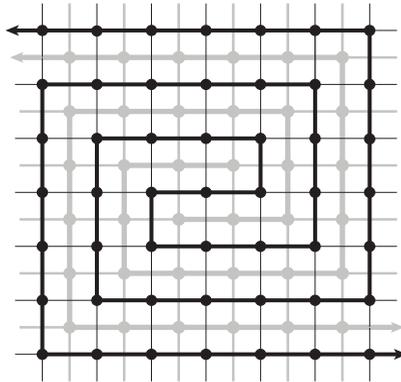


FIGURE 2.3. Dual ‘spanning trees’: the grey tree contains the shared end, the black tree does not

Let us make this precise in greater generality. Given a set $F \subseteq E$ of edges and a set $\Psi \subseteq \Omega$ of ends, let us say that the subspace $\bigcup F \cup \Psi$ of $|G|$ is *formed by* F and Ψ . Given a subspace X of $|G|$, let \tilde{X} denote the quotient space of X obtained by identifying every vertex in X with all the ends in X that it dominates (in G). (For $X = G$ this yields precisely the definition of \tilde{G} given informally earlier.) For example, if G is the fan shown on the left in Figure 5.1, we might take as X the space formed by the set F of fat edges and the unique end of G . Then \tilde{X} is a circle. However, if we take as X the space formed only by the set F of fat edges, without the end, then \tilde{X} is homeomorphic to

the half-open interval $[0, 1)$. Call \tilde{X} a *spanning tree* of \tilde{G} if it contains all the vertices of \tilde{G} , is arc-connected, and contains no circle.

Theorem 2.16. [16] *Let $G = (V, E, \Omega)$ and $G^* = (V^*, E^*, \Omega^*)$ be 2-connected multigraphs satisfying $(*)$, and let $*$: $E \rightarrow E^*$ and $*$: $\Omega \rightarrow \Omega^*$ be compatible bijections. Then the following two assertions are equivalent:*

- (i) *G and G^* are duals of each other, and this is witnessed by the maps $E \rightarrow E^*$ and $\Omega \rightarrow \Omega^*$ and their inverses.*
- (ii) *Whenever the subspace X of $|G|$ formed by subsets $F \subseteq E$ and $\Psi \subseteq \Omega$ is such that \tilde{X} is a spanning tree of \tilde{G} , the sets $E^* \setminus F^*$ and $\Omega^* \setminus \Psi^*$ form a subspace Y of $|G^*|$ such that \tilde{Y} is a spanning tree of \tilde{G}^* .*

Here, as before, $F^* := \{e^* \in E^* \mid e \in F\}$ and $\Omega^* := \{\omega^* \in \Omega^* \mid \omega \in \Omega\}$. The obvious implications for infinite matroids which Theorem 2.7, Theorem 2.15 and 2.16 suggest will be a topic in Section 5.

For finite graphs, the theory of algebraic duality outlined above corresponds to a theory of *geometric duality*: two connected graphs are duals of each other if and only if they can be drawn in the plane so that precisely the corresponding edges cross and vertices correspond bijectively to faces of the dual.¹⁶

For our graphs satisfying $(*)$, this is not so obvious: how, for instance, should we draw the two dual graphs of Figure 2.2 so that their rays converge as they do in \tilde{G} and \tilde{G}^* ? In particular, where in such a drawing can we put the two vertices of infinite degree, the limit points of those rays? For now, setting up geometric duality in a satisfactory way remains an unsolved problem:

Problem 2.17. *Define geometric duals for infinite graphs satisfying $(*)$ compatibly with their algebraic duals.*

Perhaps the solution to Problem 2.17 lies in an approach similar to that for Theorem 2.16: to apply duality ‘before identification’ rather than after, that is to say, to look for geometrically dual embeddings of $|G|$ and $|G^*|$ rather than of \tilde{G} and \tilde{G}^* . The drawback of such an approach is that $|G|$ and $|G^*|$ need not be compact as soon as we refine their topologies enough to make their embeddings in the plane continuous.¹⁷ This, then, raises issues such as how exactly to define a face.

¹⁶ In particular, graphs with a dual are planar. This is Whitney’s theorem, which extends to locally finite graphs satisfying $(*)$ if and only if duality is defined as we did: with \mathcal{C} as the topological cycle space of G based on \tilde{G} . See Section 4.

¹⁷ Since $|G|$ fails to be Hausdorff as soon as G has a dominated end, we need to refine our current topology, known as VTOP, to a topology called MTOP; see Section 5.1–2. Under MTOP, $|G|$ is compact only if G is locally finite (in which case MTOP agrees with VTOP).

3. Proof Techniques

Once more, let G be an infinite, locally finite, connected graph. The aim of this section is to bring the reader new to the field up to date with its collective memory of techniques. Although each of the deeper proofs of existing theorems naturally has its own difficulties and ways of overcoming them, there is by now a body of basic approaches that can be described—as well as some common pitfalls the novice might like to hear about before plummeting their depths themselves.

We shall concentrate on proofs that establish the existence of infinite substructures of $|G|$ such as arc, circles, or topological spanning trees. Such a proof usually has two parts: the construction of this structure, usually by a limit process, and the proof that the object thus constructed does what it is intended to do.¹⁸ We shall only discuss the construction part. If the construction is done well, it can happen that the proof of correctness becomes easy. This does not always happen, and if it does, it may simply mean that the problem is intricate enough that ‘blind’ constructions will fail: that any difficulties that may arise in the proof have to be anticipated by the construction. Still, the basic techniques that are common to many proofs tend to lie at the construction level, while difficulties arising at proof level tend to be individual to the problem. Our focus on the construction part of proofs reflects this phenomenon; it is not a deliberate restriction.

3.1 The direct approach

The direct, or naive, approach to the construction of standard subspaces with certain desired properties, such as an arc, a circle or a topological spanning tree, is to obtain them as the union of nested finite subgraphs defined explicitly, or as the intersection of a nested sequence of explicitly defined superspaces. For example, we might try to construct a spanning tree as a union of finite trees $T_0 \subseteq T_1 \subseteq \dots$ eventually covering all the vertices of G . The union T of these T_n will be a tree, but its closure in $|G|$ may well contain a circle: unlike finite cycles, this will arise only at the limit step, and it may not be clear how to choose the T_n so that this does not happen.

If we want this naive approach to work, we have to express the statement to be proved in a finitary way—for example, as statements about finite sets of edges. The characterization of topological connectedness in Lemma 1.4 is a case in point; the characterization of cycle space elements in Theorem 2.7(i) is another. Unlike the circle in our example, finitary properties do not appear or disappear at limits, which makes it easy to prove such statements by induction or an application of Zorn’s lemma.

¹⁸ For example, we might be constructing a *Hamilton circle*, a circle in $|G|$ passing through all the vertices. The construction part might define a map $\sigma: [0, 1] \rightarrow |G|$ that is, by definition, injective at vertices and inner points of edges (except that $\sigma(0) = \sigma(1)$), while we might still have to show that σ is continuous and injective at ends.

For instance, we have already seen the construction of a topological spanning tree ‘from above’, in the proof of Lemma 1.5. In that proof we used the fact that connectedness of a standard subspace is a finitary property (Lemma 1.4), and that topological spanning trees are edge-minimal connected standard subspaces of $|G|$ containing all its vertices. The proof of the following lemma, which we shall need in Section 4, is very similar:

Lemma 3.1. *Let X be a standard subspace of G .*

- (i) *If X is k -edge-connected, we can delete edges from X to obtain a standard subspace Y that is edge-minimally k -edge-connected.*
- (ii) *If X is k -vertex-connected, we can delete edges from X to obtain a standard subspace Y that is edge-minimally k -vertex-connected.*

Proof. (i) To obtain Y from X , we go through the edges of X in turn, deleting the edge under consideration if and only if this does not spoil the k -edge-connectedness of the current space. To show that Y is still k -edge-connected (in which case it will clearly be minimally so), we consider a hypothetical set F of fewer than k edges whose removal disconnects Y . By Lemma 1.4, the space $Y \setminus \overset{\circ}{F}$ meets both sides of some finite cut of G in which it has no edge. But $X \setminus \overset{\circ}{F}$ has an edge in this cut, because it is connected. All its edges in the cut, however, were deleted in the construction of Y , and the edge deleted last should not have been deleted.

(ii) The proof of the vertex-connectivity case is analogous: we again delete edges one by one, but instead of an edge set F and the space $Y \setminus \overset{\circ}{F}$ we consider a separating set U of fewer than k vertices and the standard space obtained from Y by deleting U and its incident edges. □

The analogue of Lemma 3.1 for k -connected subspaces—those that cannot be made disconnected by the deletion of fewer than k vertices or ends—cannot be proved in the same way. For arbitrary substandard subspaces it is false (consider the double ladder, without the ends, for $k = 2$), but for standard subspaces it is an open problem.¹⁹ For example, if we delete all the broken edges in the graph of Figure 3.1 in ω steps, the space will lose its 3-connectedness (since deleting the two ends will then disconnect it), but does so only at the limit step.

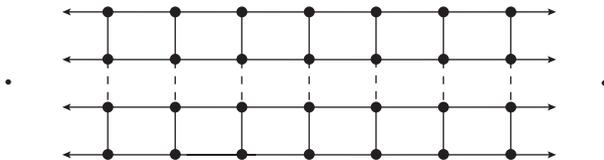


FIGURE 3.1. The subspace obtained by deleting all the broken edges is no longer 3-connected.

¹⁹ We shall discuss this problem in Section 4.1.

As we shall see in Section 4, many applications of the object $|G|$ have the following form: there is an existence statement about finite graphs that fails for infinite graphs—for example, a theorem asserting the existence of a Hamilton cycle—but a topological analogue works in $|G|$. In our example, it might be that the same conditions (such as being planar and 4-connected) that force a Hamilton cycle in a finite graph force a Hamilton circle in $|G|$. Rather than proving the infinite theorem from first principles, we might try to use the finite result for the proof of its infinite analogue. We would then want to express $|G|$ as a limit of finite graphs G_0, G_1, \dots that satisfy the assumptions (such as planarity and 4-connectedness) if G does, use the finite theorem to find the desired structures in the G_n , and then take a limit of these to obtain the analogous structure in $|G|$.

Experience has shown that defining these G_n as induced subgraphs of G , e.g. those on its first n vertices, is *not* often a good approach: a badly chosen $G_n \subseteq G$ is too oblivious of how it lies inside G . (There might, for example, exist G_n -paths²⁰ in G between vertices of G_n that are easily separated in G_n .) However, there is a standard approach using finite ‘nested’ minors rather than subgraphs as G_n , which we define now.

Let v_0, v_1, \dots be an enumeration of the vertices of G . For each $n \in \mathbb{N}$ let $S_n := \{v_0, \dots, v_n\}$, and write G_n for the minor of G obtained by contracting every component of $G - S_n$ to a vertex, deleting any loops that arise in the contraction but keeping multiple edges. The vertices of G_n contracted from components of $G - S_n$ will be called the *dummy vertices* of G_n .

Note that every cut of G_n is also a cut of G . Conversely, a cut of G is a cut of every G_n that contains all its edges. This is an important feature, which distinguishes these minors from a similar sequence of finite subgraphs exhausting G . A topologically connected standard subspace of $|G|$, for example, will induce a connected subgraph in each G_n (cf. Lemma 1.4), an element of the cycle space of G will induce a cycle space element of G_n (cf. Theorem 2.5.), and so on. We shall appeal to these properties again in our discussion of homology in Section 5.3.

As a first application of these G_n let us show that they allow us to construct a topological spanning tree directly ‘from below’, using the direct approach.²¹

Lemma 3.2. *G has a topological spanning tree.*

Proof (by explicit construction from below).

Pick spanning trees T_n of G_n recursively so that $E(T_{n+1}) \cap E(G_n) = E(T_n)$.²² We claim that $T := \bigcup_n E(T_n)$ is a topological spanning tree of G .

²⁰ See [32] for the notion of an H -path in G for a subgraph $H \subseteq G$.

²¹ The proof of Lemma 3.2 can be generalized to a proof ‘from below’ of the full statement of Lemma 1.5 for the existence of topological spanning trees in standard subspaces [16].

²² There is a unique way of doing this. Unless v_{n+1} forms a component of $G - S_n$ by itself (in which case $T_{n+1} = T_n$), the graph G_{n+1} is obtained from G_n by expanding the dummy vertex of G_n contracted from the component containing v_{n+1} to a star with centre v_{n+1} and

To show that T is connected (and hence arc-connected, by Lemma 1.2), we have to check that every finite cut of G contains an edge from T . It does, because it is also a cut of some G_n , and T_n has an edge in this cut.

Now suppose that T contains a circle, C say. Pick vertices $u, v \in C$, and choose n large enough that $u, v \in S_n$. Let F be a fundamental cut of T_n separating u from v . Since this is also a cut of G and C meets both its sides, C has an edge in F (Lemma 1.4). Deleting this edge from C leaves an arc, which still has an edge in F (by the same argument). Hence $C \subseteq T$ has two edges in F , contradicting the fact that $E(T) \cap E(G_n) = E(T_n)$ and T_n has only one edge in F . \square

Finally, there is a type of construction, very common in finite graph theory, which is rather naive in the context of $|G|$ but should not go unmentioned: the construction of paths or arcs by ‘greedily moving along’. As an example, consider the proof of Theorem 2.3 (i), that every edge set $D \in \mathcal{C}(G)$ is a disjoint union of circuits. Let us just try to find the first circuit. If G is finite, we start from any edge in D and ‘greedily move along’ further edges of D : since D has an even number of edges at every vertex, we cannot get stuck until we reach our starting vertex again. We stop, however, when we hit the first previously visited vertex, having found our first circuit. We then delete this and repeat, eventually decomposing D into circuits.

When G and D are infinite, this approach can fail quite spectacularly. For example, D might be the wild circuit of Figure 1.1. By ‘moving along’ its edges we would run into an end, from which there would be no escape, not even by an inverse ray. The linear order structure of the double rays on the wild circle indicates that the only way to obtain it as a limit is to ‘approach all its points at once’—which is exactly what we shall do using our sequence of G_n .

3.2 The use of compactness

A typical compactness proof, such as of the Erdős-De Bruijn theorem on colouring, is a way of making consistent choices for all the finite substructures of an infinite structure, and then deducing from the fact that all the finite substructures have a certain property that so does the entire structure.

In our context, only the simplest results have compactness proofs quite like this. But we often use compactness for the construction part of a proof, including the more difficult ones: to define a limit object (e.g., a subspace X) in such a way that all of certain finite substructures (e.g., the finite cuts of G) relate to the limit object in a certain desired way (contain an edge of X). This may be less than we actually want to prove (e.g. that X is arc-connected)—but it may be a good start.

The best examples for this are infinite analogues of theorems asserting for a finite graph the existence of a path or cycle that can be found by a process

some new dummy vertices of G_{n+1} as leaves. We obtain T_{n+1} from T_n by adding the edges of this star. In other words, we do not add edges incident with a vertex in S_n .

of ‘focussing in’ on some part of a graph: in the same process may converge to a limit object in $|G|$ with similar properties. Bruhn’s proof [11] of Theorem 2.3 (ii), for example, works in principle like Tutte’s own proof of his finite theorem. Or think of the ‘focussing’ in on a component in Thomassen’s proof that every finite 3-connected graph has an edge whose contraction preserves 3-connectedness. (See [32] for both.)

To illustrate how such a compactness proof typically works, or might fail, let us take another look at the proof of Theorem 2.3 (i). We were trying to find a single circuit C in a given set $D \in \mathcal{C}(G)$ (which might be the wild circle, so that we cannot find C by ‘moving along’). Our approach now is to construct C as a limit of circuits C_n in the G_n . Since D meets every cut of G_n evenly (because that cut is also a cut of G), we have $D_n := D \cap E(G_n) \in \mathcal{C}(G_n)$. By the trivial finite version of Theorem 2.3 (i), we find a circuit $C_n \subseteq D_n$. Will these C_n , or at least a subsequence, tend to a limit $C \subseteq D$ that is a circle in $|G|$?

Not necessarily. For a start, the limit might exist²³ but be empty. Indeed, depending on the structure of D , we might have chosen the C_n so badly that no edge lies eventually in C_n .²⁴ But there is a simple way to prevent this: by insisting that every C_n contain some fixed edge $e \in D$.

Assuming this, let us use the infinity lemma to construct a limit of these C_n . To apply the lemma we need to define ‘predecessors’, which is easy: let the n th finite set in the setup of the lemma contain not only C_n but also all edge sets of the form $C_n^m := C_n \cap E(G_m)$ with $m > n$. Since the cuts of G_n are also cuts of G_m , these sets will meet every such cut evenly and hence lie in $\mathcal{C}(G_m)$. We may now define the predecessor relation for the infinity lemma as inclusion (of edge sets), and obtain the desired limit C as a union of certain nested sets $C_n^m \in \mathcal{C}(G_m)$ found by the lemma, one for each n . Then C , too, meets every finite cut of G evenly, because that cut is also a cut of some G_m and the selected C_n^m meets this cut evenly. Hence $C \in \mathcal{C}(G)$, by Theorem 2.7 (i).

But is C a circuit? Unfortunately, it need not be: we cannot in general prevent the C_n^m from traversing dummy vertices more than once, in which case the closure \overline{C} of their union can be a union of circles that are vertex-disjoint but meet at ends. Figure 3.2 indicates an example where this happens: the sequence of C_n there is formed by alternately extending them in two different places, so that both extensions tend to the unique end, yielding a limit in which that end has degree 4.

So what did we achieve? As the limit of our circuits C_n we found just another algebraic cycle in $\mathcal{C}(G)$, but not a circuit. For our present problem this seems like next to nothing: after all, we started with an algebraic cycle $D \in \mathcal{C}(G)$, looking for a circuit $C \subseteq D$. But in general, the method does have

²³ We have not made this precise yet.

²⁴ For example, let G be the plane $\mathbb{Z} \times \mathbb{Z}$ grid, and choose as C_n the boundary circuit of the outer face of G_n .

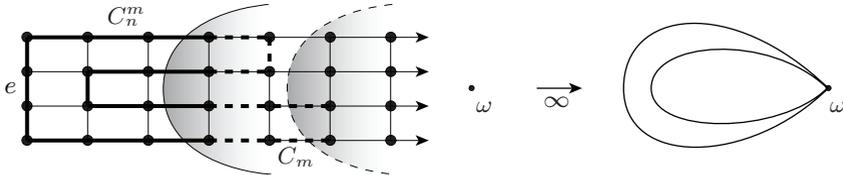


FIGURE 3.2. Finite cycles tending to a double circle. The shaded regions contract to dummy vertices.

its uses: it shows that from algebraic cycles $C_n \in \mathcal{C}(G_n)$ —we never used that they are actually circuits—we can always obtain an algebraic cycle $C \subseteq \bigcup_n C_n$ by compactness, even if the C_n were not induced by some given $D \in \mathcal{C}(G)$. And we can say a little more about C too: there is a subsequence of the C_n (consisting of those C_m that gave rise to one of the C_n^m in our nested sequence picked by the infinity lemma) such that every edge of C lies eventually on every C_n in this subsequence.

Moreover, the set $C \subseteq \mathcal{C}(G)$ we found (now using that the C_n are circuits) does have one advantage over D : the standard subspace $\overline{C} \setminus \dot{e}$ is connected. Indeed, every $C_m \setminus \{e\}$ spans a connected subgraph in G_m , and hence so does $C_n^m \setminus \{e\}$ in G_n . The connectedness of $\overline{C} \setminus \dot{e}$ now follows from $C_n^m \subseteq C$ and Lemma 1.4. We can therefore use Lemma 1.2 to find an arc in $\overline{C} \setminus \dot{e}$ between the endvertices of e , which together with e will form the desired single circuit in D containing e —solving our problem.²⁵

3.3 Constructing arcs directly

In our compactness proof discussed above, we obtained as a limit of finite cycles $C_n \subseteq G_n$ an edge set $C \in \mathcal{C}(G)$ with \overline{C} connected, and we had to invoke Lemma 1.2 in order to find inside \overline{C} an arc between the endpoints of our fixed edge e . Could we have obtained such an arc directly, by using compactness in a more subtle way—e.g. by choosing the edge sets to which we applied the infinity lemma more carefully?

Problem 3.3. *Given two vertices $x, y \in G$ and, for all n large enough, an x - y path $P_n \subseteq G_n$, construct by a direct compactness argument an x - y arc in $|G|$ with edges in $\bigcup_n E(P_n)$.*

This problem is still unsolved. What we can do is construct a topological x - y path in $|G|$ by direct compactness. This path will traverse vertices at most once, but it may be non-injective at ends (as in Figure 3.2). In order to obtain the desired arc, we therefore still have to invoke a theorem from general topology: that the image of any x - y path in a Hausdorff space contains an x - y arc.

²⁵ The proof of Theorem 2.3 (i) follows easily now: we just collect circuits $C \subseteq D$ inductively and delete them, to eventually decompose D , starting each new circuit with the next uncovered edge of D in some fixed enumeration.

Before we apply this technique to obtain such a partial solution to Problem 3.3, let us illustrate it by an exercise. Let T be the infinite binary tree. The task is to construct a topological path $\sigma: [0, 1] \rightarrow |T|$ that starts and ends at the root and traverses every edge exactly twice, once in each direction. For $k \in \mathbb{N}$, let T^k denote the finite subtree of T consisting of its first k levels. We shall obtain σ as a limit of analogous paths σ_k around these T^k , as follows.

Let σ_0 be the constant map from $[0, 1]$ to the 1-vertex graph T^0 . Assume inductively that we have defined for $i = 0, \dots, k$ paths $\sigma_i: [0, 1] \rightarrow T^i$ such that the inverse image under σ_i of each leaf of T^i is a non-trivial interval, that the inverse image of any other point of T^i is a single point, and that σ_{i+1} agrees with σ_i on all points of $[0, 1]$ that σ_i maps to points of T^i other than leaves. (Thinking of σ_i as a journey with $[0, 1]$ measuring ‘time’, we thus ask that σ_i should pause for a while at every leaf of T^i , but not elsewhere.) In order to obtain σ_{k+1} from σ_k , we just expand every pause at a leaf t of T^k to a little path $tt'tt''t$ in T^{k+1} from t to its two successors t', t'' and back, pausing again at each of t' and t'' .

To define our desired limit σ of these σ_k , we first look at all those $x \in [0, 1]$ for which the σ_k eventually agree, and let $\sigma(x) := \sigma_k(x)$ for these x (with k large enough). For any other x , every $\sigma_k(x)$ will be a leaf of T^k , and as $k \rightarrow \infty$ these $\sigma_k(x)$ form a ray in T ; we let $\sigma(x)$ be the end of that ray. It is straightforward from the definition of the topology of $|T|$, but still instructive, to check that this map $\sigma: [0, 1] \rightarrow |T|$ is continuous at these latter points x mapped to ends. Indeed, every basic open neighbourhood of an end $\omega = \sigma(x)$ has the following form: for some vertex t on the ray R in ω starting at the root, it is the set of all points in $|T|$ strictly above t . Let t' be the upper neighbour of t on R . If k is the level of t' in T , then t' is a leaf of T^k , and all the points in the subinterval of $[0, 1]$ that σ_k maps to this leaf t' will be mapped by all σ_ℓ with $\ell > k$ to points of G that still lie above t . Hence σ will map all such points to the closure of this set (but not to t), i.e. to the open neighbourhood of ω we started with.

This technique for constructing a topological path directly as a limit of maps can be combined with compactness, as follows. Consider Problem 3.3. Every x - y walk in G_m , including the given path P_m , induces for every $n < m$ an x - y walk in G_n consisting of those of its edges that already lie in G_n , traversed in the same order and direction. In order to obtain an x - y arc in $|G|$, we first use the infinity lemma to select from these walks a sequence $(W_n)_{n \in \mathbb{N}}$ of x - y walks W_n each inducing the previous. We then parametrize these walks in turn for $n = 1, 2, \dots$, as in our binary tree example: by topological paths σ_n ‘pausing’ at every dummy vertex of G_n (but only there), and so as to agree with σ_{n-1} at all points that σ_{n-1} does not map to a dummy vertex. The limit of these σ_n will be a topological x - y path σ in $|G|$, defined as as earlier, since sequences of dummy vertices of the G_n ($n \rightarrow \infty$) that correspond to nested components of $G - S_n$ define a unique end, and these components form a neighbourhood basis for that end.

Note that if we insist that the lengths of pauses tend to zero as n tends

to infinity, the inverse image of an end under the limit path σ will be a totally disconnected subset of $[0, 1]$. But I do not know how to set up the infinity lemma in such a way that σ will be injective at ends: we have found an x - y path in $|G|$, but not the desired arc.

The construction of a topological Euler tour in $|G|$ for the proof of Theorem 2.5 (ii) works exactly like this, too.

4. Applications

A particularly satisfying feature of the topological approach to infinite graphs via ends is its applicability. Considering arcs and circles in $|G|$ instead of paths and cycles in G enables us to extend classical theorems from finite graph theory whose naive extension to infinite graphs fails.

In this section we give an overview of three areas in which this phenomenon is particularly striking: an emerging ‘extremal’ theory of infinite graphs; applications of the cycle space (eg. to characterizations of planarity), and infinite electrical networks. Let G be a locally finite, connected infinite graph.

4.1 Extremal infinite graph theory

Following Bollobás [10], let us take the wider notion of *extremal* graph theory to refer to how graph invariants (usually numerical) interact. A typical question is how one invariant can be forced up or down in an arbitrary graph by making assumptions about what values another invariant takes on that graph. Particular emphasis is usually given to pairs of invariants of which one is ‘local’ and the other ‘global’. The following variant of Turán’s original such question is a good case in point: for which function $f(r)$, if any, can we force the presence of a K_r minor in an arbitrary graph by assuming only that it has average degree at least $f(r)$?

This type of question makes no immediate sense for an infinite graph, because there is no obvious notion of ‘average degree’. So let us replace it by ‘minimum degree’. Now the question does make sense, but unlike in the finite case it has a trivial negative answer: there are locally finite *trees* of arbitrarily large minimum degree. Clearly, these have no K_r minor for $r > 2$.

What, then, makes infinite graphs so different in this respect? Since there are finite graphs of both minimum degree and girth arbitrarily large [32], also finite graphs of large minimum degree can look *locally* like trees. However, a large finite tree also has many leaves, whose low degrees push its average degree back below 2. Requiring a high average or minimum degree in a finite graph therefore also has another effect, in addition to just fast growth: it forces us to ‘wrap up’ those leaves. The classical theorem of Mader [61] that a large enough average degree forces any desired minor can thus be restated as saying that, no matter how we do this ‘wrapping up’ (whether by adding edges, say, or by identifying vertices), we will always create such a minor.

The key idea for infinite graphs now is that compactifying a locally finite graph by adding its ends wraps it up in a similar way, albeit at infinity. And this wrapping-up can help us restore properties of finite graphs that are lost in the naive transition from finite to infinite.

In fact, most of the time this restoration happens automatically: often, all we have to do is allow circles when we used to wish for a cycle, and arcs when we used to desire a path. These objects tend to exist as limits (or inside the limits) of the cycles or paths whose existence was asserted by the finite theorems that now fail. More precisely, this works whenever the cycles or paths whose existence is claimed in the finite theorem can be found there by a process of ‘focusing in’, as discussed in Section 3. We shall see some more examples of this phenomenon later in this section.

First, however, let us take advantage of ends as ‘wrapping’ in a more direct way: by explicitly requiring them to have large degrees and seeing what substructures this can force.

Consider again the aim of forcing a K_r minor, for $r = 4$ say. In other words, we wish to force a TK_4 subgraph by assuming some large enough minimum degree. Assume for simplicity that our graph G is 2-connected. If G is finite, a minimum degree of $\delta(G) \geq 3$ will force a TK_4 subgraph. Indeed, as $\kappa(G) \geq 2$ we can construct G starting from a cycle and adding paths, one at a time, that share only their endvertices with the graph constructed so far. If $G \not\supseteq TK_4$, those two vertices lie on a common earlier path, and they do not ‘cross’ other paths grafted later on to this same earlier path; see Figure 4.1. This makes it easy to see that after every construction step there is still a vertex of degree 2.

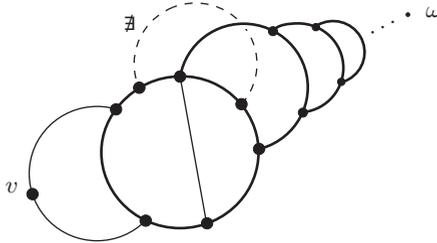


FIGURE 4.1. A vertex or end of degree 2 in a graph $G \not\supseteq TK_4$

If G is infinite, however, we can go on grafting new paths on to any vertices of degree 2, and after ω steps all such vertices have disappeared. But any graph obtained in this way will have an end of vertex-degree 2. (In fact, every end will have vertex-degree 2.) Indeed, using the infinity lemma we can find a sequence of paths each grafted on to the previous one. These paths converge to an end, which is easily seen to have vertex-degree 2. (Its edge-degree may be higher.) We can therefore get an analogue of the finite theorem, that $\delta \geq 3$ forces a TK_4 subgraph, after all: every graph whose vertices *and ends* all have (vertex-) degree at least 3 contains a TK_4 .

So what kind of substructures can be forced by assuming that both vertices and ends have large degree? The first theorem in this vein is due to Stein:

Theorem 4.1. [77] *Let G be a locally finite graph.*

- (i) *If $\delta(G) \geq 2k^2 + 6k$ and every end of G has vertex-degree at least $2k^2 + 2k + 1$, then G has a $(k + 1)$ -connected subgraph.*
- (ii) *If $\delta(G) \geq 2k$ and every end of G has edge-degree $\geq 2k$, then G has a $(k + 1)$ -edge-connected subgraph.*

The bounds in Theorem 4.1 are close to best-possible; see [77].

Theorem 4.1 uses vertex and end degrees to force not a concrete desired subgraph but just some subgraph from a desired class. This is an interesting variant of the original extremal problem, which is to force a concrete subgraph (or minor etc.) by global assumptions such as on average or minimum degrees. Forcing a concrete finite minor in an infinite graph by assuming that its vertex and end degrees are large, unfortunately, takes us little further than the K_4 example we saw earlier: already a K_5 minor cannot be forced in this way. This is because there are planar graphs with arbitrarily large degrees and only one end, of infinite degree: just take a regular tree of large degree, and add edges forming circuits D_i , one for every $i \in \mathbb{N}$, through all the vertices at distance i from the root.

Forcing K_r minors for $r \geq 5$ is possible, however, with another notion of end degrees: one is motivated less by topological considerations (such as to give every end on a circle degree 2 there) than extremal ones. Call an induced subgraph C of G a *region* of G if C is connected and its *edge-boundary* $B_e(C)$, the set of edges of G between C and $G - C$, is finite. Then also the *vertex boundary* $B_v(C)$ of C , the set of vertices of C incident with an edge in $B_e(C)$, is finite. Let us say that a nested sequence $C_1 \supseteq C_2 \supseteq \dots$ of regions *defines* an end ω of G if the sets \hat{C}_i form a neighbourhood basis of ω in $|G|$, i.e., if ω belongs to every C_i and $\bigcap_i C_i = \emptyset$. Now let the *relative degree* of ω in G be defined as the infimum, taken over all its defining sequences C_1, C_2, \dots of regions, of the numbers

$$\liminf_{i \rightarrow \infty} \frac{|B_e(C_i)|}{|B_v(C_i)|}.$$

Intuitively, the relative degree measures how many edges per boundary vertex the regions around an end ω will at least send away from ω eventually, just as the degree of a vertex measures the number of edges it sends away. Note that the edge-degree of an end is always at least its relative degree [78], but the vertex- and the relative degree of an end are not abstractly related. Relative end degrees in subspaces can be defined analogously, taking the end-defining regions in G but their vertex or edge boundaries in the subspace.

The notion of relative end degree was introduced by Stein. It works so well that we get the following blanket adaptation of results from finite sparse extremal graph theory:

Theorem 4.2. [78] *Let \mathcal{H} be a class of finite graphs, and assume that $d \in \mathbb{R}$ is such that every finite graph of average degree at least d has a minor in \mathcal{H} . Then every locally finite graph of minimum degree at least d whose ends all have relative degree at least d has a minor in \mathcal{H} .*

Note that, despite the use of end degrees, the finite minors found by Theorem 4.2 will be minors of finite subgraphs of G .²⁶ This motivates the following problem, where this will not be possible. Given a finite graph H and a standard subspace X of $|G|$, define the minor relation $H \preceq X$ just as for graphs, with closed and connected subspaces of X as branch sets (which may contain ends).

Problem 4.3. *Does Theorem 4.2 extend to subspaces?*

What about forcing infinite minors? From our earlier example we know that by assuming large vertex and end degrees (non-relative) we cannot force non-planar minors. But we can force most planar minors. Indeed, by a result of Halin [51, 32] every graph G with an end of infinite vertex-degree contains the half-grid $\mathbb{N} \times \mathbb{Z}$ as a minor, and similarly every graph with an end of large enough finite vertex-degree contains an $n \times \mathbb{N}$ grid (and hence any given finite planar graph) as a minor. On the other hand, there are planar graphs of arbitrarily large (finite) minimum degree and minimum vertex-degree for ends that do not contain the half-grid as a minor [78]. It would be interesting to find some natural strengthening of the degree assumption on ends that would force a planar graph to contain the full grid, or even to contain every locally finite planar graph, as a minor (cf. [33]). More generally:

Problem 4.4. *What infinite minors can be forced by assuming large vertex and end degrees (of any type)?*

Let us return to our theme of how certain paths or cycles (with some desired properties) whose existence in a finite graph is proved by some focusing process can fail to exist in an infinite graph G , because (in G itself) that process need not converge. Our aim will then be to show that such paths or cycles in finite minors G_n of G can tend to a limit that is an arc or circle in $|G|$ with the desired properties.

For example, consider in a finite 3-connected plane graph a maximal sequence of nested cycles (not necessarily disjoint). This sequence will end with a cycle that bounds a face. When we delete this cycle, it will not disconnect the graph. Cycles, in any graph, whose deletion does not reduce the connectivity of a graph by more than 3 are called *connectivity-preserving*. Cycles with the property that deleting their edges does not reduce the edge-connectivity of the graph by more than 2 are *edge-connectivity-preserving*. Such cycles exist in

²⁶ It is easy to see that *any* finite minor of a graph G is also a minor of a finite subgraph of G : just replace any infinite branchsets by finite subtrees that link the edges needed.

every finite graph, and they can be found by a process of ‘focusing in’, just as in the planar case; this was proved by Thomassen [81] for connectivity and by Mader [65, 66] for edge-connectivity.²⁷

In an infinite graph such cycles need not exist. Let us show this by constructing a counterexample, due to Aharoni and Thomassen [4]. This graph will be locally finite, and will depend on a given integer k ; so let us call it $AT(k)$. The graph $AT(k)$ combines two properties that no finite graph can have at the same time:

- it is k -connected (where k is as large as we like);
- deleting any cycle, or the edges of any cycle, disconnects the graph.

As a result of these two properties, the graphs $AT(k)$ are counterexamples to a number of statements which, for finite graphs, are well-known theorems. The existence of connectivity-preserving cycles is one of these, and we shall meet another below.

We shall construct $AT(k)$ inductively from copies of some fixed finite k -connected graph H . Let us choose H of girth at least k^2 ; then H has a set X of k vertices at distance at least k from each other. (For example, spread X around a shortest cycle in H .) The idea now is to build $AT(k)$ as a union of finite graphs $G_0 \subseteq G_1 \subseteq \dots$, where each G_{n+1} is obtained from G_n by grafting on to any non-separating cycles of G_n some new copies of H , to make them separating.

Formally, we begin with a copy G_0 of H . Let us assume inductively that we have constructed G_n in such a way that it is k -connected, that any G_{n-1} -path²⁸ in G_n has length at least k , and that the edges of any cycle contained in G_{n-1} separate G_n . We now consider separately every cycle C in G_n that does not lie in G_{n-1} . By our second assumption about G_n , the cycle C has some k edges that do not lie in G_{n-1} ; subdivide these once. (Thus, on an edge that lies on 3 such cycles of G_n we insert 3 subdividing vertices.) We now take k fresh copies of H specific to our choice of C , and identify their k -vertex sets X with those k subdividing vertices inserted on C . It is easy to check that G_{n+1} is again k -connected, that every G_n -path has length at least k (because it links two vertices from X in a copy of H), and that the edges of any cycle contained in the newly subdivided G_n (including those inside G_{n-1}) separate G_{n+1} . Clearly, $AT(k) = G_0 \cup G_1 \cup \dots$ has the two desired properties.

The Aharoni-Thomassen graph $AT(k)$ for $k \geq 4$ has neither connectivity-preserving nor edge-connectivity-preserving cycles. However, it has edge-connectivity-preserving circles, as indeed does every locally finite graph:

²⁷ We remark that the connectivity-preserving cycle C found by Thomassen [81] is induced. Hence if $C \subseteq G$ with $\kappa(G) = k + 3$, say, then every vertex of C sends at least $k + 1$ edges to $G - C$, which is k -connected. Thus, we also have the ‘mixed connectivity’ result that deleting only the *edges* of C reduces the (*vertex*-) connectivity of G by at most 3.

²⁸ See [32] for the notion of an H -path in G for a subgraph $H \subseteq G$.

Theorem 4.5. [16] *Let G be a locally finite graph, and $k \in \mathbb{N}$. If G is $(k+2)$ -edge-connected, then $|G|$ contains a circle C such that the subspace of $|G|$ obtained by deleting the edges in $E(C)$ is k -edge-connected.*

(Recall that, by Lemma 1.4, the graph G is $(k+2)$ -edge-connected if and only if $|G|$ is, and likewise for vertex-connectivity and general connectivity.) Theorem 4.5 extends to standard subspaces [16].

The analogues of Theorem 4.5 for deleting vertices and/or ends are open:

Problem 4.6.

- (i) *If G is $(k+3)$ -connected, does $|G|$ contain a circle C such that the standard subspace of $|G|$ obtained by deleting the vertices of C and all their incident edges is k -vertex-connected?*
- (ii) *If G is $(k+3)$ -connected, does $|G|$ contain a circle C such that the subspace of $|G|$ obtained by deleting C and the edges incident with its vertices is k -connected?*

There are also versions of these theorems and problems for deleting paths and arcs. Mader [65] proved that any two vertices of a finite graph G are linked by a path whose edges we can delete without reducing the edge-connectivity of G by more than 2. A construction very similar to that of $AT(k)$ provides a counterexample to this statement for locally finite graphs, but the corresponding statement for arcs in $|G|$ (joining two given vertices or ends) is true [16]. For vertex-connectivity, it is a well-known open problem of Lovász whether or not there is even a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that any two vertices of any $f(k)$ -connected finite graph can be linked by a path whose deletion leaves a k -connected graph. The corresponding statement for arcs in $|G|$, in the spirit of Problem 4.6, is also unknown.

Two classical theorems from finite extremal graph theory, due, respectively, to Halin and Mader, say that if a finite graph is edge-minimal with the property of being k -connected [52], or of being k -edge-connected [62], then it has a vertex of degree only k . Since every finite k -connected or k -edge-connected graph contains an edge-minimal such graph, these are fundamental results about the structure of all finite k -connected or k -edge-connected graphs.

Unlike its edge-connectivity version, the vertex-connectivity version of the above theorem remains true for infinite graphs [52] and $k \geq 2$: every edge-minimal k -connected graph has a vertex of degree k . However, it is no longer that interesting: as the double ladder shows for $k = 2$, an infinite k -connected graph need not contain an edge-minimal such graph. The justification for studying these minimal graphs, therefore, collapses.

However, if we extend our class of objects from graphs to all their standard

subspaces, then edge-minimal objects exist by Lemma 3.1.²⁹ For example, we can now delete all the rungs in the double ladder: what remains is disconnected as a subgraph, but 2-connected in both senses as a subspace that includes the ladder’s ends. The basic objects to investigate for a study of the k -connected or k -edge-connected locally finite graphs, therefore, are their edge-minimal k -vertex-connected or k -edge-connected standard subspaces—or their edge-minimal k -connected subspaces if they exist:

Problem 4.7. *Can we delete edges from any k -connected standard subspace of $|G|$ to obtain an edge-minimal k -connected standard subspace?*

The graph of Figure 3.1, for example, does have an edge-minimal k -connected standard subspace: just delete all but one of the broken edges in Figure 3.1, or delete all the broken edges in Figure 4.2.

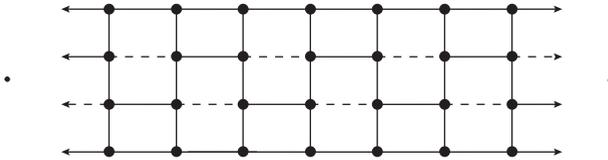


FIGURE 4.2. Deleting all the broken edges yields an edge-minimal 3-connected standard subspace.

Edge-minimal k -edge-connected or k -vertex-connected standard subspaces need not have a vertex of degree only k . Indeed, consider for $k = 2$ the cartesian product of a 3-regular tree T with an edge e ; deleting all the ‘rungs’ (the edges projecting to e) leaves a subspace consisting of the two copies of T glued together at the ends. This subspace X is edge-minimally 2-edge-connected and 2-vertex-connected, but every vertex is incident with 3 edges of X . Note, however, that all the ends have edge- and vertex-degree 2 in X .

This observation suggests the following infinite analogues to Halin’s and Mader’s finite theorems. Given k , let ‘ k -highly connected’ mean any one of ‘ k -edge-connected’, ‘ k -vertex-connected’ or ‘ k -connected’, and let ‘degree’ (for an end) mean any one of ‘edge-degree’, ‘vertex-degree’ or ‘relative degree’.

Problem 4.8. *Given $k \in \mathbb{N}$, does every edge-minimal k -highly connected standard subspace of $|G|$ contain a vertex or end of degree at most k ?*

²⁹ Having noticed that a given k -connected or k -edge-connected graphs need not contain an edge-minimal such subgraph, Halin [53, 54, 55] posed various problems to determine and study those that do. The shift to subspaces, coupled with Lemma 3.1, solves this problem in a much more satisfactory way.

Halin's finite theorem has been strengthened in various ways, and one can ask about infinite analogues also of those strengthenings. For example, Mader [63, 64] proved that in an edge-minimal k -connected graph every cycle contains a vertex of degree k . Among other things this implies that every subgraph of an edge-minimal k -connected graph has a vertex of degree at most k : either on a cycle, or else as a leaf. If Problem 4.8 has a positive answer, it will be natural to ask (in the same informal terminology as above):

Problem 4.9. *Given k , does every circle in an edge-minimal k -highly connected standard subspace X of $|G|$ contain a vertex or end whose degree in X is at most k ?*

For finite cycles, Problem 4.9 has been answered positively by Stein [78].

Little is known about graphs that are minimally k -connected with respect to deleting vertices rather than edges (let alone vertices or ends in subspaces). Does every k -connected locally finite graph have a k -connected subgraph that is minimal in the sense of not having a k -connected proper subgraph?

Perhaps this notion of minimality is too strong. Let us call a k -connected graph G *minimal* if every k -connected $H \subseteq G$ also satisfies $H \supseteq G$, that is, has a subgraph isomorphic to G . An analogous definition can be adopted for subspaces.³⁰

Once we do have a minimal k -connected graph or standard subspace, what can we say about its structure? Must such a graph be finite? Can it be $(k+1)$ -connected? Must it have a vertex or end whose degree is small in terms of k ? (By a theorem of Lick [60], every minimal k -connected finite graph has a vertex of degree at most $(3k-1)/2$.)

Problem 4.10. *Does every k -connected graph or standard subspace contain a minimal such object, for some suitable notion of 'minimal'? If so, what are its properties?*

See Stein [78] for more on this topic.

Another nice example of how arcs and circles in $|G|$ provide the natural setting for a classical finite theorem is tree-packing. The finite theorem here, due independently to Nash-Williams [68] and Tutte [84], says that a finite graph contains k edge-disjoint spanning trees unless its vertex set admits a partition, into ℓ sets say, such that G has fewer than $k(\ell-1)$ *cross-edges*, edges between different partitions sets. (See [32].)

For infinite graphs, the Aharoni-Thomassen graph $AT(2k)$ provides a counterexample to this statement, even to its corollary that $2k$ -edge-connected graphs have k edge-disjoint spanning trees [32]. Indeed, since $AT(2k)$ is $2k$ -edge-connected but the edge set of every cycle separates it, there can be no more

³⁰ Embeddings between subspaces should map vertices to vertices.

than two such trees: the edges of a fundamental circuit of one tree obtained by adding an edge of another tree would separate the graph, so no third tree could be spanning.

However, the finite tree-packing theorem has a topological analogue:

Theorem 4.11. [32] *The following statements are equivalent for all $k \in \mathbb{N}$ and locally finite multigraphs G :*

- (i) G has k edge-disjoint topological spanning trees.
- (ii) For every finite partition of $V(G)$, into ℓ sets say, G has at least $k(\ell - 1)$ cross-edges.

In particular, if G is $2k$ -edge-connected it has k edge-disjoint topological spanning trees.

Theorem 4.11 has an interesting history: while Nash-Williams had conjectured (incorrectly) that the finite tree-packing theorem ought to extend to countable graphs verbatim, Tutte anticipated Theorem 4.11, even though he could not express it in the now natural topological language. See [30] for details of the story.

The problem of tree-packing is closely related to the *arboricity* of a graph: the least number of forests that will cover its edges. For a finite graph G , another classical theorem of Nash-Williams [69] says that the edges of G can be covered by k forests if no set of ℓ vertices spans more than $k(\ell - 1)$ edges, the number of edges on an ℓ -set that k forests can at most provide. This theorem extends verbatim to infinite graphs, by compactness.

In our topological setting, however, it is natural to ask for more: that we can cover the edges of G by k *topological forests*, standard subspaces of $|G|$ that contain no circle. Interestingly, the above local sparseness condition no longer implies this; see [76] for a counterexample of Bruhn. However, if we require that our graphs are also ‘sparse at infinity’, by bounding their end degrees from above, we get the following result of Stein:

Theorem 4.12. [76] *Let $k \in \mathbb{N}$, and let G be a locally finite graph. If no set of (say) ℓ vertices of G induces more than $k(\ell - 1)$ edges and every end of G has edge-degree $< 2k$, then $|G|$ contains k topological forests covering all its edges.*

The bound of $2k$ in Theorem 4.12 is sharp [76].

Before we leave the subject of connectivity, let us briefly summarize what is known about Menger’s theorem in $|G|$. When a, b are vertices in a locally finite graph G , it is well known and easy to see that if we cannot separate a from b by a set S of fewer than k vertices then there are k independent a - b paths in G . This result does not gain from admitting a - b arcs in $|G|$ instead of just paths in G , since those arcs will still meet S (Lemma 1.3). Hence neither the minimum size of an a - b separator nor the maximum number of

a - b paths (arcs) changes. The problem becomes more interesting for standard subspaces of $|G|$, and has been solved: since standard subspaces are locally connected (Lemma 1.2), results of Whyburn [87] imply that both the point-to-point and the set-to-set version of Menger’s theorem hold in them, as long as k is finite. See [83] for what is known about Menger’s theorem in more general 1-dimensional spaces.

For infinite k , the correct version of Menger’s theorem to consider is the set-to-set version in the form suggested by Erdős (see [32]). This states that given two sets A, B of vertices there is a set of disjoint A - B paths and an A - B separator consisting of a choice of one vertex from each of these paths. This was proved for countable graphs by Aharoni [1], and for arbitrary graphs by Aharoni and Berger [2]. These results have been extended to versions where A and B are allowed to contain ends as well as vertices [18], and these have to be connected by paths, rays, or double rays (not arbitrary arcs in $|G|$). With these assumptions, Menger’s theorem holds if $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, and this condition is also necessary. If we allow arbitrary arcs in $|G|$, an example of Kühn (see [29]) shows that one even has to require that $\overline{A} \cap \overline{B} = \emptyset$. In that case, however, we once more have the situation that arcs cannot avoid separators that meet all connecting paths, rays or double rays, so again the topological version offers nothing new [29].

No versions of Menger’s theorem are known for standard subspaces of $|G|$ with k infinite.

A popular area of finite graph theory which, traditionally, has no infinite counterpart is the theory of Hamilton cycles.³¹ When we replace ‘cycle’ with ‘circle’, however, hamiltonicity problems immediately make sense. So let us call a circle in $|G|$ a *Hamilton circle* of G if it contains every vertex of G . Since circles are compact and hence closed in $|G|$, Hamilton circles also contain every end.

What does a Hamilton circle look like? The answer to this question is somewhat daunting: as soon as the graph has uncountably many ends (which is the rule rather than the exception), any Hamilton circle must be wild [16], as in Figure 1.1. Still, the notion of a Hamilton circle seems to be just the right one to generalize hamiltonicity problems to infinite graphs. Let us look at some of these.

From the extremal graph theory point of view, a particularly interesting

³¹ This is not to say that there have been no attempts. Nash-Williams and others sought to replace Hamilton cycles by spanning rays or double rays. This approach works to some extent for graphs with only one end (in which spanning double rays form circles). But it runs into difficulties as soon as the graph has more than two ends, since no ray or double ray can pass through a finite separator infinitely often. Realizing these difficulties, Halin suggested replacing Hamilton cycles with ‘end-faithful’ spanning trees: spanning trees that contain from every end exactly one ray starting at the root (which can be chosen arbitrarily). This notion has led to some interesting problems that are still open—see e.g. [37]—but not to any theory related to that of finite Hamilton cycles.

problem is how local conditions can force the (global) existence of a Hamilton cycle. Most popular among these are minimum degree conditions. When the degrees needed are large not just in absolute terms but in terms of the order of G (as in most classical results such as Dirac’s theorem [32]), such theorems are hard to generalize to infinite G .

But there are also local degree conditions that force a Hamilton cycle in a finite graph. For example, Asratian and Khachatryan [6] found a number of local Hamiltonicity conditions, all implying of Dirac’s theorem, whose the simplest version has the following infinite analogue:

Conjecture 4.13. *A connected locally finite graph G of order at least 3 has a Hamilton circle if $d(u) + d(w) \geq |N(u) \cup N(v) \cup N(w)|$ for every induced path uvw in G .*

A classical sufficient local Hamiltonicity condition that generalizes to Hamilton circles is Fleischner’s theorem: the *square* G^2 of a 2-connected finite graph G has a Hamilton cycle. (The n th power G^n of G is the graph on $V(G)$ with edges joining any pairs of vertices that have distance at most n in G .) While this finite theorem is not easy, its infinite counterpart, conjectured in [30] and proved by Georgakopoulos, is perhaps the deepest result about $|G|$ to date:

Theorem 4.14. [45] *Let G be a locally finite connected graph.*

- (i) G^3 has a Hamilton circle.
- (ii) If G is 2-connected, then G^2 has a Hamilton circle.

Thomassen [79] had previously proved (ii) for 1-ended graphs (in which a Hamilton circle is a spanning double ray).

Georgakopoulos [43] conjectured that Theorem 4.14 should extend to countable graphs that are not locally finite (see Section 5 for subtleties about $|G|$). This is interesting, because at first glance it seems impossible. For since $\Omega(G)$ will be a closed subset of any Hamilton circle of G , it must be compact for the conjecture to be true. But $\Omega(G)$ is compact (if and) only if no finite separator $S \subseteq V(G)$ splits G into infinitely many components containing rays—a property of locally finite graphs that usually fails in a countable graph. But, fortuitously, it always holds in G^2 (and in G^3): since G is connected, any component of $G^2 - S$ sends a G -edge to S , but no two components can send a G -edge to the same vertex of S , since this would create an edge of G^2 between those components. Hence $G^2 - S$ has at most $|S|$ components.

There are numerous other local density conditions that force a Hamilton cycle in a finite graph. For example, line graphs are ‘locally dense’. Thomassen [82] conjectured that every 4-connected line graph is hamiltonian, and 7-connected finite line graphs indeed are [90]. Also, the line graph of a 4-edge connected finite graph is hamiltonian [26].

Problem 4.15. *Does sufficient connectivity force the line graph of a locally finite connected graph to have a Hamilton circle?*

Georgakopoulos [43] conjectured that this should be true with the same connectivity assumptions as are currently known for finite graphs.

Another classical result about finite Hamilton cycles is Tutte's hamiltonicity theorem for 4-connected planar graphs. It was the following conjecture of Bruhn (see [30]) that first advanced the notion of a Hamilton circle:

Conjecture 4.16. *Every 4-connected locally finite planar graph has a Hamilton circle.*

Like the extension of Fleischner's theorem, this appears to be a hard problem. Partial results have been obtained by Bruhn & Yu [24] and by Cui, Wang & Yu [28].

There are also very interesting results and conjectures about Hamilton cycles in finite graphs asserting the sufficiency of conditions for their existence that are themselves global. Think of hamiltonicity problems for Cayley graphs, for sparse expanders, or for products of graphs. Some of these may be extendable to Hamilton circles, perhaps under additional assumption.

Let us close this section by remarking that most of the theorems and problems we discussed have meaningful analogues in arbitrary standard subspaces of $|G|$ rather than just $|G|$ itself. Sometimes, these extensions are easy and can be obtained by imitating the proof for $|G|$. But at other times they can be challenging. The reader is invited to explore this further.

4.2 Cycle space applications

The earliest, and so far the most successful, applications of our topological approach to locally finite graphs have been results which, for finite graphs, relate the cycle space of G to its structural properties. We have seen one such example: Theorem 2.5, which says that a standard subspace of $|G|$ can be covered by a topological Euler tour if and only if its edge set lies in $\mathcal{C}(G)$. There had been a number of earlier attempts to generalize Euler's finite theorem to infinite graphs, based on double rays as infinite analogues of finite cycles, but these attempts were hampered from the outset by the handicap that a double ray cannot visit more than $|S| + 1$ components of $G - S$ for any finite set S of vertices, and hence could not really succeed.

Another early application is MacLane's planarity criterion. Call a set $\mathcal{D} \subseteq \mathcal{C}(G)$ *sparse* if no edge of G lies in more than two elements of \mathcal{D} . MacLane's theorem says that a finite graph is planar if and only if its cycle space has a sparse generating subset. If the graph is 3-connected, these generators will necessarily be its peripheral circuits (all but at most one), the face boundaries in any drawing. (See [32], and [14] for a generalization to arbitrary surfaces.)

For our infinite G , this fails unless we allow infinite circuits: the 3-connected graph G in Figure 4.3, for example, has no finite face boundary containing the edge e , so its finite peripheral circuits do not even generate $\mathcal{C}_{\text{fin}}(G)$.

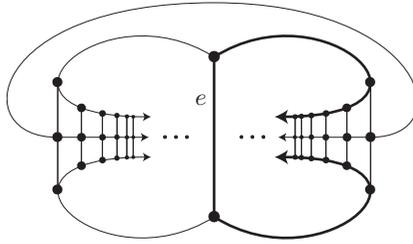


FIGURE 4.3. The edge e lies on no finite face boundary, so these do not generate the cycle space

Solving a long-standing problem of Wagner [86], Bruhn and Stein extended MacLane’s theorem to infinite graphs, using $|G|$:

Theorem 4.17. [21] G is planar if and only if $\mathcal{C}(G)$ has a sparse generating subset.

Just as MacLane turned the algebraic properties of the face boundaries of a plane graph into a planarity criterion, Archdeacon, Bonnington and Little [5] found an algebraic planarity criterion in terms of the ‘left-right-tours’ of finite plane graphs. Bruhn et al [20] extended this to locally finite graphs, based on possibly infinite ‘left-right-tours’ in $|G|$.

The planarity criterion of Kelmans and Tutte says that a finite 3-connected graph is planar if and only if its set of peripheral circuits is sparse. This follows from MacLane’s theorem and the theorem of Tutte that, in any 3-connected finite graph (planar or not), the peripheral circuits generate the cycle space.

The Kelmans-Tutte theorem, too, fails for infinite graphs unless we allow circuits to be infinite: sparseness of the finite peripheral circuits alone does not imply planarity [21]. But Bruhn [11] extended Tutte’s generating theorem to the topological cycles space (Theorem 2.3 (ii)). Hence for $\mathcal{C}(G)$ rather than $\mathcal{C}_{\text{fin}}(G)$, the infinite MacLane theorem implies the Kelmans-Tutte criterion:

Theorem 4.18. If G is 3-connected, then G is planar if and only if every edge lies in at most two peripheral circuits.

Another classical result in this context is Whitney’s duality theorem for finite graphs. It is often thought of as a planarity criterion, but can equally be viewed as a topological characterization of the graphs that have an (algebraic) dual: that these are precisely the planar graphs. As explained in Section 2.5,

duality for infinite graphs, and in particular any analogue of Whitney’s theorem, will work only in the class of graphs G satisfying condition $(*)$ from Section 2.5 (that every two vertices can be separated by finitely many edges), and for the topological cycle and cut spaces taken of the compactifications \tilde{G} obtained by adding only the undominated ends as new points (and making any other rays converge to their dominating vertex). In that setting, however, Whitney’s theorem does generalize smoothly:

Theorem 4.19. [13] *A graph satisfying $(*)$ has a dual if and only if it is planar.*

By colouring-flow duality (see [32]), the four-colour theorem can be rephrased as saying that the edge set of any finite planar bridgeless graph G is the union of two algebraic cycles: this is equivalent to the existence of a $(\mathbb{Z}_2 \times \mathbb{Z}_2)$ -flow on G , which in turn is equivalent to the 4-colourability of any dual of G . Since infinite planar graphs are 4-colourable by the Erdős – de Bruijn theorem [32], Theorem 4.19 has the following double cover corollary:

Corollary 4.20. [13] *Assume that G satisfies $(*)$, and let $\mathcal{C}(G)$ be based on the space \tilde{G} . If G is planar and bridgeless, then $E(G)$ is the union of two algebraic cycles.*

4.3 Flows in infinite graphs and networks

Most of this section deals with electrical flow in infinite networks: an exciting field with rich connections to other branches of mathematics, and one where the study of $|G|$ has a natural well-recognized place. Towards the end of the section we also discuss algebraic (group-valued) and (non-electrical) network flows. Basing these on $|G|$ rather than just G is likely to lead to an extension of most of the known finite theory to locally finite graphs, just as in Sections 4.1–2.

A finite *electrical network* is a connected finite graph (V, E) whose (undirected) edges e have real *resistances* $r(e) > 0$ assigned to them, and which has two specified vertices s and t , called the *source* and the *sink*. A *flow* in this network is a real function $f: \vec{E} \rightarrow \mathbb{R}$ on the set of orientations \vec{e} and \bar{e} of these edges e such that $f(\vec{e}) = -f(\bar{e})$ for every e and

$$\sum_v f(\vec{e}) = 0 \text{ for every vertex } v \notin \{s, t\}, \quad (\text{KH1})$$

where the sum $\sum_v f(\vec{e})$ ranges over all edges \vec{e} at v oriented away from v .

The *value* of this flow f is the number $\sum_s f(\vec{e})$. Using (KH1), one shows easily that this number is equal to $\sum_{\vec{e} \in \vec{F}} f(\vec{e})$ for every cut \vec{F} separating s from t , oriented from the side containing s to that containing t .

For every $i \in \mathbb{R}$, there is a unique flow f on G of value i that also satisfies

$$\sum_{\vec{e} \in \vec{C}} f(\vec{e})r(e) = 0 \text{ around every oriented circuit } \vec{C} \text{ in } G. \quad (\text{KH2})$$

The existence of some flow of value i is clear: we just send a flow of value i along an s - t path and let $f(\vec{e}) = 0$ elsewhere. To find a flow of value i that also satisfies (KH2), one can use the *total energy* of f : the number

$$w(f) := \sum_{e \in E(G)} f^2(e)r(e),$$

where $f^2(e) := f^2(\vec{e}) = f^2(\bar{e})$. It is easy to show that among all the flows of value i there is one of minimum total energy, and a short calculation shows that this flow satisfies (KH2). Uniqueness now comes for free, since there cannot be two flows f, f' of the same value that both satisfy (KH2). Indeed, their difference $g = f - f'$ would be a flow of value 0 such that $g(\vec{e}_0) > 0$ for some oriented edge \vec{e}_0 . Using that $\sum_v g(\vec{e}) = 0$ at every vertex v (including s and t now, since g has value zero), we ‘move greedily along’ oriented edges of G , starting at \vec{e}_0 , to find an oriented circuit \vec{C} with $g(\vec{e}) > 0$ for all $\vec{e} \in \vec{C}$. Hence $g = f - f'$ does not satisfy (KH2), but it should, since f and f' do.

We remark that condition (KH2) is equivalent to the existence of a *potential function* inducing f , a function $p: V \rightarrow \mathbb{R}$ related to f via *Ohm’s law* that $f(\vec{e})r(e) = p(v) - p(u)$ for every oriented edge $\vec{e} = uv$ of G : this is easily seen by considering a spanning tree.

For the case that $r(e) = 1$ for all e , Ohm’s law says in topological terms (cf. Section 5.3) that $f = \delta^0 p$, where δ^0 is the coboundary operator on 0-chains. Then being ‘induced by a potential function via Ohm’s law’ is the same as lying in the image of δ^0 . In our terminology, the image of δ^0 is the oriented cut space: every oriented cut $f = \lambda \vec{E}(X, Y)$ has the form $\delta^0 p$ (let p assign λ to the vertices in Y and 0 to those in X), and conversely we can write any function of the form $f = \delta^0 p$ as a sum of oriented atomic bonds $\lambda \vec{E}(v)$. The fact that it is precisely these functions $f: \vec{E} \rightarrow \mathbb{R}$ that satisfy (KH2) then is the oriented version of Theorem 2.7 (ii). Finally, such a function $f = \delta^0 p$ satisfies (KH1) if and only if p is harmonic at every vertex $v \notin \{s, t\}$: as

$$\sum_v f(\vec{e}) = \sum_v (\delta^0 p)(\vec{e}) = \left(\sum_{u \in N(v)} p(u) \right) - d(v)p(v),$$

the net flow out of a vertex v is zero if and only if $p(v)$ equals the average of the p -values at its neighbours.

So what about infinite networks? Let us specify two vertices s, t in our locally finite connected graph G , and let a function r assign positive real resistances to its edges. How do we define a ‘flow’ in G ? Traditionally, this is done as for finite networks, perhaps adding that $\sum_s f(\vec{e}) = -\sum_t f(\vec{e})$, since this no longer follows from the other assumptions.

Let us see how far this definition will carry. Given any $i \in \mathbb{R}$, there will still be a flow of value i (just along an s - t path), there will be a flow of value i with minimum total energy (defined as before) [47], and this flow will satisfy (KH2) for finite circuits. However, there may now be more than one such flow of

value i : Figure 4.4 shows a particularly striking example, which satisfies (KH2) because the graph contains no circuit at all. (Note that its second flow ‘from s to t ’ would even exist if we deleted the edge st , disconnecting s from t .)

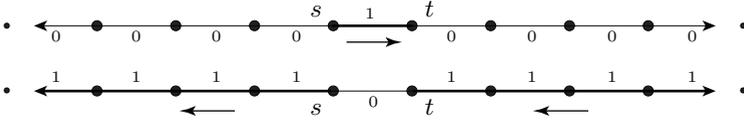


FIGURE 4.4. Two flows of value 1 satisfying (KH2)

Clearly, such ‘flows’ are of little interest, and one would like to amend the definition so as to exclude them. But it has not been clear until recently how this can be done without also excluding desirable flows. From our topological perspective, however, this is much clearer: the requirement of $\sum_v f(\vec{e}) = 0$ should be applied not only to vertices v , but also, somehow, to ends. (In the second flow of Figure 4.4, a net flow of 1 disappears into the left end, while a net flow of 1 emerges from the right end, violating the intended requirement at both ends.)

The simplest way to make such an additional requirement for ends is without mentioning ends directly: instead of (KH1) we require a condition that is equivalent to (KH1) for finite networks, but not for locally finite ones: that f should sum to zero across any *finite* oriented cut not separating s from t . Formally:

$$\sum_{\vec{e} \in \vec{F}} f(\vec{e}) = 0 \text{ for every finite oriented cut } \vec{F} \text{ of } G \text{ that does not separate } s \text{ from } t. \quad (\text{KH1}')$$

For reference, let us also restate (KH2) with explicit reference to finite circuits only:

$$\sum_{\vec{e} \in \vec{C}} f(\vec{e})r(e) = 0 \text{ around every finite oriented circuit } \vec{C} \text{ in } G. \quad (\text{KH2}')$$

A *flow* in our infinite network (G, r, s, t) , then, is a function $f: \vec{E}(G) \rightarrow \mathbb{R}$ of finite total energy that satisfies (KH1') and satisfies $f(\vec{e}) = -f(\overleftarrow{e})$ for all edges e . The *value* of this flow is the number $\sum_s f(\vec{e})$; using (KH1') one can show that this number equals the total value $\sum_{\vec{e} \in \vec{F}} f(\vec{e})$ of f across any finite oriented cut \vec{F} that does separate s from t , oriented from its side containing s to that containing t .

With this definition of a flow, and assuming in addition that the *total resistance* $\sum_{e \in E(G)} r(e)$ of the network is finite, Georgakopoulos [47] proved that flows of a given value and satisfying (KH2') are indeed unique. This is interesting, since (KH2') makes demands only on finite circuits, while the difference between two flows of the same value might have an infinite circuit as its

carrier.³² In particular, we cannot prove uniqueness directly from (KH2'), as for finite graphs, but have to use a limit construction as discussed in Section 3.

Theorem 4.21. [47] *Let G be a locally finite network, and let $i \in \mathbb{R}$. If G has finite total resistance, there is a unique flow in G of value i that satisfies (KH2'). These flows also satisfy (KH2) for infinite circuits \vec{C} .³³*

Theorem 4.21 is best possible in that uniqueness fails if we do not require finite total resistance. And the assumption of finite total resistance is not, in fact, unnatural. Indeed, if we interpret the numbers $r(e)$ as edge lengths (as in Figure 4.5), then r makes G into a metric space. Let \overline{G} be its completion. Georgakopoulos [46] showed that if $\sum r(e) < \infty$ then \overline{G} is homeomorphic to $|G|$, by a homeomorphism that is the identity on G (Theorem 5.4). We shall look at metric completions of G in more detail in Section 5.

Consider any locally finite connected graph G , a function $r: E(G) \rightarrow \mathbb{R}^+$ such that $\sum_{e \in E(G)} r(e) < \infty$, and a function $p: V(G) \rightarrow \mathbb{R}$ such that $f := \delta^0 p/r$ satisfies

$$\sum_{\vec{e} \in \vec{F}} f(\vec{e}) = 0 \text{ for every finite oriented cut } \vec{F} \text{ of } G. \quad (\text{KH1}'')$$

Then for any choice of $s, t \in V(G)$, our function f is a flow in the network (G, r, s, t) of value 0 that satisfies (KH2'). Since the constant function $\vec{E} \rightarrow \{0\}$ is also such a flow, Theorem 4.21 implies that our f is in fact constant with value 0, and hence p too must have been constant.

However if we relax (KH1'') to (KH1), there may be more such functions f and p . And since we did not specify a source or sink among the vertices, these functions are more natural than the example of Figure 4.4: f may be viewed as a flow with (possibly many) sources and sinks at infinity, as in Figure 4.5.

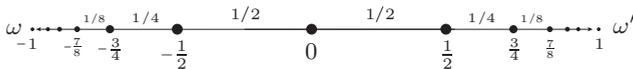


FIGURE 4.5. A flow of value 1 from ω to ω' , with resistances (top) and potentials (bottom)

To rule out other pathological examples, however, we should require now that p extends continuously to the ends:

Problem 4.22. *Let G , r and p be given as above, with $f = \delta^0 p/r$ satisfying (KH1) but not necessarily (KH1''). Does p extend continuously to $\Omega(G)$?*

Some results in this direction can be found in the last section of Georgakopoulos [47].

³² By the oriented version of Theorem 2.7 (i), it will lie in the oriented cycle space $\vec{C}(G)$.

³³ In particular, the infinite sum $\sum_{\vec{e} \in \vec{C}} f(\vec{e}) r(e)$ is well defined and finite.

Problem 4.23. Which continuous functions $p: \Omega(G) \rightarrow \mathbb{R}$ extend continuously to $V(G)$ so that $\delta^0 p/r$ satisfies (KH1)?

The Dirichlet problem has been widely studied; see, e.g., [7, 88, 89].

Let us close this section with a glance at how non-electrical network and algebraic flow theory extend to locally finite graphs. As before, we wish to exclude ‘flows’ that issue from or dissipate to infinity, and will adjust our definitions accordingly. Let G be a locally finite connected graph. Call a function f on $\vec{E}(G)$ *symmetrical* if $f(\vec{e}) = -f(\bar{e})$ for all oriented edges $\vec{e} \in \vec{E}(G)$.

Specifying two vertices s, t as source and sink turns G into a (non-electrical) network. A *flow* through this network is a symmetrical function on $\vec{E}(G)$ satisfying (KH1’).

Problem 4.24. Extend the non-algorithmic aspects of network flow theory to locally finite graphs, based on the above definition of a network flow.

The max-flow min-cut theorem has been extended in this way; see the last few pages of [3].

Turning now to algebraic flows, let H be any abelian group. Call a symmetrical function $f: \vec{E}(G) \rightarrow H$ an *H -flow* on G if it satisfies (KH1’’) and $f(\vec{e}) \neq 0$ for all $\vec{e} \in \vec{E}$. Such a function induces H -flows also on the finite minors G_n of G defined in Section 3.1. It seems that, using compactness as explained in Section 3.2. (and topological spanning trees wherever spanning trees are needed), one can extend most—but not all—the standard results about H -flows in finite graphs to locally finite graphs:

Problem 4.25. Extend the algebraic flow theory of finite graphs to locally finite graphs, based on the above definition of an H -flow.

5. Outlook

The topological approach to studying locally finite graph discussed in this paper can be taken further, in various directions. In this section we address some of these:

- graphs that are not locally finite;
- locally finite graphs with finer ‘boundaries’ than ends, such as boundaries of (Cayley graphs of) hyperbolic groups in the sense of Gromov;
- homology of non-compact spaces other than graphs;
- infinite matroids.

5.1 Graphs with infinite degrees

Consider the two graphs shown in Figure 5.1, the *fan* and the *double fan*. In each of these graphs there is a ray R that is dominated by a vertex (respectively, by two vertices).³⁴

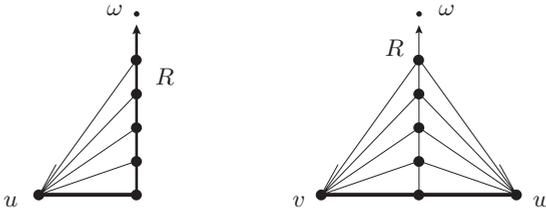


FIGURE 5.1. Should the fat edge sets be circuits?

In both these graphs, the sum (mod-2) of all the triangle circuits equals the set of heavy edges. These edges do not look like circuits: they are the edge sets of a ray, or of a path of length 2. However, as thin sums of finite circuits, these edge sets ought to be elements of the cycles space $\mathcal{C}(G)$: remember that we need to allow infinite thin sums in \mathcal{C} in order to make important applications of \mathcal{C} work, such as Mac Lane’s theorem. But if these sets of edges lie in \mathcal{C} , they ought to be disjoint unions of circuits (if we want Theorem 2.3 (i) to extend to these graphs). Hence we feel compelled to admit those edge sets as circuits—or to abandon much of the cycle space theory we just established for locally finite graphs.

So can we make sense of the notion that those edge sets should be circuits? Indeed it seems we can; and the answer has to do with what topology we choose for $|G|$ when G is not locally finite.

Let us be more formal. Let G be any connected graph. The ends of G are defined as earlier, as equivalence classes of rays. As the topology on G itself we take slightly fewer basic open sets than the identification topology for a 1-complex does:³⁵ around every vertex v we take as basic open sets only the open stars $E_\epsilon(v)$ of length ϵ for arbitrary but fixed $\epsilon < 0$, the same ϵ for every edge at v . (Then the sets $E_{1/n}(v)$ form a countable neighbourhood basis of v . In the 1-complex topology, vertices of infinite degree have no countable neighbourhood basis.) As basic open sets around ends ω we take for every finite set S of vertices and every ϵ with $0 < \epsilon \leq 1$ the open ϵ -collar $\hat{C}_\epsilon(S, \omega)$ of $\overline{C(S, \omega)}$.³⁶ Note that, as for vertices, these ϵ -collars have uniform width.

³⁴ See the end of Section 2 for the formal definition of ‘dominate’.

³⁵ That topology leads to a topology for $|G|$ called TOP. Historically, this was the first topology for $|G|$ to be considered, but it has few advantages over the topologies MTOP and VTOP discussed here. In particular, $|G|$ under TOP is neither compact nor metrizable as soon as a vertex has infinite degree. See [31, 34] for more.

³⁶ More formally, let $\hat{C}_\epsilon(S, \omega)$ be obtained from the set $\hat{C}(S, \omega)$ defined in Section 1 by replacing each open C - S edge $\hat{e} = (c, s)$ with its initial open segment of length ϵ , assuming

Let us call this topology MTOP.³⁷ Under this topology, $|G|$ is clearly Hausdorff, in fact normal [75]. However, unless G is locally finite, $|G|$ will no longer be compact. (Indeed, consider an open cover that consists the middle half of e for every edge e at some fixed vertex v , and one further open set to cover the rest of $|G|$ but none of the mid-points of edges at v . If v has infinite degree, this cover has no finite subcover.)

In order to give $|G|$ a chance to be compact, we have to take a coarser topology still: we take the same open stars $E_\epsilon(v)$ around vertices, but around ends ω we take only the open sets $\hat{C}_\epsilon(S, \omega)$ with $\epsilon = 1$, our original sets $\hat{C}(S, \omega)$. This topology is called VTOP.

Theorem 5.1. [31] *For every graph $G = (V, E)$, not necessarily connected or locally finite, the following statements are equivalent under VTOP:*

- (i) $|G|$ is compact.
- (ii) For every finite $S \subseteq V$, the graph $G - S$ has only finitely many components.
- (iii) Every closed set of vertices is finite.

For G countable and connected, Theorem 5.1 implies that $|G|$ is compact under VTOP if and only if G has a locally finite spanning tree, which can be chosen normal [31]. Note that since VTOP agrees with TOP and MTOP on its closed subspace $\Omega(G)$, condition (ii) of Theorem 5.1 characterizes, for any of these topologies, the graphs whose end space Ω is compact.

We complete our excursion to the various topologies of $|G|$ by mentioning one more, known as ETOP. Strictly speaking, this is not a topology on $|G|$,³⁸ but on a similar space obtained by adding to G its *edge-ends* rather than its ends: the equivalence classes of rays under which two rays are equivalent if no finite set of *edges* separates them. The basic open sets of this space are the components of $G - X$ left by deleting a finite set X of inner points of edges (and an edge-end belongs to the component in which all its rays have a tail).³⁹ If G is connected, as we assume, then this space is compact [74].

On the set of vertices and edge-ends, being separated by such a finite set X (i.e., by a finite set of edges) is an equivalence relation. The open sets in ETOP cannot distinguish such equivalent points. But identifying equivalent points (including different vertices, such as the vertices v and w in Figure 5.1) yields a compact Hausdorff space, which is metrizable if G was countable. See [50, 74, 46] for more on edge-ends and ETOP.

that e itself has length 1.

³⁷ The ‘M’ comes from fact that MTOP makes $|G|$ metrizable, at least for countable connected G (see below).

³⁸ ... unless G is locally finite, in which case its edge-ends coincide with its usual ends.

³⁹ Such a ‘component’ is a graph-theoretical component left by deleting from G the edges containing points from X , together with the segments of these edges (after deleting X) that have an endpoint in such a component.

5.2 The identification topology

Let us turn back to our two graphs from Figure 5.1. Under MTOP, the sets of fat edges are still far from being circuits. Under VTOP and ETOP, however, they almost are: these topologies cannot distinguish an end from vertices dominating it,⁴⁰ and if we identified the end with their dominating vertices, the fat edge sets would indeed become circuits. In the single fan on the left, this does indeed look right: why introduce, when forming $|G|$ from G , a new limit point for the ray R if a natural limit point, the vertex u , already exists? In the double fan on the right, however, one may have more qualms: identifying both v and w with the end ω would result in the identification of two vertices, changing the graph.

So shall we identify ends with their dominating vertices, or shall we not? Before we describe a radical and intriguing answer to this question due to Georgakopoulos [46], let us deal with the case when the dilemma does not arise: the case that no end is dominated by more than one vertex. This is the case, for example, for the graphs satisfying condition (*) of Section 2: a natural class of graphs that occurs quite independently of this problem. For such graphs G , let \tilde{G} denote the identification space obtained from $|G|$ by identifying each vertex with all the ends it dominates. This is a Hausdorff space—even under VTOP when $|G|$ fails to be Hausdorff [36]. Other features of the topology of \tilde{G} , such as compactness, depend on the topology used for $|G|$. If VTOP was used on $|G|$, the identification topology on \tilde{G} is known as ITOP.

When G is 2-connected, then (*) implies that G is countable, and that no finite set of vertices separates G into infinitely many components. Then $|G|$ is compact under VTOP (Theorem 5.1), and hence so is \tilde{G} under ITOP. Answering a question of Sprüssel [75], Richter and Vella [73] extended this by showing that \tilde{G} is in fact ‘Peano’:

Theorem 5.2. [73] *Let G be a 2-connected graph in which every two vertices can be separated by finitely many edges.⁴¹ Then \tilde{G} under ITOP is a compact, connected, locally connected, metric space.*

For more on \tilde{G} , see [36, 73, 83].

5.2 Compactification versus metric completion

As mentioned already in Theorem 1.1, the compactification $|G|$ of a connected and locally finite graph G is metrizable. When G is not locally finite, $|G|$ with our usual topology, VTOP, fails to be even Hausdorff (as soon as there is a

⁴⁰ More generally, it is easy to see that a space $|G|$ is Hausdorff under VTOP if and only if no end of G is dominated.

⁴¹ This assumption can be weakened: it is only necessary to assume that no two vertices can be linked by infinitely many independent paths [73]. Under this weaker assumption, however, \tilde{G} can contain arcs and circles consisting entirely of vertices or ends. Such spaces were studied in [73].

dominated end), and no vertex of infinite degree has a countable neighbourhood basis in the usual 1-complex topology. However, if we consider $|G|$ with MTOP, we have the following result:

Theorem 5.3. [31] *For a connected graph G , and $|G|$ endowed with MTOP, the space $|G|$ is metrizable if and only if G has a normal spanning tree.*

Unless otherwise mentioned, the space $|G|$ will in this section always carry the topology MTOP.

Theorem 5.3 is interesting in both directions: as a characterization of the graphs G for which $|G|$ is metrizable, but also as an unexpected topological characterization of the graphs admitting a normal spanning tree. These are known to include all countable graphs, and they have been characterized by two types of forbidden substructure [37]. But these substructures themselves are not fully understood, and classifying them more accurately remains a challenging open problem.

The metric which a normal spanning tree T of G induces on $|G|$ is easy to describe. We first assign length $\ell(e) := 2^{-n}$ to every tree-edge e from level $n - 1$ to level n , the root being at level 0. This defines a metric on $V(G) \cup \Omega(G)$, via

$$d_\ell(x, y) := \sum_{e \in xTy} \ell(e),$$

where xTy is the path, ray or double ray in \bar{T} from x to y . We then extend this metric to one on $|G|$ by mapping every edge xy homeomorphically to a real interval of length $d_\ell(x, y)$. As one can check [31], this metric induces the topology MTOP on $|G|$.

Instead of defining distances between ends and other points explicitly, we could also have defined them only for points of G itself (as above), and then taken the completion of this metric space (G, d_ℓ) . Since two Cauchy sequences of points in (G, d_ℓ) are equivalent if and only if in $|G|$ they converge to the same end, this would have yielded the same result: the complete metric space on $G \cup \Omega(G)$ we defined explicitly above, and whose metric induces MTOP on $|G|$.

Georgakopoulos [46] showed that this is not just a feature of our particular metric. Given any function $\ell: E(G) \rightarrow \mathbb{R}^+$ assigning positive ‘edge lengths’, defining distances between vertices u, v via

$$d_\ell(u, v) = \inf \sum_{e \in P} \ell(e) \text{ over all } u\text{-}v \text{ paths } P \text{ in } G$$

and identifying any u, v with $d_\ell(u, v) = 0$ defines a metric on the resulting quotient space of $V(G)$, which extends to a metric d_ℓ on the corresponding quotient of the entire graph G once we fix homeomorphisms between the edges

and real intervals of their respective lengths. Let us denote this metric space as (G, d_ℓ) , and the topology it induces on its completion as ℓ -TOP. The metric subspace induced by the completion points we added is the ℓ -TOP *boundary* of G .

Theorem 5.4. [46] *If G is any countable connected graph and $\ell: E(G) \rightarrow \mathbb{R}^+$ satisfies $\sum_{e \in G} \ell(e) < \infty$, then completing the metric space (G, d_ℓ) yields the edge-ends of G as completion points and ETOP as the induced topology of the completion. If G is locally finite, this coincides with $|G|$.*

In view of our earlier normal spanning tree example one should expect that, for G locally finite, the finiteness condition in Theorem 5.4 can be relaxed considerably without losing that the metric completion of G coincides with $|G|$:

Problem 5.5. *Given a countable connected graph G , characterize the functions $\ell: E(G) \rightarrow \mathbb{R}^+$ for which the completion of (G, d_ℓ) coincides with $|G|$.*

For functions ℓ not satisfying $\sum_{e \in G} \ell(e) < \infty$, the metric-completion approach opens up a wide range of new possibilities over the purely topological compactification approach, even for locally finite graphs. For example, the *hyperbolic compactification* of a locally finite hyperbolic graph G , as introduced by Gromov [49], is sometimes defined in a somewhat roundabout way: in an intermediate step one endows G with with a different (bounded) *Gromov metric* (which on G itself is somewhat odd⁴²), whose sole purpose is that its completion has a compact boundary that can then be combined with G in its *original* topology to yield an interesting compact space.⁴³ But Gromov [49] also proved that the hyperbolic compactification of G can, alternatively, be obtained as the completion of (G, d_ℓ) for a natural function ℓ of edge lengths (see [27] for a detailed proof, or [46] for an indication), a result that makes its use—e.g. for the study of hyperbolic groups—more accessible to a graph-theoretical approach. Figure 5.2 shows a hyperbolic graph G with an assignment ℓ of edge lengths for which the ℓ -TOP boundary of (G, d_ℓ) is its hyperbolic boundary.

The flexibility of the metric-completion approach to defining a boundary of an infinite graph G allows us to keep our options open regarding the opening question of Section 5.1, the problem of how to define circuits—and the resulting homology—in a graph with infinite degrees. Depending on the graphs under investigation, we can choose our edge length function ℓ in whatever way seems best to describe those graphs: we can *choose* whether or not we want two rays from the same end to converge to a common point at infinity or not, or not even to converge at all.

⁴² For example, it induces the discrete topology on G , including edges!

⁴³ For example, this *hyperbolic boundary* refines the ‘end boundary’ $\Omega(G)$ of G induced by $|G|$, and unlike $\Omega(G)$ it can have non-trivial connected components.

there may be circles in this completion in which the edges they contain are not dense; we shall therefore not be able to represent homology classes by sets of edges, as we did in the case of $\mathcal{C}(G)$.

Problem 5.6. *Extend the topological cycle space theory of locally finite graphs to a homology theory for metric spaces that can arise as completions of spaces of the form (G, d_ℓ) .*

As a step towards this goal, it will help to recast the topological cycle space theory of Sections 1 and 2, even for locally finite graphs, as a homology theory in the usual terms of algebraic topology [40]. We shall look at this problem next. Some first steps towards Problem 5.6 itself are already taken in [46].

5.3 Homology

In this section, any undefined graph G will again be an arbitrary connected and locally finite graph.

When the topological cycle space was first introduced [30, 34], the motivation was to extend the standard notion of the cycle space of a finite graph in such a way that it could play a similar role for infinite graphs. In particular, the various theorems relating $\mathcal{C}(G)$ to other structural properties of G , such as planarity, should extend to the new notion. The topological cycle space $\mathcal{C}(G)$, for locally finite G , has achieved this goal superbly—see Section 4—but at a price: due to the ad-hoc manner of its definition, this new $\mathcal{C}(G)$, unlike $\mathcal{C}_{\text{fin}}(G)$, is no longer obviously a special case of the first homology of either G or $|G|$ in any sense that would be standard in algebraic topology.

On the other hand, it seems desirable to recast our notion of $\mathcal{C}(G)$ in terms of standard homology. This is for two reasons: our understanding of $\mathcal{C}(G)$ might benefit from the vast body of knowledge available for standard homology theories; and conversely, the ideas that went into the notion of $\mathcal{C}(G)$ might throw a new light on the homology of more general spaces once we can rephrase $\mathcal{C}(G)$ in such a way that it makes sense when G is not a graph.

For a finite or infinite graph G , its first simplicial homology group is the same as, in our notation, $\mathcal{C}_{\text{fin}}(G)$. The elements of this group or vector space are finite sets of edges.⁴⁵ Hence as soon as G contains an infinite circuit, $\mathcal{C}(G)$ differs from $\mathcal{C}_{\text{fin}}(G)$. But $\mathcal{C}(G)$ also differs from the simplicial homology we would obtain from just allowing infinite (thin) sums of edges as 1-chains: in this homology, the edges of a double ray would form an algebraic cycle, since every vertex lies on two such edges and hence the boundary of this edge set is zero, but this is not a cycle in our topological cycle space. Thus, $\mathcal{C}(G)$ cannot

⁴⁵ We still take all our coefficients from \mathbb{Z}_2 , for easier compatibility with our earlier treatment of $\mathcal{C}(G)$. However all the results we describe hold with integer coefficients too, which are indeed the setting in the papers we refer to.

be described by an extension of simplicial homology obtained by just allowing infinite chains.⁴⁶

Another way to extend the simplicial homology of finite complexes to infinite ones is to take limits. In our situation, we could use the techniques from Section 3 to do this. Indeed, recall that for the finite minors $G_0 \preceq G_1 \preceq \dots$ of G defined in Section 3.1 the algebraic cycles of G_m induce algebraic cycles in G_n , for all $n < m$. The first simplicial homology groups of the G_n thus form an inverse system of homomorphisms, and we could define the first homology group of G as their inverse limit.⁴⁷

Theorem 5.7. *The maps $D \mapsto D \cap E(G_n)$ define a group isomorphism from $\mathcal{C}(G)$ to $\varprojlim (\mathcal{C}_{\text{fin}}(G_n))_{n \in \mathbb{N}}$.*

While Theorem 5.7 points out an interesting aspect of $\mathcal{C}(G)$, it does not capture all its relevant aspects. For even if we assume that representing $\mathcal{C}(G)$ as $\varprojlim \mathcal{C}(G_n)$ captures all its relevant features as a group, the (abstract) homology group of a graph does not capture all the relevant aspects of its homology.⁴⁸ For example, recall that if G is 3-connected, then its peripheral circuits generate $\mathcal{C}(G)$ (Theorem 2.3 (ii)), and G is planar if and only if these circuits form a sparse set (Theorem 4.18). In order to prove this for infinite G , we thus need to know when a given edge set $D \in \mathcal{C}(G)$ is a peripheral circuit. If all we know of D are its intersections with the edge sets of the finite minors G_n (as we would if we viewed D as an element of $\varprojlim (\mathcal{C}_{\text{fin}}(G_n))_{n \in \mathbb{N}}$), this would be hard to decide; indeed it would not be obvious from this information whether D is a circuit at all (see Section 3.2).

A more subtle approach, which has been pursued in [38, 40], 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 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

⁴⁶ By contrast, $\mathcal{B}(G)$, the set of all cuts, remains the image of the coboundary operator δ^0 if we allow arbitrary infinite 0-cochains.

⁴⁷ This is tantamount to taking the Čech homology of $|G|$; see [38].

⁴⁸ This is reflected by the fact that, as abstract groups, $\mathcal{C}(G)$ hardly depends on G : by Theorem 2.1, it is the direct product of as many copies of the coefficient ring as a topological spanning tree of G has chords—i.e., of \aleph_0 copies for most G .

as σ^+ , then $[\sigma^+ + \sigma^-] = 0 \in H_1$.⁴⁹ 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 .

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 $H_1(|G|) \rightarrow \mathcal{E}(G)$,⁵⁰ 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 Section 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 more than finitely 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 $v = v_i$. Using the methods from Section 3.3 it is not hard to show that σ is continuous [38], and clearly it traverses every edge of G the desired number of times.

Equally surprisingly, perhaps, f is not normally 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 5.8. [38] *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.*

⁴⁹ To 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 .

⁵⁰ For 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 $H_1(S^1; \mathbb{Z}_2) \rightarrow \mathbb{Z}_2$. Given $h \in H_1(|G|)$, let $f(h) := \{ e \mid (\pi \circ (f_e)_*)(h) = 1 \in \mathbb{Z}_2 \}$. See [38] for details.

An example of a non-null-homologous loop in $|G|$ whose homology class maps to the empty 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 L 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 L 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 twice, once in each direction. Clearly, f maps (the homology class of) this loop ρ to $0 \in \mathcal{C}(G)$.

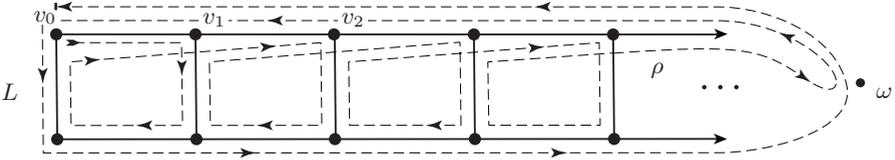


FIGURE 5.3. The loop ρ is not null-homologous

The loop ρ is indeed not null-homologous [38], 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 , two for every 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 our infinite ladder L , 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 $\Lambda(\partial\tau) = \Lambda(\sigma_1\sigma_2) - \Lambda(\sigma) = 0$. We did not

quite find such an invariant Λ , but a collection of similar invariants which, together, can distinguish loops like ρ from boundaries.

In order to find such functions on 1-chains that are invariant under homotopies, it was necessary to find a combinatorial description for the homotopy types of loops—that is, for the fundamental group of $|G|$. Such a combinatorial characterization is given in [39], where $\pi_1(|G|)$ is characterized as a group of (infinite) words of chords of a topological spanning tree—in the spirit of the usual description of $\pi_1(G)$ for a finite graph as the free group generated by such chords—and as a subgroup of the inverse limit of these finitely generated free groups. (The group $\pi_1(|G|)$ itself is not free, unless G is essentially finite.)

Not surprisingly, our topological cycle space $\mathcal{C}(G)$ can, as a group, be viewed as the *infinite abelianization* of $\pi_1(|G|)$: the factor group of $\pi_1(|G|)$ obtained by declaring two (reduced) words as equivalent if each letter occurs in both the same number of times (regardless of position).

Let us return to our original goal: to see to what extent $\mathcal{C}(G)$ can be captured by the singular homology of $|G|$. In view of Theorem 5.8, the goal might be phrased as follows:

Problem 5.9. *Devise a singular-type homology theory for locally compact spaces with ends that coincides with $\mathcal{C}(G)$ when applied to $|G|$ in dimension 1.*

Some first steps in this direction were already taken in [38]. The approach there was to allow only singular simplices whose vertices lie in G (i.e., are not ends); to allow infinite chains that are locally finite at points of G —i.e., every point in G has a neighbourhood meeting only finitely many simplices in the given chain—but not necessarily at ends (reflecting the fact that, for example, the union of all the vertical double rays in the grid defines an algebraic cycle of which infinitely many 1-simplices contain the end); and to restrict the set of cycles to those chains in the kernel of the boundary operator that could be written as a (possibly infinite) sum of finite cycles.

This approach solves Problem 5.9 in an ad-hoc sort of way: it permits the definition of homology groups for arbitrary locally compact spaces with ends, it defaults to the topological cycle space for graphs in dimension 1—but it is not a homology theory in the sense of the usual axioms [56, 41]. To achieve the latter, one has to find a way of implementing the required restrictions as conditions on chains rather than on cycles. This was done in [40]. However, this is no more than a beginning, and more translation work remains to be done—for example, of the duality theory indicated for $\mathcal{C}(G)$ in Section 2.

5.4 Infinite matroids

Traditionally, infinite matroids are defined like finite ones, with the additional axiom that an infinite set is independent as soon as all its finite subsets are independent. This reflects the notion of linear independence in vector spaces, and also the absence of the usual (finite) cycles in a graph: the bases of the cycle

matroid of an infinite graph are then the edge sets of its (ordinary) spanning trees. We shall call such matroids *finitary*. Note that the circuits in a finitary matroid, the minimal dependent sets, are necessarily finite.

An important and regrettable feature of such finitary matroids is that the additional axiom restricting the infinite independent sets spoils duality, one of the key features of matroid theory. For example, any bond of a graph ought to be a circuit in its dual matroid: a set of edges that is minimal with the property of not lying in the complement of a spanning tree, or equivalently, with the property of containing an edge from every spanning tree. Since finitary matroids have no infinite circuits, the circuit matroid of a graph with an infinite bound thus cannot have a finitary dual.

Our theory, however, suggests an obvious solution to this problem: shouldn't infinite matroids be defined in such a way that infinite circuits in a graph can become matroid circuits, and topological spanning trees become bases? Indeed, infinite circuits are not contained in the edge set of any topological spanning tree (although they are contained in the edge set of an ordinary spanning tree), while if we delete any edge from an infinite circuit, then its remaining edges can be extended to a topological spanning tree by Lemma 1.5.

There are two main challenges in devising axioms for such a non-finitary theory of infinite matroids: to avoid the mention of cardinalities, and to take care of limits.⁵¹ In [15], such a theory has been proposed in terms of independence axioms, and equivalent sets of basis and circuit axioms. There is also an equivalent set of closure axioms, which was proposed by Higgs [57, 58, 59] long before [15] but had gone largely unnoticed (perhaps due to a lack of examples), and a combination of independence and exchange axioms developed from these by Oxley [70, 71].⁵²

With these axioms, duality works as expected from finite matroids: it is exemplified by dual planar graphs (see below), there is a well-defined notion of minors with contraction and deletion as dual operations, and so on.

If one takes the view that these matroids deserve further study, there are two things to do: to complete the collection of axiom systems, and—perhaps with their help—find more natural examples. An aspect of finite matroids that is particularly resistant to a reformulation without cardinalities is rank:

Problem 5.10. *Characterize the infinite matroids axiomatized in [15] in terms of suitably defined ‘qualitative’ rank functions.*

⁵¹ If one wants to have bases and circuits, one has to ensure that maximal independent sets or minimal dependent sets exist: with infinite sets, this is no longer clear.

⁵² These matroids were known as *B-matroids*. Perhaps the fact that such very different sets of axioms describing what are fundamentally the same objects were developed independently speaks for the notion underlying both: it does appear that the infinite matroids defined in [15], or equivalently the B-matroids studied by Higgs and Oxley, are the ‘right’ notion of an infinite matroid that deserves further study.

So what non-finitary matroids are there? As had been our motivation, the edge sets of the topological spanning trees of a locally finite graph G do indeed form the bases of an infinite matroid in this theory, as do those of the ordinary spanning trees. Let us denote these matroids by $M_{\text{TST}}(G)$ and $M_{\text{ST}}(G)$, respectively. Similarly, the sets of edges of G that do not contain a bond of G form the independent sets of a matroid $M_{\text{B}}(G)$, and the sets of edges that do not contain a *finite* bond of G form the independent sets of a matroid $M_{\text{FB}}(G)$. Clearly, $M_{\text{ST}}(G)$ and $M_{\text{B}}(G)$ form a pair of dual matroids, and by Lemmas 1.2 and 1.4 so do $M_{\text{TST}}(G)$ and $M_{\text{FB}}(G)$.

The same is true for the slightly more general identification spaces \tilde{G} under ITOP, defined in Section 5.2. Thus if G and G^* are a pair of dual graphs satisfying condition (*) from Section 2.5, and we consider circuits and topological spanning trees in the spaces \tilde{G} and \tilde{G}^* , then $M_{\text{ST}}(G)$ and $M_{\text{TST}}(G^*)$ form a dual pair of matroids [15].

Much of the attractiveness of finite matroids stems from the fact that they provide a unified framework for some essential common aspects of otherwise disparate branches of mathematics. Whether or not the same can be said for infinite matroids, axiomatized in this way, will depend on concrete examples that have yet to be found—if possible, from as different areas of mathematics as possible.

There is no doubt that non-finitary matroids are plentiful: as proved in [15], the dual of any finitary connected matroid is non-finitary. However, perhaps there are natural ‘primary’ matroids that are non-finitary and have therefore gone unnoticed?

Problem 5.10.

- (i) *Find more natural examples of non-finitary matroids.*
- (ii) *Find some natural examples of non-finitary matroids whose duals are also non-finitary.*

Theorem 2.16 seems to suggest that, given a locally finite graph G (or one that satisfies (*) from Section 2.5), there might be a matroid on the set $E(G) \cup \Omega(G)$ in which the sets $F \cup \Psi$ with $F \cup E(G)$ and $\Psi \subseteq \Omega(G)$ for which $\bigcup F \cup \Psi$ contains no circle form the independent sets. However, as soon as we try to apply basis or circuit exchange axioms to ends, we see that this fails. For example, if two circles meet in exactly one end, we cannot delete the end and find another circle in the rest (the union of all the edges of the two circles). Similarly, in the double ladder we could choose as a basis B_1 the union of one double ray, all the rungs, and both ends, and as another basis B_2 all the edges that are not rungs and one of the two ends. If we delete the other end, ω say, from B_1 , we cannot find an element of $B_2 \setminus B_1$ (which would be an edge) that we could add to $B_1 - \omega$ to form another basis.

More generally, the duality which Theorem 2.16 expresses for graphs with ends cannot be expressed in terms of matroid duality. Indeed, suppose that,

with the notation of Theorem 2.16, there is a matroid M on $E \cup \Omega$ whose bases are the sets $F \cup \Psi$ (with $F \subseteq E$ and $\Psi \subseteq \Omega$) that form the subspaces X of $|G|$ for which \tilde{X} is a spanning tree of \tilde{G} . Then $M/\Omega = M_{\text{TST}}(G)$ and $M \setminus \Omega = M_{\text{ST}}(G)$. But one can show that, given any 2-connected graph G satisfying condition (*) of Section 2.5, there is no matroid M on any set $E(G) \cup X$ with $X \cap E(G) = \emptyset$ such that $M/X = M_{\text{TST}}(G)$ and $M \setminus X = M_{\text{ST}}(G)$. See [16] for details.

Richter et al. used a non-finitary matroid analogous to $M_{\text{TST}}(G)$ for their proof of an extension of Theorem 4.19 (Whitney) to ‘graph-like spaces’; see [72] in this volume.

References

- [1] R. Aharoni, Menger’s theorem for countable graphs, *J. Comb. Theory B* **43** (1987), 303–313.
- [2] R. Aharoni & E. Berger, Menger’s theorem for infinite graphs, *Invent. math.* (to appear).
- [3] R. Aharoni, E. Berger, A. Georgakopoulos, A. Perlstein & P. Sprssel, The max-flow min-cut theorem for countable networks, *J. Comb. Theory B* (to appear).
- [4] R. Aharoni & C. Thomassen, Infinite highly connected digraphs with no two arc-disjoint spanning trees, *J. Graph Theory* **13** (1989)
- [5] D. Archdeacon, P. Bonnington, and C. Little, An algebraic characterization of planar graphs, *J. Graph Theory* **19** (1995), 237–250.
- [6] A.S. Asratian & N.K. Khachatryan, Some localization theorems on hamiltonian circuits, *J. Comb. Theory B* **49** (1990), 287–294.
- [7] I. Benjamini & O. Schramm, Harmonic functions on planar and almost planar graphs and manifolds, via circle packings, *Invent. math.* **126** (1996), 565–587.
- [8] E. Berger & H. Bruhn, Eulerian edge sets in locally finite graphs, preprint 2008.
- [9] N. Biggs, Algebraic potential theory on graphs, *Bull. Lond. Math. Soc.* **29** (1997), 641–682.
- [10] B. Bollobás, *Extremal Graph Theory*, Academic Press, London 1978.
- [11] H. Bruhn, The cycle space of a 3-connected locally finite graph is generated by its finite and infinite peripheral circuits, *J. Comb. Theory B* **92** (2004), 235–256.
- [12] H. Bruhn, personal communication 2009.
- [13] H. Bruhn and R. Diestel, Duality in infinite graphs, *Comb. Probab. Comput.* **15** (2006), 75–90.
- [14] H. Bruhn and R. Diestel, MacLane’s theorem for arbitrary surfaces, *J. Comb. Theory B* **99** (2009), 275–286.
- [15] H. Bruhn, R. Diestel, M. Kriesell & P. Wollan, A new look at infinite matroids, preprint 2009.
- [16] H. Bruhn, R. Diestel, and J. Pott, Dual trees can share their ends, preprint 2009.
- [17] H. Bruhn, R. Diestel & M. Stein, Cycle-cocycle partitions and faithful cycle covers for locally finite graphs, *J. Graph Theory* **50** (2005), 150–161.
- [18] H. Bruhn, R. Diestel & M. Stein, Menger’s theorem for infinite graphs with ends, *J. Graph Theory* **50** (2005), 199–211.
- [19] H. Bruhn & A. Georgakopoulos, Bases and closures under infinite sums, preprint 2008.
- [20] H. Bruhn, S. Kosuch & M. Win Myint, Bicycles and left-right tours in locally finite graphs, *Eur. J. Comb.* **30** (2009), 356–371.

- [21] H. Bruhn and M. Stein, MacLane’s planarity criterion for locally finite graphs, *J. Comb. Theory B* **96** (2006), 225–239.
- [22] H. Bruhn and M. Stein, On end degrees and infinite circuits in locally finite graphs, *Combinatorica* **27** (2007), 269–291.
- [23] H. Bruhn and M. Stein, Duality of ends, *Comb. Probab. Comput.* (to appear).
- [24] H. Bruhn and X. Yu, Hamilton circles in planar locally finite graphs, *SIAM J. Discrete Math.* **22** (2008), 1381–1392.
- [25] K. Casteels and B. Richter, The bond and cycle spaces of an infinite graph, *J. Graph Theory* **59** (2008), 126–176.
- [26] P.A. Catlin, Supereulerian graphs: a survey, *J. Graph Theory* **16** (1992), 177–196.
- [27] M. Coornaert, T. Delzant & A. Papadopoulos, Geometrie et theorie des groupes. Les groupes hyperboliques de Gromov. *Springer Lecture Notes in Mathematics* **1441**, Springer 1990.
- [28] Q. Cui, J. Wang and X. Yu, Hamilton circles in infinite planar graphs, *J. Comb. Theory B* (to appear).
- [29] R. Diestel, The countable Erdős-Menger conjecture with ends, *J. Comb. Theory B* **87** (2003), 145–161.
- [30] R. Diestel, The cycle space of an infinite graph, *Comb. Probab. Comput.* **14** (2005), 59–79.
- [31] R. Diestel, End spaces and spanning trees, *J. Comb. Theory B* **96** (2006), 846–854.
- [32] R. Diestel, *Graph theory*, 4th edition, Springer-Verlag 2010 and <http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html>
- [33] R. Diestel & D. Kühn, A universal planar graph under the minor relation, *J. Graph Theory* **32** (1999), 191–206.
- [34] R. Diestel & D. Kühn, On infinite cycles I, *Combinatorica* **24** (2004), 69–89.
- [35] R. Diestel & D. Kühn, On infinite cycles II, *Combinatorica* **24** (2004), 91–116.
- [36] R. Diestel & D. Kühn, Topological paths, cycles and spanning trees in infinite graphs, *Eur. J. Comb.* **25** (2004), 835–862.
- [37] R. Diestel & I.B. Leader, Normal spanning trees, Aronszajn trees and excluded minors, *J. Lond. Math. Soc.* **63** (2001), 16–32.
- [38] R. Diestel & P. Sprüssel, The homology of a locally finite graph with ends, preprint 2007
- [39] R. Diestel & P. Sprüssel, The fundamental group of a locally finite graph with ends, preprint 2008
- [40] R. Diestel & P. Sprüssel, On the homology of non-compact spaces with ends, preprint 2009
- [41] S. Eilenberg and N. Steenrod, *Foundations of Algebraic Topology*, Princeton University Press 1952.
- [42] A. Georgakopoulos, Connected but not path-connected subspaces of infinite graphs, *Combinatorica* **27** (2007), 683–698.
- [43] A. Georgakopoulos, Fleischner’s theorem for infinite graphs, *Oberwolfach reports* **4** (2007).
- [44] A. Georgakopoulos, Topological circles and Euler tours in locally finite graphs, *Electronic J. Comb.* **16** (2009), #R40.
- [45] A. Georgakopoulos, Infinite Hamilton cycles in squares of locally finite graphs, *Adv. Math.* **220** (2009), 670–705.
- [46] A. Georgakopoulos, Graph topologies induced by edge lengths, preprint arXiv:0903.1744 (2009).
- [47] A. Georgakopoulos, Uniqueness of electrical currents in a network of finite total resistance, preprint arXiv:0906.4080 (2009).

- [48] A. Georgakopoulos & Ph. Sprüssel, Geodesic topological cycles in locally finite graphs, preprint 2006.
- [49] M. Gromov, Hyperbolic Groups, in: *Essays in group theory* (S.M. Gersten, ed), MSRI series vol. 8, pp. 75–263, Springer, New York, 1987.
- [50] G. Hahn, F. Laviolette and J. Siran, Edge-Ends in countable graphs, *J. Comb. Theory B* **70** (1997), 225–244.
- [51] R. Halin, Über die Maximalzahl fremder unendlicher Wege, *Math. Nachr.* **30** (1965), 63–85.
- [52] R. Halin, A theorem on n -connected graphs, *J. Comb. Theory B* **7** (1969), 150–154.
- [53] R. Halin, Unendliche minimale n -fach zusammenhängende Graphen, *Abh. Math. Sem. Univ. Hamburg* **36** (1971), 75–88.
- [54] R. Halin, Minimization problems for infinite n -connected graphs, *Comb. Probab. Comput.* **2** (1993), 417–436.
- [55] R. Halin, Miscellaneous problems on infinite graphs, *J. Graph Theory* **35** (2000), 128–151.
- [56] A. Hatcher, *Algebraic Topology*, Cambridge University Press 2002.
- [57] D.A. Higgs, Matroids and duality, *Colloq. Math.* **20** (1969), 215–220.
- [58] D.A. Higgs, Equicardinality of bases in B -matroids, *Can. Math. Bull.* **12** (1969), 861–862.
- [59] D.A. Higgs, Infinite graphs and matroids, *Recent Prog. Comb., Proc. 3rd Waterloo Conf.* (1969), 245–253.
- [60] D. Lick, Critically and minimally n -connected graphs, in (G. Chartrand & S.F. Kapoor, eds.): The many facets of graph theory, *Lecture Notes in Mathematics* **110** (Springer-Verlag 1969), 199–205.
- [61] W. Mader, Homomorphieeigenschaften und mittlere Kantendichte von Graphen, *Math. Ann.* **174** (1967), 265–268.
- [62] W. Mader, Minimale n -fach zusammenhängende Graphen, *Math. Ann.* **191** (1971), 21–28.
- [63] W. Mader, Ecken vom Grad n in minimalen n -fach zusammenhängenden Graphen, *Arch. Math.* **23** (1972), 219–224.
- [64] W. Mader, Über minimal n -fach zusammenhängende, unendliche Graphen und ein Extremalproblem, *Arch. Math.* **23** (1972), 553–560.
- [65] W. Mader, Eine Reduktionsmethode für den Kantenzusammenhang in Graphen, *Math. Nachr.* **93** (1979), 187–204.
- [66] W. Mader, Paths in graphs reducing the edge-connectivity only by two, *Graphs Comb.* **1** (1985), 81–89.
- [67] B. Nadler, *Continuum theory*, Dekker 1992.
- [68] C.St.J.A. Nash-Williams, Edge-disjoint spanning trees of finite graphs, *J. Lond. Math. Soc.* **36** (1961), 445–450.
- [69] C.St.J.A. Nash-Williams, Decompositions of finite graphs into forests, *J. Lond. Math. Soc.* **39** (1964), 12.
- [70] J.G. Oxley, Infinite matroids, *Proc. Lond. Math. Soc.* **37** (1978), 259–272.
- [71] J.G. Oxley, Infinite matroids, in: (N. White, ed) *Matroid Applications*, Cambridge Univ. Press 1992.
- [72] B. Richter, Graph-like spaces: an introduction, in this volume.
- [73] B. Richter and A. Vella, Cycle spaces in topological spaces, preprint 2006.
- [74] M. Schulz, Der Zyklusraum nicht lokal-endlicher Graphen, *Diplomarbeit*, Univ. Hamburg 2005.
- [75] P. Sprüssel, End spaces of graphs are normal, *J. Comb. Theory B* **98** (2008), 798–804.

- [76] M. Stein, Arboricity and tree-packing in locally finite graphs, *J. Comb. Theory B* **96** (2006), p.302–312.
- [77] M. Stein, Forcing highly connected subgraphs in locally finite graphs, *J. Graph Theory* **54** (2007), 331–349.
- [78] M. Stein, Extremal infinite graph theory, preprint 2009.
- [79] C. Thomassen, Hamiltonian paths in squares of infinite locally finite blocks, *Ann. Discrete Math.* **3** (1978), 269–277.
- [80] C. Thomassen, Planarity and duality of finite and infinite graphs, *J. Comb. Theory B* **29** (1980), 244–271.
- [81] C. Thomassen, Nonseparating cycles in k -connected graphs, *J. Graph Theory* **5** (1981), 351–354.
- [82] C. Thomassen, Reflections on graph theory, *J. Graph Theory* **10** (1986), 309–324.
- [83] C. Thomassen and A. Vella, Graph-like continua and Menger’s theorem, *Combinatorica* **28** (2009).
- [84] W.T. Tutte, On the problem of decomposing a graph into n connected factors, *J. Lond. Math. Soc.* **36** (1961), 221–230.
- [85] A. Vella, A fundamentally topological perspective on graph theory, PhD thesis, Waterloo 2004.
- [86] K. Wagner, *Graphentheorie*, BI-Hochschultaschenbücher, Bibliographisches Institut, Mannheim 1970.
- [87] G.T. Whyburn, On n -arc connectedness, *Trans. Amer. Math. Soc.* **63** (1948), 452–456.
- [88] W. Woess, Dirichlet problem at infinity for harmonic functions on graphs, in (J. Kral et al, eds): International conference on potential theory (1994), Proceedings, de Gruyter (1996), 189–217.
- [89] W. Woess, *Random walks on infinite graphs and groups*, Cambridge University Press 2000.
- [90] S. Zhan, On hamiltonian line graphs and connectivity, *Discrete Math.* **89** (1991), 89–95.

Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 D - 20146 Hamburg
 Germany

12.7.2009